UNIVERSITY OF CALIFORNIA
SANTA CRUZ

RESPONSIVENESS IN NARRATIVE SYSTEMS

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

COMPUTER SCIENCE

by

Stacey Mason

September 2021

The Dissertation of Stacey Mason
is approved:

Michael Mateas, Chair

Noah Wardrip-Fruin

Emily Short

Mark Bernstein

Peter F. Biehl
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures vii

List of Tables ix

Abstract x

Acknowledgments xii

1 Introduction 1

1.1 Building an Approach to Understanding Responsiveness 5
1.2 Why Do We Evaluate? 8
1.3 Art Practice as Research Practice 11
1.4 Theoretical Contributions 13
1.5 Design & Technical Contributions 15
1.6 Overview of This Dissertation 17

2 Frameworks and Assumptions 20

2.1 Lenses, Not Boxes 21
2.2 A Broader Understanding of Game Narrative 22
2.3 Games and Narrative Are Not at Odds 25
2.4 Games Offer a Variety of Pleasure-Experiences 27
2.5 Intended Pleasure-Experience Matters to Analysis 30

3 Models of Player Agency 34

3.1 Agency is a Central Pleasure of Interactive Media 36
3.2 Disambiguating Agency and Interactivity 38
3.3 Understanding the Experience of Agency 45
3.3.1 Agency Requires Intention 46
3.3.2 Agency is Bounded by System Affordances 48
3.3.3 Agency Requires Buy-In and Trust 53
3.3.4 Agency is Felt Through Feedback 59
3.3.5 Agency is Dependent on Feedback, Not Consequences 62

3.3.6 Agency Exists at Different Levels of Abstraction Simultaneously 66
3.4 Reconciling Varieties of Agency 69
3.4.1 Categorizations of Narrative Agency: Story and Telling 70
3.4.2 Bestiary of Player Agency 74
3.5 Situating Agency Within a Model of Interactivity 77
3.5.1 Player LTS 77

4 Models of Affordances 81
4.1 Gibson's Affordances 83
4.2 Norman's Affordances 87
4.3 Reconciling Gibson and Norman 90
4.4 Formal and Material Affordances 92
4.5 Affordances and the Mental Model 96
4.6 Perceived Affordances, Real Affordances, and Feedback 101
4.7 Agency as a Loop 104

5 Responsive Systems 108
5.1 Introducing Responsiveness 108
5.2 A Model of Changing Affordances 112
5.2.1 Changing Feedback 115
5.2.2 Changing Formal Affordances 116
5.2.3 Changing Material Affordances 120
5.2.4 Changing Perceived Affordances 123
5.3 Nested Affordances and Levels of Abstraction 128
5.4 Responsiveness and Emergence 131
5.5 Separating Responsiveness and Agency 132
5.6 Responsiveness as a Lens 134

6 Aesthetics of Responsive Systems 138
6.1 Tightly-Coupled Feedback 139
6.1.1 "Clementine will remember that." 140
6.1.2 "The Kid just raged for a while." 144
6.1.3 "Stanley stepped into the broom closet." 146
6.1.4 Feedback and Counterfactuality 148
6.2 Responsiveness Through Affordance Changes 151
6.2.1 Directly-Authored Changes to Affordances 151
6.2.2 Combinatorial Affordances 153
6.2.3 Contextual Affordance Changes 155
6.2.4 Community-Driven Affordance Changes 157
6.2.5 Improvisation as Changes to Formal Affordances 159
6.2.6 Toward the Holodeck 161
6.2.7 The Chicken or the Egg: Formal or Material Affordances First? . 163
6.3 Responsiveness Through Signalled Counterfactuality 164

6.4 Changing affordances at a high level of abstraction 169
6.5 A Matter of Taste 170

7 Interventions Toward System Responsiveness 172
7.1 How Systems Respond: The System-Side LTS Loop 174
7.1.1 How Systems Listen 175
7.1.2 How Systems Think 177
7.1.3 How Systems Speak 179
7.2 Technical Interventions in the LTS Loop 181
7.2.1 System Listening and Responsiveness 182
7.2.2 Toward Listening Improvements 184
7.2.3 Considerations When Designing Systems That Listen 185
7.2.4 System Thinking and Responsiveness 187
7.2.5 Toward Thinking Improvements 188
7.2.6 Considerations When Designing Narrative Systems That Think . 194
7.2.7 Keep NPC Models Focused 195
7.2.8 System Speaking and Responsiveness 196
7.2.9 Toward Speaking Improvements 197
7.2.10 Considerations When Designing Systems That Speak 198

8 The Lume System 199
8.1 Motivations 199
8.1.1 Initial Motivations 199
8.1.2 Related Work 200
8.1.3 Motivations Toward Responsiveness 203
8.2 System Overview 205
8.2.1 Narrative Model 205
8.2.2 Scene Node Trees 207
8.3 Procedurality Within Scenes 216
8.3.1 Definite Clause Grammars 216
8.3.2 Pronoun Replacement and Point of View 217
8.4 Recalls and the event list 219
8.5 Characters and Relationships 220
8.6 Balancing Narrative Dynamism and Coherence 222
8.6.1 Toward Narrative Dynamism 222
8.6.2 Toward Narrative Coherence 223
8.7 Technical Interventions Toward Responsiveness 226
8.7.1 Think: New Approach to Story Modeling 226
8.7.2 Speak: New Approach to Feedback 228
8.8 Authoring with Lume 229

9 Case Study of Rumina Woods **231**
9.1 Overview of Rumina Woods . . . 232
9.2 Development: Integrating Lume & Unity . . . 233
9.3 Narrative Structure of Rumina Woods . . . 235
9.3.1 Mirror World . . . 236
9.3.2 Fires in the Desert . . . 236
9.4 Example Scene Authoring . . . 238
9.4.1 Ensuring Proper Selection . . . 240
9.4.2 Signalling Consequence and Counterfactuality . . . 241
9.4.3 How Different Characters React . . . 244
9.4.4 Foreshadowing and Seeding Future Drama . . . 246
9.4.5 Adding Variety and Voice . . . 249
9.5 Authorial Affordances of the Lume System . . . 253
9.5.1 What Lume Makes Easy . . . 254
9.5.2 What Lume Makes Difficult . . . 258
9.6 Design Insights . . . 262
9.6.1 Stationary vs Location-Based: Trade-offs in Scene Bucketing . . . 262
9.6.2 Size of content unit . . . 263
9.6.3 The Uncanny Valley of Representative Fidelity . . . 264
9.6.4 Work With The Generator Rather Than Against It . . . 264
9.6.5 Narrative and the Oatmeal Problem . . . 265
9.6.6 Put Oatmeal to Work . . . 267
9.6.7 Apophenia is Your Friend . . . 267
9.7 Looking Ahead . . . 268

10 Conclusions **269**
10.1 Contributions . . . 270
10.2 Open Problems in Responsive Narrative Systems . . . 271
10.3 Future Work . . . 276

List of Figures

2.1 Image from Bartle's taxonomy of player types [13] 29

3.1 Depiction of Crawford's Listen-Think-Speak (LTS) model of interactivity [35] 42
3.2 Reformulation of Murray's agency as enacting one side of the LTS model from the player's point of view 44
3.3 Revised model of agency as half of the LTS interactivity loop. Agency is bounded by system affordances, which color the player's expectations and understandings. Her actions are mediated through the material affordances availabe to her and she understands her agency is confirmed by the feedback she receives. 79

4.1 Box art for *Super Mario Bros.*. The box art and other marketing materials offer formal affordances to start the process of building the player's mental model of the game before she even begins playing. 95

6.1 Screenshot from *Long Live the Queen*. The player has just made a choice and now a string of skill checks suggest that the choice had counterfactual outcomes the player is not seeing. 167
6.2 Screenshot of level choice map from *Detroit: Become Human*. Choice maps at the end of each level make narrative counterfactuality very explicit. 168

7.1 Responsiveness as a component phenomenon in the LTS interactivity loop on the system's side. 174
7.2 Diagram of responsiveness that depicts various narrative models at the think stage. 180

8.1 Example of a scene node-tree 208
8.2 Eligibility in Lume scene node trees. This is a valid scene, because at least one continuation or at least two choice nodes are available at each level, thus a viable scene can be presented. 214

8.3 Example narrative emerging from individually-selected scenes. The combination of preconditions, bindings, and recall phrases help ensure coherence. . . . 224

9.1 Screenshot from *Rumina Woods* prototype . . . 233
9.2 Screenshot of Lume's Scene Builder tool in Unity environment. . . . 234
9.3 Diagram of the narrative structure of *Rumina Woods*. . . . 239
9.4 Screenshot of the Lume Scene Editor depicting a single node and the Node editor. . . . 242
9.5 Screenshot of the Lume Scene Editor. The example scene is in progress (zoomed in). . . . 244
9.6 Screenshot from the Lume Scene Editor showing the final structure of our example scene (zoomed out). . . . 246

List of Tables

3.1 The distinction between choice, feedback, and consequence in a narrative choice scenario. 65
3.2 Explanation of the player's experience of agency at multiple levels of abstraction simultaneously while playing *Street Fighter V*. 67
3.3 Classifications of agency at different levels of abstraction. 69
3.4 Marie-Laure Ryan's categorizations of types of interactivity. 72

5.1 Explanation of system responsiveness at multiple levels of abstraction simultaneously while playing *Street Fighter V*. 129

Abstract

Responsiveness in Narrative Systems

by

Stacey Mason

While game studies and interactive narrative communities have developed a rich body of research around agency in our interactions with narrative systems, comparatively little research effort has been devoted to studying the system's role in fostering, supporting, and reinforcing that agency.

This dissertation offers a combined theory-design-technical exploration of a system's ability to respond to player agency. Building upon Chris Crawford's model of interactivity as a loop between two participants who each Listen, Think, and Speak in turn, this work considers the interactivity loop as a dynamic between a player and a system. It formulates agency as the experience of this loop from the perspective of the player, and incorporates current theories of agency into this model. It explores how affordances and *feedback* act as the means of communication between the player and system, and defines a system's responsiveness as the degree to which a system changes its affordances and feedback as a result of player actions. The theoretical lens of responsiveness is then applied to a range of games as a design analysis tool.

The technical contributions of this dissertation include a technical design analysis of the inner workings of Lume—an AI storylet system designed to offer highly responsive narrative experiences—with an eye toward examining how its technical design and component systems foster responsiveness. Finally, the dissertation offers a case study of a prototype game created with the Lume system. The case study discusses the authorial affordances of the Lume system, a range of procedural narrative design techniques, and how the components of responsiveness outlined in this dissertation can be leveraged to create a compelling narrative experience.

Acknowledgments

No research occurs in a vacuum, and I have been particularly fortunate to be surrounded by an incredible community of amazing people without whom this dissertation could not have happened. The work presented in this document grew out of numerous conversations with many insightful people around me. I am immeasurably grateful to all of them.

I owe an immeasurable debt to the reading committee of this dissertation, who have generously provided invaluable feedback, suggestions, conversations, and encouragement. My advisors Michael Mateas and Noah Wardrip-Fruin have supported me in every possible way throughout years of mentorship, first as a fledgling designer in an MFA program, where they helped me find my voice as a creator, and now over the course of this PhD. They have taught me so many things, but most importantly, they have helped me learn how to navigate cross-disciplinary spaces of innovation and break down boundaries—between industry and academia, between humanities and computer science, between research and creation, and between convention and innovation. They have helped me grow as a researcher, a creator, and a person more than they will ever know, and their influence will be felt in my life and work hereafter. I am extraordinarily grateful to Emily Short for her incredible insights, feedback, mentorship, and friendship over the years. Emily was one of my first role models for what a career in interactive fiction might look like and her polymathy has been a huge inspiration for the theory-design-technical approach to research that this dissertation adopts. And finally, I owe

more than I can put into words to the mentorship of Mark Bernstein, who took a chance on a tenacious southern girl determined to leave a mark on the future of storytelling. I owe every shred of success I have had since then to the opportunities Mark gave me at Eastgate. Over the years, he has pushed me to be a better writer, researcher, speaker, and thinker, and his unwavering standards have always arisen from a steadfast belief in potential I did not recognize within myself. The fact that this particular combination of excellent mentors gave their support to this work has improved it to an incalculable degree, and I can say without any hyperbole that they have all changed my life.

I was fortunate to have several academic mentors with whom I studied on my way to this degree program, including Soraya Murray, Terry Harpold, Greg Ulmer, and Brandon Kershner. Each of them went above and beyond in guiding my studies, offering me encouragement and opportunities, and shaping the path I have ultimately pursued. Additionally, a superlative academic research community welcomed me well before I even began my graduate studies, and in particular the encouragement and kindness of Mark Anderson, Bill Bly, Les Carr, Diane Greco, Dene Grigar, Dame Wendy Hall, Charlie Hargood, Michael Joyce, George Landow, Cathy Marshall, Nathan Matias, David Millard, Stuart Mouthroup, Nick Montfort, Rikki Prince, and Su White were especially influential in helping me find my place within the growing landscape of research on interactive narrative as it relates to narrative and critical theory, literature and media studies, agency and choice, links and hypertext theories, interaction design, and play.

For the last 5 years, I have simultaneously worked full-time on research and development teams within the commercial videogame industry while working toward

my PhD, and my experiences there have shaped my research in countless ways. The commercial game space is notoriously difficult to break into, especially for someone still actively pursuing an education. I owe many of my opportunities in industry to the IGDA Foundation, their IGDA Scholars program, and the excellent mentorship I received from Like Dicken as a part of that program. During my time in industry, I have particularly appreciated the people who have taken special effort to teach me, review my work, and offer their wisdom and advice. Among them, Robin Hunicke, Mitu Khandaker, Robin Yang, Leah Hoyer, Dan Connors, Ryan Kaufman, Stephen McManus, Molly Maloney, Meredith Ainsworth, Shanon Ingles, and Nicole Martinez have each taken particular care to help me grow as a narrative designer and industry professional. Conversations with all of them have influenced this dissertation in some way.

We learn more from our close collaborators than from anyone, and I am particularly fortunate to work with some of the brightest minds and visionaries for narrative technology in the games industry. None of this work would be possible without the insights and talents of my collaborator Ceri Stagg, a brilliant AI researcher who has been an ideal working partner throughout Lume's development. Likewise, our collaborator Carl Muckenhoupt's talent for transforming difficult technical problems into beautifully designed tools and solutions has taken Lume from a promising but unwieldy backend prototype to a robust authoring environment in a major development engine. And all of our efforts would be invisible without Nikola Odic's contributions to art direction and visual communications that have allowed us the ability to present affordances to players and solve the visual problems inherent to procedural narrative. I could not ask for a better development team.

In addition to my current collaborators, I have been extremely grateful to learn from a variety of collaborators across various research and innovation efforts during my time in industry. Those who have impacted me most include Brett Douville, Lyndsey Gallant, Erica Harrell, James Lantz, Gabe McGill, Megan Muir, Desiree Proctor, Mallory Schleif, and Julia Palatovska. Each of them has given me some insight or learning that has impacted my thinking on responsiveness in narrative design.

All of my work has been influenced by the incredible intellectual community offered by the Expressive Intelligence Studio and the Center for Games and Playable Media at UC Santa Cruz. My thinking on agency, narrative, AI, and game design has been deeply touched by that community over the years, particularly by colleagues Devi Acharya, Duncan Bowsman, Sherol Chen, Melanie Dickinson, Cyril Focht, Teale Fristoe, Katie Green, John Grey, April Grow, Rehaf Jammaz, Shi Johnson-Bey, Nick Junius, Eric Kaltman, Max Kreminski, Sarah Fay Krom, Larry LeBron, Dylan Lederle-Ensign, Heather Logas, John Mawhorter, Peter Mawhorter, Stella Mazeika, Josh McCoy, Wes Modes, John Murray, Joe Osborn, Adrian Phillips, Allie Riggs, James Ryan, Ben Samuel, Ben Spalding, Adam Summerville, Mike Treanor, and Marcelo Viana Neto. In addition to these incredible researchers, I want to extend a special thanks to Kate Compton, Jacob Garbe, Jonathan Pagnutti, Aaron Reed, and Dietrich Squinkifer, each of whom have deeply influenced me as incredible colleagues and even better friends and personal confidantes.

Outside of UC Santa Cruz, I have also received an incredible amount of intellectual and emotional support over the course of this research from my dear friends Kristal Dougherty, Rachel Chai, Erika Dempsey, Laura Libert, Earl McKee, and Patrick McSweeney. I am additionally grateful for the help and support of Lance Simons, who is an incredible friend and human, and who spent many hours tutoring me in advanced math and algorithms so that I, an artist loyal to her humanities roots, could grow into a card-carrying computer scientist.

And of course, none of this would have been possible without the love and support from my family, who have always encouraged my curiosity and insatiable desire to find knowledge. I offer endless love and gratitude to my partner, collaborator, and mentor Zacariah Litton, who has been my biggest fan and a source of strength throughout this process. I would not be the person I am today if not for your belief in me. And finally, my boundless love and gratitude to my husband, Michael Brich, who has been my partner, my rock, my confidante, and my cheerleader; I never could have done this without you.

The text of this dissertation includes reprint from the following previously published material: Mason, Stacey & Stagg, Ceri & Wardrip-Fruin, Noah. (2019). Lume: a system for procedural story generation. *FDG '19: Proceedings of the 14th International Conference on the Foundations of Digital Games.* 1-9. 10.1145/3337722.3337759. Noah Wardip-Fruin, one of the co-authors listed in this publication, served as an advisor who supervised and directed the research for this dissertation; the other co-author, Ceri Stagg, has granted express written permission to include this material in this dissertation. The Lume System presented in Chapter 8 was developed collaboratively; Ceri and

I both made substantial, deeply-intertwined contributions to the technical design and implementation of the system.

Part of this research was made possible through support from the Eugene Cota-Robles Fellowship program.

Chapter 1

Introduction

In defining the essence of what we do and why it is important, we have named the field "interactive narrative", indicating that interactivity is central to the aesthetic goals of our medium.

Brenda Laurel notes that even in the early days of computing, scientists identified the interaction between computer and user as a key potential for the new technology. She opens her book *Computers as Theatre* [87] with an anecdote about the invention of the canonical first computer game, *Spacewar!* [145], which the scientists who created it saw as an obvious use for the machine. But why, she asks, was a game the obvious thing to do with this powerful new technology?

> Why was *Spacewar!* the "natural" thing to build with this new technology? Why not build a pie chart, an automated kaleidoscope, or a desktop? Its designers identified *action* as the key ingredient, and conceived *Spacewar!* as a game that could provide a good balance between thinking and doing for its players. They regarded the computer as a machine naturally suited to representing things that you could see, control, and play with. Its interesting potential lay not simply in its ability to perform calculations, but in its capacity to co-create and represent actions with human participants.

(original emphasis) [87]

For Laurel, one of the magical properties of computers from their outset was their ability to work *with* human users in a phenomenon we now know as interactivity. The act of human and computer exchanging actions back and forth created the sense that unlike other tools, computers have the capacity to understand us and respond intelligently.

How interactive media foster the pleasures of interactivity became a natural area of inquisition for researchers and theorists. Janet Murray's *Hamlet on the Holodeck*, one of the early explorations of the aesthetic properties of the interactive narrative, opened several areas of inquiry that remain relevant today [118]. Among other aesthetic pleasures like *immersion*, and *transformation*, Murray identifies *agency* as a key component of interactive narratives, which she defines as "the satisfying power to take meaningful action and see the results of our decisions and choices." In the 2016 updated edition of the book, she further emphasizes the importance of agency, arguing that "the experience of agency by the interactor is the key design value for all digital artifacts... Dramatic agency should be the goal of design for interactive narrative in any form." [119]

Agency has become a robust area of research within interactive media, and games in particular[99][193][179][95][80][48]. We have developed a rich vocabulary for exploring an audience's desires, intentions, and power to enact its will. This work tends to frame agency in terms of the user's actions, decisions, and choices, which places the focus on her as the subject, and much of the work in this space explores the nuances and bounds of this framing. However, the promise of interactivity as the computer's ability to "co-create and represent actions with human participants" [87] also implies another participant: the computer. While current theories of agency represent an attempt to model interactions with a computer from the user's perspective, comparatively few attempts have been made to model how the system's actions factor into this conversation. We do not have a clear picture of the system's role in fostering, supporting, and reinforcing the pleasure of agency, though such a role must surely exist since many of these experiences are ultimately *designed*. After all, the designer does not control the player; she controls the system with which the player will interact. Understanding the system's role in the player's agency is critical to designing for it. When a player exerts agency, what can the system do to show it has understood and is working with her? How do systems *respond*?

Game industry practitioners, who have been successfully designing for agency for decades, are a natural place to turn for a starting point for such knowledge. Unfortunately, academia is historically miserable at recognizing the knowledge of practitioners if that knowledge is not documented in "the academic literature." Some of the knowledge that is elementary to industry narrative designers has not been incorporated into the research canon, nor for that matter published in papers or books. Consequently, much of it is rediscovered or relearned years later by the academic community and presented as new. Meanwhile, some of the most groundbreaking work lives across blogs, Youtube videos, and Twitter threads, and these informal sources represent state-of-the-art narrative design techniques.

Yet even among the industry state-of-the art, knowledge is often based on intuition, trial-and-error, and word-of-mouth techniques. Terms like *agency* are used among industry practitioners in conflicting ways, many of which are arguably-correct and equally-valid despite their seeming contradictions. The industry knowledge around designing a system to foster agency, while often serving as a prerequisite for what is understood as "good design", lacks a unifying formalism at which narrative designers can point—something akin to the popular MDA framework [73]. This design framework, published in academic literature and adopted widely by industry design practitioners, offers a blueprint for the process of designing games with a particular aesthetic experience at their core, while also providing academics with language to analyze the Mechanics, Dynamics, and Aesthetics of such experiences. The MDA framework is inspired in its simplicity and insightfulness; though an analogous framework for designing for agency may be more elaborate or complex, its potential for usefulness is indisputable.

Meanwhile, the technical capabilities of our interactive systems have become more and more advanced since Murray and Laurel's theoretical conceptions. Computer scientists have approached the problem of how to make interactive stories feel more responsive by providing technical solutions that improve individual components of these narrative experiences: character AI and agent believability, story structure, input recognition, and so forth. Much of the current development into interactive narrative generation presumes that increased technical capabilities will automatically lead to more pleasurable narrative experiences. The implicit argument seems to be that if we only had better NPC AI, better drama managers, better social simulation, better authoring tools, and so forth our narratives would be so much richer and more responsive. Yet many of these advances never become integrated into released products. There are many reasons for why this might be that fall outside of the scope of this inquiry, including systemic incentives of academia that favor novelty over usability, lack of incentive to release code, negative incentive to maintain active codebases or documentation, and the use of a tally of publications (read: peer-reviewed papers) as a metric for success within computer science departments. But within the scope of our inquiry, I also believe that much of the technical research in the space of interactive narrative focuses on very local problems without much thought toward how these solutions might be integrated into larger experiences. My experience as a creator has taught me that this approach, while useful in informing other research, is unlikely to be integrated into successful creative experiences if it does not understand how it needs to fit into them. Our technical researchers need a better understanding of why games feel responsive so they can figure out how to situate their goals within that framework. If we don't know the goal, how can we possibly know how to get there?

1.1 Building an Approach to Understanding Responsiveness

As illustrated above, there are at least three ways we might approach the problem of understanding an interactive narrative system's response to player agency, and each comes with its own traditions, strengths, and drawbacks:

1. **The humanistic approach** analyzes the problem with theory. This approach leans on the learnings of past theoreticians, analyzes available media artifacts through its theoretical lenses, and provides theoretical models that account for trends and patterns observed in those media. Humanists might also write speculatively about technical solutions or areas that would reward the creation of new techniques or artifacts, but rarely create this technical solution themselves.

2. **The design approach** tries to offer the best solution to a given problem while working within the constraints at hand, such as the current tools available, the current realities of budget, time, or team composition, etc.. This approach offers invaluable insights for how to achieve certain aesthetic effects and solutions to problems. Design approaches might also offer speculative solutions, such as to say "I should be able to do X, but can't; someone should invent a technical solution." The design approach is the one we see most among industry talks and creative practitioners, and these conversations often happen independently from academic research efforts.

   Among academics, foregrounding a design contribution is difficult, especially in computer science (CS) conferences and publication venues. The contributions that do find publication often take the form of human-computer interaction (HCI) work, but these venues have their own evaluative norms that make certain discoveries difficult to publish. Groundbreaking design discoveries are hard to quantify, and thus difficult to evaluate in most CS contexts, but we all recognize retrospectively that something novel has happened once we see a representative creative product, even if we tend to lean on the other approaches to articulate exactly what is novel or impactful about the discovery.

3. **The technical approach** proposes innovations in tools and the technical capabilities of systems. When applied well, this approach provides solutions that change the constraints that designers and creators must work within. New technologies open up the possibility space for creative output and give designers the ability to create experiences that have never existed before. Unfortunately, in practice, much technical research is undertaken without enough input from the other two disciplines, so new solutions become a hammer with which all problems are solved. Often an empirical study is run to evaluate the usefulness of the hammer, but without working closely and collaboratively with an actual creator who is using the tools and solutions, how much the solutions actually advance the field remains unproven.

The question then becomes, "how can we evaluate this work?" If I am effectively creating the metrics by which I will judge myself, how can you, my dear reader, be sure that I'm not begging the question? To answer this, I hope you will permit a brief aside about evaluation.

1.2 Why Do We Evaluate?

The difficulties of evaluation in the research of games and other computational systems is well-documented [192]. Gatekeepers of computer science, fearful that "the quality of CS research might be inferior to other disciplines, in particular the natural sciences, the engineering sciences, and applied mathematics" argue that research that lacks quantifiable experiments to validate claims represents a "weakness that should be rectified for the long-term health of the field" [184], but this approach is myopic and fails to account for contributions to knowledge that are not best-evaluated by an experiment. When some of the most foundational papers in computer science would not pass muster in this academic climate, we have to ask ourselves whether this approach is the correct one. The insecurities of being seen as "lesser" than other sciences and the clamoring for quantitative analysis to validate the position of the field as legitimate has been taken to an extreme in many computer science publication venues, where it incentivizes the publication of research that can be measured to the exclusion of research that represents invention.

This leads us to take a step back and ask, why do computer scientists want experiments in the first place? What is the purpose of evaluation?

The purpose of research is ultimately the creation of knowledge through discovery. If knowledge is added to the corpus of human understanding, research has succeeded. We ask a question and search for an answer, and when we find an answer, we need a way to determine whether our answer is a good one. Ultimately, all answers are making an argument, and the nature of that argument will determine the best way to evaluate its validity. For scientific discoveries, the argument that altering this particular variable caused that particular result lends itself to an experiment. The argument that a new software design approach yields a different kind of product, however, may benefit more from a qualitative analysis, an ethnographic study, a post-mortem analysis, and so on. The argument that creative works can be analyzed differently if we think of them through a new theoretical lens benefits from an example of the lens being applied to a subject to see what new insights are gained. The answer to the question, "How do we know this is a good solution?" will vary based on the claim at hand.

Many of the greatest advancements in interactive narrative research to date have been made by people working on creative artifacts, theory, and technical solutions simultaneously. Judy Malloy, one of the earliest and most impactful figures in hypertext literature [11], created foundational works of hypertext fiction including *Uncle Roger* [91] and *its name was Penelope* [90], while also writing prolifically about the theory and practice of interactive literature and internet art [92] [89]; and while spearheading a variety of early critical initiatives, including serving as the founding editor of *Leonardo Electronic News*, now *Leonardo Electronic Almanac*, a premiere journal of interactive art. Michael Joyce, often cited as the "granddaddy of hypertext fiction" [34], not only created one of the most beautiful and prolifically-studied works of hypertext fiction—*afternoon, a story*—and co-created Storyspace, the system used to create *afternoon* which then became widely used among prominent hypertext creators and educators of the 1990s [16], but he was also a prominent figure in the design, theory, and craft of hypertext authoring throughout the 1990s [77] [78]. Among current interactive fiction circles, Emily Short's simultaneous work on groundbreaking works of interactive fiction—including *Galatea* [155], *Alabaster* [162], and *Blood & Laurels* [156]—as well the technical systems that produce them [120], has informed her writings on design craft and theory, and her blog [160] is regarded as one of the most robust and insightful resources on interactive fiction today. Michael Mateas's work on *Façade* [106], for which he pioneered several technological advancements including ABL [103], drama management techniques, and natural language processing, also informed his theoretical model for agency in interactive narrative [99]; his contributions to design, technology, and theory that resulted from his simultaneous approach to all three [100] are still considered state-of-the-art among interactive narrative research communities almost 20 years later [81].

Clearly the approach to a combined theory-design-technical solution is a fruitful one if many of the major advances in interactive narrative—hypertext literature, Storyspace, Inform 7, *Façade*, to name only a few—resulted from such an approach. Yet due to the interconnected nature of this approach in which theory, design, and tech are all influencing each other in an interconnected and iterative way, evaluation is difficult. It almost certainly does not take the form of a quantitative study. The existence of a new genre of artifact is itself proof that an invention has occurred.

1.3 Art Practice as Research Practice

In his dissertation, Michael Mateas offers something of a manifesto toward the validity of a combined art and AI research practice. He writes,

> I call this intertwined combination of AI research and art practice *Expressive AI*. Expressive AI has two major, interrelated thrusts:
>
> 1. exploring the expressive possibilities of AI architectures - posing and answering AI research questions that wouldn't be raised unless doing AI research in the context of an art practice, and
> 2. pushing the boundaries of the conceivable and possible in art - creating artwork that would be impossible to conceive of or build unless making art in the context of an AI research practice.
>
> [...] In Expressive AI, technical research and artistic exploration are intertwined. Building an AI-based artwork such as *Façade* requires deep changes in both AI research and art practice; neither the researcher nor the artist can continue in a "business as usual" way. [100]

Similarly, I find that my own artistic practice within an Expressive AI context has allowed me to explore the theoretical assumptions around agency I carried when I entered the work. Over the course of creating a work with the goal of "maximizing agency," I soon realized that to "maximize agency" was itself a problematic approach, and what I was really trying to do was to *surface the impact that the player's agency had on the system*, which is a slightly different problem. In order to carry out this work, I needed to reformulate the framework in which I thought about agency. The narrative AI architecture toward which I was working was, in fact, a case of "posing and answering AI research questions that wouldn't be raised unless doing AI research in the context of an art practice" [100].

The major research questions that arose for me in this practice were thus:

1. How do the technical capabilities of a system create or support opportunities for agency? What makes a system feel responsive?

2. Does a responsive system differ in its design? If so, how?

3. What kinds of technical improvements can help in making a player feel heard by the system?

The framework that I present here arose from my work in building an AI system that could respond to players' interventions into a story in a way that felt more like a conversation. As I began to explore industry best-practices toward this goal, I realized that those conversations also lacked a framework for the thing I was trying to optimize. If the conversations around agency were difficult to coalesce around a single definition, the thing I was trying to optimize for was in an even sorrier discursive state. Occasionally it was called "reactivity" or "response", but specific definitions of those terms were difficult to come by, much less how to optimize for them or how to improve them.

Ultimately what I was looking for was a better understanding of how a system responds to a player's agency.

Although it seems like much of the AI research into agent believability and story structure also have an unspoken underlying goal of improving this property, the property itself remained elusive. I argue here that many of the research efforts in narrative AI are trying to improve upon the system's ability to recognize and respond to players' input, though we currently lack a framework to discuss this.

Many approaches to narrative design practice seem similarly in need of a better understanding of agency. I found that as I started the work on the intellectual foundations of the Lume system, new approaches to narrative design in the commercial games space were striking an excellent middle ground between the coherence of well-authored embedded narrative and the variety and personalization of emergent narrative. Games like *Sunless Sea* [55], *80 Days* [74], and *The Ice-Bound Concordance* [139] were demonstrating that there was a real appetite and fascination with storylet systems, systems in which narrative is constructed from small bits of content. Excitement around the idea of procedural narrative was starting to build [164], though few people in the industry seemed to know what procedural narrative actually meant, and almost nobody seemed to understand what the design considerations of procedural narrative might be. The design of these games was so different, many industry narrative designers I knew, who were still largely working in branches and state-machines, concluded that the new learnings offered by these new approaches to narrative design could not possibly be applied to work *they* were doing. Again, a framework that shows how these narrative techniques could apply to more traditional narrative design approaches might bridge that gap.

1.4 Theoretical Contributions

As Laurel and Murray demonstrated in the quotes at the beginning of this chapter, one of the truly magical allowances of our medium is the unique pleasure of feeling like a system has heard and recognized your input, adapted to your actions, and is playing *with* you. This property—that of a system intelligently responding to a player and adapting to them—is a property that I want to call *responsiveness.*

Many players have had an experience with an unexpectedly responsive system. It is the property that gives us feelings of "Wow, I didn't expect that to work!" or "Wow, I didn't realize how consequential my actions were!" Responsiveness is fostered by, supports, and reinforces our feelings of agency. Over this dissertation, I will also argue that it is a property of narrative systems that can be maximized and for which we can design.

Toward this aim, the primary theoretical contribution of this dissertation is a new model of responsiveness. In technical terms, responsiveness is the degree to which the system alters its affordances and feedback as a result of the player's actions; from an experiential perspective, it is the degree to which a system communicates to the player that it has heard and understood her. A good model of responsiveness should achieve two key goals:

1. The model should incorporate our current understanding of agency. It should be able to accommodate existing academic frameworks as well as support design "common knowledge" from industry best-practices.

2. Additionally, it should be able to explain why different games that feel responsive do. Ideally, it should be able to articulate how games that feel responsive in different ways achieve their distinct aesthetic experiences, as well as how games with different pleasure-experiences are able to arrive at responsiveness from different

mechanics.

The model I propose fulfills these criteria and articulates how we might improve narrative agency through technical interventions at various points in the model. It provides greater context into which we might fit existing research efforts. Further, it suggests new technical solutions we might explore.

Building upon Chris Crawford's model [35] of interactivity as a loop between two participants who each Listen, Think, and Speak in turn, I consider this loop as a dynamic between a player and a system. I then formulate agency as the experience of this loop (the LTS loop) from the perspective of the player, and incorporate current theories of agency into this model. I explore how *affordances* and *feedback* act as the means of communication between the player and system, and finally formulate a system's responsiveness as its ability to change its affordances and feedback in response to player actions. I then demonstrate how improvements in the system's ability to Listen, Think, and Speak are communicated through affordances and feedback, and how these different technical components affect the aesthetics of responsiveness.

1.5 Design & Technical Contributions

In recent years, the games industry has seen more and more games adopting approaches to narrative design that embrace procedural narrative, yet we still have relatively few design insights around these techniques and no unifying understanding that unites games across the spectrum of procedural narrative. As we move toward more procedural designs for narrative, the ability for narrative designers to understand the systemic elements of their craft is going to become ever more critical.

While I hope to offer practical design insights throughout this dissertation, a design *contribution* is achieved when design knowledge is discovered that previously did not exist. A successful design contribution should achieve two goals:

1. It should open up the possibility of creating a new experience.

2. Additionally, it should be actionable and work in concert with possible tools.

The approach provided here offers the possibility of a narrative that is uniquely responsive to player choice. In the tradition of Expressive AI's combined approach to art and AI research practice [101], many of the realizations that formulate my understanding of responsiveness grew out of the creation of Lume [98], an AI storylet system I built in collaboration with Ceri Stagg to offer highly responsive narrative experiences.

The technical contributions of this dissertation include a deep dive into the inner workings of Lume, with an eye toward examining how its technical design and component systems foster responsiveness. I discuss the authorial affordances provided by the system, and consider the tradeoffs of the narrative model we have chosen.

In addition to a discussion of Lume, I also offer a case study on a prototype game I have created with it. Mateas notes that "The issues of interactive story design are at least as important as the architecture." [100]. As the final technical contribution of this dissertation, I offer a case study of the technical design of a prototype game *Rumina Woods*, along with an exploration of some of the tradeoffs I have considered in its design. Inherent to this discussion is some reflection on the authorial affordances of the Lume system, the kinds of things it makes easy and difficult to achieve, the kinds of narrative scaffolding we have added to create a prototype that balances combinatorial possibility space with rules to enforce coherence, and finally how the components of responsiveness I have outlined can be used to create a compelling narrative experience.

1.6 Overview of This Dissertation

Chapter 2 outlines key assumptions from which the rest of the dissertation will be operating. I provide some context for interdisciplinary interpretations of the terms I use, and try to make explicit the assumptions from which my thinking operates.

Chapter 3 reviews existing models of agency. In this chapter, I look at the foundational model of interactivity from which we will examine responsiveness, with special care to disambiguate *interactivity* as a whole from the component pleasure of *agency*. In this endeavor, I will consider different meanings currently in use for the term and will try to reconcile them within academic models. I examine the model of agency as the balance of formal and material affordances, and offer that the systemic counterpart to agency is the system's ability to change those affordances.

Chapter 4 offers a deeper dive into affordances as the system's means of communication. In this chapter I examine the underlying assumptions inherent in our formulations of affordances—from their initial introduction to psychology literature by Gibson [67] through their introduction to the design and HCI communities by Norman

[125]. This chapter demonstrates how these assumptions manifest in notions of formal and material affordances carried into the LTS loop.

Chapter 5 builds upon this detailed understanding of affordances and feedback by proposing a definition of system responsiveness as the degree to which a system can change its affordances and feedback in response to player actions. I offer illustrative examples of different types of affordances and feedback changes alongside discussion of the aesthetic impacts of those changes. I also situate this definition within our working model of interactivity and demonstrate places where it overlaps with agency and places where it departs.

Chapter 6 examines the aesthetics of responsiveness in more detail. We look at places where technical design intervention creates responsiveness, situate current research efforts into some of these improvements, and suggest places where additional research efforts may yield increased responsiveness in a narrative context.

Chapter 7 interrogates how systems respond via the LTS loop. In looking forward toward how various technical design interventions in the LTS loop might be leveraged toward responsiveness, this chapter offers design considerations toward this end.

Chapter 8 details the technical specifications of Lume, an interactive narrative engine that I built with some very talented collaborators. Our goal with the design of the system was to facilitate some of the increases to narrative responsiveness that I laid out in the previous chapter. I discuss how the system's technical design tries to capture some of the design strategies discussed.

Chapter 9 explores a case study of a prototype game in progress being built using the Lume system. In this chapter, I discuss the authorial affordances provided by the Lume system and examine a prototype narrative game's structure through the lens of responsiveness introduced in Chapter 4.

Finally, Chapter 10 offers conclusions about how this research might be carried forward in the future.

Chapter 2

Frameworks and Assumptions

In this section, I make explicit some of the biases and assumptions that underlie many of my arguments to follow. While a critical reader might find these to be digressive or extraneous, I find myself growing increasingly passionate about understanding and exposing one's biases and ideologies in an effort to provide context for perspectives. For such a young field in which many of its most prominent figures span disciplines and ideological perspectives, such context can be critical to avoid misunderstandings or misread motivations.[1]

[1] On that point, please note that I use "she" throughout this document as the generic third-person pronoun. Most of my current work employs "they" as the gender-neutral pronoun, but in the authoring of this document, my repeated need to refer to a player in reference to a system made the use of "they" confusing, as my intended reference to the player as "they" read instead as if "they" referred to both the player and the system. Thus for readability I have returned to "she" as the default pronoun for this document. I hope the reader will forgive the departure from my usual conventions and will not read this as an erasure of other genders.

2.1 Lenses, Not Boxes

I wondered whether this section was necessary, but games discourse seems to repeatedly find itself adrift in questions of whether ____ counts as a ____. We have seen the question repeatedly: Do games count as art? [41] Do games without goals count as games? [79] Is *The Sims* a game? [190] Are hypertexts interactive fiction? [157] Are Twine pieces literature? [151]. In fairness to the researchers cited here, all of whom I respect immensely, many of these sources pose or imply the question rhetorically. Still, the necessity of such a framing frustrates me.

"Is ____ a ____?" is, in my opinion, a fruitless question. The question can never be answered to anyone's satisfaction, and it often serves as either a provocation to espouse politics or as a cover to institute gatekeeping. Its agenda is usually opaque to onlookers and bewildering to those trying to advance the field from across disciplines.

We should be thinking in lenses, not boxes.

It is always more interesting to ask "If we think of ____ through the lens of ____, what do we learn?" We know it is more interesting to ask "What do we learn by considering Tom Stoppard through the lens of poststructuralism?" than to ask "Is Stoppard's work poststructuralist?" And while some works may be more or less fruitful when examined through different lenses, sometimes the most unlikely pairings of lenses and objects are the most rewarding. Certainly any of the questions above would be much better served by reframing them in this way.

The most useful lenses will offer a new way to view a work and provide a path toward new insights through its application. The primary contributions of this dissertation include both the theoretical framework surrounding the introduction of responsiveness in computational systems as well as a technical system that spotlights the aesthetic components of responsiveness that we will discuss and the design insights gained therein.

My hope in introducing language around system responsiveness is not to create a box of responsiveness into which we can place some experiences and exclude others. Rather, I hope to create a lens that we can use to discuss how a game's aesthetic decisions around responsiveness contribute to its pleasure-experience: What kinds of affordances are changing and at what level of abstraction? Does the system adapt to the player's inputs? How are those changes signalled to the player? I hope to use this framework to offer a new way to discuss the aesthetics of playable media, in both its experience and its design.

In places, I also enumerate properties of prototypical examples of a term or class of creative works. My aim here is to establish a clear idea of illustrative examples of a concept, not to politically exclude any particular media or genre of works from interesting discussions by way of essentialist oversimplification.

2.2 A Broader Understanding of Game Narrative

Whole books have been written defining narrative as it is presented in games, not to mention the entire disciplines devoted to the study and critical theory of narrative more broadly. We have come a long way over the past couple of decades since researchers argued things like "stories are just uninteresting ornaments of gift-wrappings to games, and laying any emphasis on studying these kinds of marketing tools is just a waste of time and energy." (Eskelinen 2001 as quoted in [177]). Since then, we have developed a much more nuanced understanding of narrative and games and a language to support that nuance. Now, when critics argue that "Games are Better Without Stories" [19], attacking the dream of the Holodeck in the process, we can dismiss such claims with more specific arguments about how those critics are probably taking too narrow of a view of what "stories" entail. [140]

One distinction that is extremely important for our purposes is the distinction between *embedded narrative*, the narratives of the game world as told through visuals, cutscenes, dialogue, text and other in-game elements; and *emergent narrative*, the stories we can tell about a play session which emerge after-the-fact. My view of these two forms presumes that they are not distinct forms of narrative, but actually exist along a continuum. Some games offer a pleasure-experience in which the traversal of an embedded narrative is core to that experience. For example, in games like *Gone Home* [51], the exploration of space and the traversal of an embedded narrative are the player's sole focus, where other games like *Kingdom Hearts* [165] offer dual pleasure-experiences wherein a player character might be immersed in a narrative environment, but the moment-to-moment gameplay involves combat mechanics; in this series, cutscenes are used chiefly as rewards and pacing mechanisms to offer a moment of relief after a tense section such as a boss battle. On the opposite end of the spectrum, a game like *Chess*, while not focused on a narrative pleasure-experience, still offers the ability to tell a story of the game after the fact, thus containing the capacity for emergent narrative. And then a game like *Crusader Kings 3* [171] offers something in between, where the game's resources, locations, characters, and the traits that govern those characters' behavior are given theming that maps to the kinds of narrative components we expect in a story of political intrigue between rulers in Europe in the Middle Ages. Thus, most of the interesting stories players tell will involve the actions they took and the counteractions their enemies took against them as we would expect from the embedded narrative of a *Chess* match, but the narrative elements from which those stories arose were dipped in fiction; they provide specific roleplay hooks, offer the kinds of character traits we might expect from a particular narrative genre, and facilitate the kinds of character arcs we expect from a fictional narrative, and the emergent stories will reflect those elements. Even games like *Super Mario Bros.* [137], whose pleasure-experience is more focused on action-oriented gameplay and players' mastery of moving a character through a challenging space, use embedded narrative elements to motivate the player into certain actions to facilitate the intended mechanical experience. We are not just moving a box through space; we are moving a character through a castle, jumping on turtles, defeating a foe, to save a princess.

For our purposes, while I want to preserve the distinction between the embedded narrative elements of a game and the emergent ones, I also use "embedded narrative" as a holistic term that encompasses all of the aspects of a work that are embedded in a narrative world and convey narrative meaning. I argue that embedded narrative includes more elements than we currently accept—it is not just the story and plot, cutscenes, dialogue, text and so forth. Setting is conveyed through the specific art assets chosen for a particular environment, its lighting, the design of movement through the space, and so forth. Character is conveyed through AI behavior, how a character looks and stands, character placement in a scene, animations of a character swinging her feet while she waits for the player without dialogue [50], the character tag in *Crusader Kings 3* that also signals how an NPC typically behaves and will move against the player, the specific inflections in the vocal performance when the player character says "Dammit!" as she falls to her death [31], and so forth.

2.3 Games and Narrative Are Not at Odds

I believe in the narrative potential of playable media, while still recognizing the importance of mechanics, simulation, system modeling, and rules to that potential. I believe that agency *within* a narrative is the true magic of our medium—this is why I have devoted such a significant research effort to its pursuit.[2]

Games and narrative are not at odds.

Or at least, they do not *need* to be. Early games studies spent years hashing out the particulars of whether and how narrative and gameplay can coexist [169]. That

[2] As a professional narrative designer in the game industry, I have tended to fall into the more systemic side of narrative research efforts in my positions. I was the first person at Telltale Games to have the title of Technical Designer (though Telltale's Gameplay Programmers did quite a lot of design work). This was largely, I think, to distinguish that my skill set was very different from the other narrative designers there, and my approach was much more systems- and technically-focused. In other industry roles, I have found that narrative designers usually think of me as a programmer and programmers usually think of me as a narrative designer. I prefer lenses to boxes. I am often the person who bridges things that are presumed to be in tension by reframing the problem so that the conflict disappears.

conversation continues today [177].

This argument alone has spawned many entire books, and I cannot hope to be exhaustive on this topic. In honesty, I am not interested in revisiting old debates here. Instead, I am including my position to make my biases explicit.

Much of the industry design knowledge on narrative design is meant to offer approaches to the integration of narrative and player action. This knowledge usually situates itself along various axes of properties that are believed to be in tension. These axes include:

- Narrative vs mechanics
- Narrative vs agency
- Embedded vs emergent narrative
- Authored vs procedural narrative

At their core, these feel like reiterations of the same point: a fundamental struggle between author and player. Much ink has been spilled over the conflict between author and player, but this tension is not usually felt in the way it is presented (e.g., that the player is "ruining" the story) [129]. Chapter 3 will explore questions of agency in more detail, and particularly the points that agency is both bounded by a system and requires player buy-in and reframing will provide more insights as to why this is not as much a problem in practice as the amount of hand wringing over it might lead outsiders to think. But also, if we take a more inclusive look at "narrative" that moves beyond "plot" and "character", we can see that narrative soaks into all aspects of a game that takes place in a fictional world. It is visible in every art asset, every enemy behavior script, the names and placement of items, the framing of the resources players protect or strive for, and actions the player is asked to take.

2.4 Games Offer a Variety of Pleasure-Experiences

To move away from words like "fun" and "play," I use the term *pleasure-experience* to refer to the pleasurable facet(s) of a work, noting that this might also include evoking such emotions as catharsis, disgust, or cerebral reflection. Works of tragedy or horror for example offer a pleasure-experience by inciting pleasures such as tension/relief cycles or escapism. All playable media—and arguably all art—strive for some kind of pleasure-experience. Even the most opaque, pedantic piece might offer the pleasure-experience of fostering intellectual banter.

The idea that different games offer different pleasure-experiences has been captured since the beginning of games research. In his foundational research on games and play across cultures *Homo Ludens: A Study of the Play-Element in Culture*, Johan Huizinga [72] identifies that the word for "play" is divided in interesting ways across languages, with some languages separating the concept and word for "play" from related words like "toy", "game", "child's game", "contest", etc. and some combining those concepts into a single word in interesting ways. The variety of words for adjacent concepts and pleasures indicates that humans have long-recognized that different forms of play carry different pleasures.

Similarly, Roger Callois argues [23] for four distinct forms of play that fall along a spectrum from *ludus* (structured play with explicit rules) to *paidia* (unstructured play):

1. Agon - competition
2. Alea - chance
3. Mimicry - make-believe or role-play
4. Ilynx - vertigo or the bodily sensations of dizziness, altered perception, or loss of control (e.g., the pleasures of a rollercoaster or drugs)

We see analogous arguments made in more recent game studies in the enumeration of different player types, such as Richard Bartle's famous taxonomy of player types [13], which categorizes players of Multi-User Dungeons (MUDs) along two axes: those who enjoy acting autonomously vs those who enjoy interacting with others, and those who enjoy acting/interacting with other players vs those who enjoy doing so with the environment. These axes then split players into four quadrants or "types" (see also Figure 2.1):

1. Achievers
2. Explorers
3. Socializers

4. Killers

These categories have been contested and expanded ever since. [69] And naturally if we can concede that games offer different pleasures to different player types, we can surmise that there are a range of pleasure-experiences toward which any individual game might be striving.

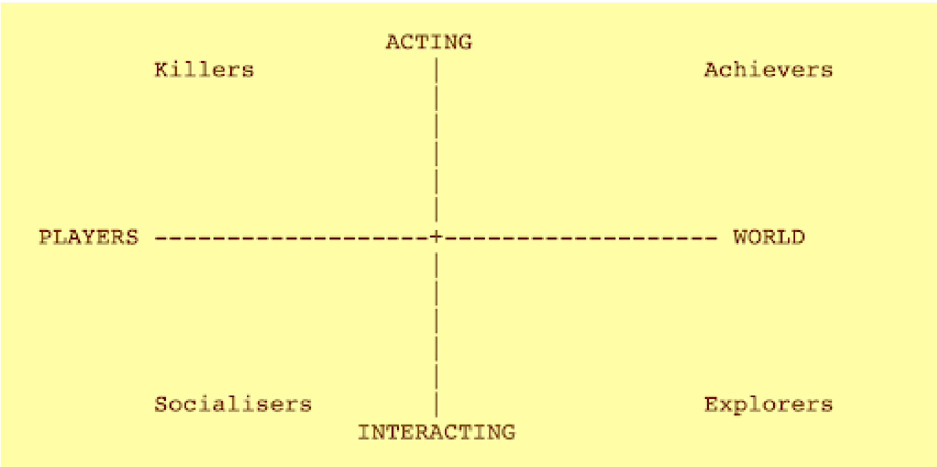

Figure 2.1: Image from Bartle's taxonomy of player types [13]

Identifying that all games are striving for some pleasure-experience, and that games striving for different pleasures will have different aesthetic focuses, is one of those notions that is superficially obvious, but carries implications that are often overlooked when we actually critique games . *Chess* is striving for a very different pleasure-experience from *The Sims* [108]. We might call these different forms of play in the tradition of Caillois—*agon* and *mimicry* respectively. [23] Or we might lean into motivation theory, which suggests that all humans have three psychological needs—autonomy, competence, and relatedness—that govern our motivation to engage with games [49], thus positioning *Chess* more as a game of competence whereas *The Sims* might be a game more about autonomy through the customization of a character and environment or relatedness in the sharing of narrative situations.[3] An ontology of pleasure-experiences is outside the scope of this dissertation, but I offer these two different lenses here to point out that a game does not only offer a single pleasure-experience, pleasures may be nested or exist along different dimensions, and different lenses may yield different interpretations of what pleasure-experiences are important.

2.5 Intended Pleasure-Experience Matters to Analysis

One of the reasons games are so prone to the rehashing of essentialist arguments of which elements we should assess in what makes a game "good," is that game criticism as a field has not come to recognize as a matter of course that different games offer different pleasure experiences. Though many critics reference Huizinga or Callois or Bartle, comparatively few of them stop to ask, "Does my analysis make sense for the intended pleasure-experience at hand?"

In my role as a critic of interactive narratives, I ask myself a series of questions when analyzing a work:

1. What is the piece trying to do?
2. What am I seeing that led me to this conclusion?
3. What techniques does the work use to try to achieve this goal?

[3]In truth, both of these characterizations are reductive and both games arguably utilize all three motivators in different ways.

4. Do those techniques achieve the goal? Why or why not? Do they suggest other pleasure-experiences as well? Which ones?

5. What, if anything, could the work have done better to try to achieve its goal?

We do a lot of games criticism without actually acknowledging that different games are trying to do different things. The narrative in a first-person shooter (FPS) does not have to be great narrative if the game is really about competence, mastery, and the thrill of fast-paced action. In such a case, the narrative is only there to provide enough motivation and quickly-understood metaphor for players to understand how the systems will work. If a gun has "ammo" to pick up, for example, that provides an easy-to-understand shorthand that my weapon will run out of ammunition at some point and I will potentially need to find more. Arguing that FPSs should have better narratives is a fine argument to make; holding a particular FPS to the standards of narrative we would expect from a tragic film is not. The former argument advances the field, while the latter prevents a useful discussion about the aesthetic offerings of the work at hand.

Of course, the death of the author is relevant here [12]. We cannot actually know what the author intends, and I do not mean for this approach to advocate for presumptions based on a creator's persona, politics, or even her past work. Instead, I mean that when analyzing a work, critics should judge a work by the standards of the pleasure-experience the work seems to be trying to offer.

Games might have more than one pleasure-experience; a first-person shooter might want to have both thrilling fast-paced combat and an excellent story, but often a game is not actually interested in narrative as a core pleasure-experience. Good criticism will judge those games relative to the aesthetics to which they aspire. We would not judge a Shakespearean sonnet by the same metrics as a pulpy domestic thriller, even if both are ostensibly about the agonies of love. Yet in games scholarship we tend to do this surprisingly often. Many of the essentialist arguments of early game studies feel very strange today in this regard.

Meeting a game (or any art, for that matter) at the point of its aesthetic intention, as best as we can understand it, is an important and foundational assumption to critical analysis. By talking about what elements we find pleasurable in a particular game and grouping games that offer similar pleasures, rather than starting from a unilateral (and fraught) definition of "what a game is", or "what a game of type X is" allows us to have more nuanced discussions of the aesthetics of these particular pleasures. If we start by asking the question of what a game's intended pleasure-experience might be, especially when comparing it to another game, is a good first-line defense against falling into fruitless, essentialist arguments.

As I construct an argument toward a theory of responsiveness in games, I draw attention to these assumptions for the following reasons:

1. I hope for my theory of responsiveness to be used as a lens rather than a box.

2. I want to acknowledge that not all games are particularly aiming for responsiveness as a core pleasure-experience. This theory is not intended to be a value-judgement against such experiences.

3. In discussing the narrative aspects of games across genres, I want to acknowledge that many of the games under discussion carry an additional pleasure-experience at their core. For some games, this experience supports its narrative ambitions, for some the narrative experience is secondary. It would be irresponsible to critique narrative aspects of such games without acknowledging that the narrative is supporting a fundamentally mechanical experience.

Chapter 3

Models of Player Agency

Agency is a prolific topic, both within the study of games and digital media, and across other disciplines, including philosophy, cognitive science, psychology, political science, economics, and cultural studies, to name only a few. The topic of agency, approached from various critical and scholarly traditions, could easily fill several dissertations. Therefore, for scope and logistical limitations, I will focus our inquiry on agency as it is discussed in game studies with a lineage largely attributed to Janet Murrays' *Hamlet on the Holodeck*.[1]

Even within this strict definition, agency has been such a central focus of game studies and digital media since early theorizations, that even a limited focus on agency as it is understood across games studies and digital media research invites at least a whole dissertation to be written on agency, the evolution of our understanding of it, and its impact on approaches to the design of digital media. For the sake of brevity, I must

[1] Wardrip-Fruin et al rightly note that the concept has been present under different terminology since the earliest known dissertation on interactive narrative, that of Mary-Ann Buckles in 1985. [193]

scope my area of study here to current understandings of agency that will be relevant to the formation of a theory of responsiveness.

Still, a review of current literature on agency in games is necessary. To understand how a system responds to a player's actions and intentions, we must first understand how those actions and intentions work. Over this chapter, I will briefly revisit relevant theories on agency and interactivity, specifically in the context of how players recognize their own agency in game narratives, with an eye toward synthesizing much of the current research and design knowledge. My goal in this endeavor is not to exhaustively review the evolution of our conceptions of agency, but rather to draw on current understandings with an eye toward building a model of what it means for a system to be responsive. In the process, we will look at the role agency plays in interactivity, and how that role opens up a space of exploration into how a system signals to a player that she has been heard, that her agency has been recognized.

Interactive narrative by definition considers interactivity to be foundational, and agency is a crucial component of fostering interactivity. Though the two terms have often been used interchangeably in the research literature (see section 3.2), I will adopt definitions of the two that draw important distinctions between them in an effort to help separate the phenomenon that a player experiences from the elements of the system's design that foster and reinforce that experience. In my conception, interactivity is a cycle of engagement between two agents, while agency is a phenomenon that occurs from the perspective of one agent under specific conditions.

In addition to interactivity, agency is often mentioned alongside several related concepts. In reviewing the literature on agency, I hope to separate agency as a concept from the surrounding concepts with which it is often mentioned, to show where notions like choice, consequence, and roleplay overlap with it, and to bridge current understandings of agency between interdisciplinary communities of game studies research and the understandings and "common knowledge" of game development practitioners.

In order to understand how to create responsive systems—that is, systems that deeply react to player input and allow for a wide degree of agency—we must first build toward a common understanding of agency.

3.1 Agency is a Central Pleasure of Interactive Media

Agency and interactivity are central pleasures of our medium. The components of how interactive media foster the pleasures of interactivity have thus become a natural area of inquisition within the study of games and other computational media.

Much of the current thinking on agency in interactive narrative descends from Janet Murray's *Hamlet on the Holodeck*, [118], which makes it a natural place to begin our exploration. Among other aesthetic pleasures like immersion, and transformation, Murray identified agency as a key component of interactive narratives as early as 1997, defining it as "the satisfying power to take meaningful action and see the results of our decisions and choices". Emphasizing its continued relevance in the 2016 reprint of Hamlet on the Holodeck, Murray writes, "The experience of agency by the interactor is the key design value for all digital artifacts… Dramatic agency should be the goal of design for interactive narrative in any form.” [119]

Since Murray’s introduction, the nature of agency in interactive media has been debated, and the term used in different ways. While Murray’s original conception centered the player’s perceptions as the point at which agency occurs, Mateas argued that agency is bounded and expressed through the affordances of a system [99]. Wardrip-Fruin et al later expanded this notion to argue that agency is relative to the player’s mental model of the computational system within which their choices are expressed [193]. This key difference in centering—whether agency occurs as a phenomenon within the player’s perceptions or as a cycle between both the player and the system together—represents one of the more interesting and promising discussions in the field, which I will address later in this chapter.

In framing agency as a foundational pleasure of the medium, however, it also becomes tempting to reframe other tangential pleasures as agency, and the term has been used in ways very separate from the discussion above. Because of the variations in the use of the term, it has become difficult to universally characterize agency, and I would argue that this difficulty has resulted from using it as a catch-all to describe various phenomena: from the ability to successfully navigate a web page to the ability to employ a preferred tactical strategy in a puzzle game, to the ability to take intended actions in a story, and the identification with a player character in a way that feels natural and pleasurable. Many of the things we call “agency” are operating at different levels of abstraction and granularity of player action, so comparisons of different “types” of agency have often further confused what we mean by the term. In the worst cases, agency becomes a fallback for "everything that makes a playable media experience pleasurable." This ambiguity not only makes conversations about agency confusing and fraught, but it also subsumes other pleasurable aspects of games—like enactment, roleplay, narrative or mechanical immersion, and strategic planning—that deserve recognition and exploration on their own terms.

While I follow Murray's proposed use of the term, our understanding of agency has expanded dramatically since her initial exploration and now includes more nuance and caveats, many of which we will explore over the rest of this chapter.

3.2 Disambiguating Agency and Interactivity

As our field was beginning to define itself in descriptive and demarcating terms, "interactivity", a term "widely used in common parlance and in advertising" [150] was contrasted against other terms to try to understand its nature. Despite Murray's attempts to distance the idea of agency from interactivity, the terms have been intermingled in confusing ways. Indeed even recently "interactivity" and "agency" are still used somewhat interchangeably.

Several game studies authors note the ambiguity of "interactivity." Game developer Chris Crawford, in his book *Chris Crawford on Interactive Narrative*[2], begins a chapter devoted to interactivity with a critique of how the word is used in academic

[2]A note here to expressly say that this book has not aged well. Crawford seems to be writing to bridge gaps between two audiences—commercial game developers and academics. However, the tone of the book smacks of the juvenile boys'-club culture endemic to games culture at the time of its writing. Nevertheless, the model Crawford presents and its applications to interactive narratives remain a useful lens through which to explore interactivity.

circles to poor effect. He writes:

> Interactivity is without doubt the most grossly misunderstood and callously misused term associated with computers. Everybody has been using the term for so long that people are quite sure of their appreciation of interactivity. The problem is that everybody seems to have a different conception of interactivity, and most descriptions are fuzzy and accompanied with lots of arm-waving. [35]

He goes on to draw distinctions between interactivity and several related concepts, arguing that *interactive* is not synonymous with words like *engaging*, *immersive*, *participatory*, *reactive*, or *responsive*. His thoughts on the last two—*reactive* and *responsive*—are of particular interest to our endeavor. In comparing the words to interactivity, he writes of them:

> **Responsive:** This attribute hits on a fragment of the truth, but remains too vague to illuminate the nature of interactivity. Yes, interactivity requires responsiveness, but it's the character of that responsiveness that conveys the value of interactivity.
>
> **Reactive:** Again, this attribute is partially correct—but don't make the mistake of confusing reaction with interaction. Reaction is a one-way process; interaction is a two-way process. Two people interacting are engaging in a series of reactions to each other. Reaction alone is only a subset of interaction, however. And intense reaction is not the same as interaction; you can't turn reaction into interaction merely by turning up the volume. Reaction is to interaction as moving your left foot is to dancing. [35]

We will return to some of these ideas: that responsiveness conveys the value of interactivity, that it constitutes a component piece of the interactivity model, and that responsiveness is distinct from *reaction*. For now, it is worth noting that agency is not among the terms Crawford tries to distance from interactivity. In fact, later in the book he argues that the two words cover roughly the same concept, writing "Academics prefer to use the term agency to refer to what most people call 'interactivity.'". This conflation seems to lose the player-centering in Murray's original definition, and feels contradictory to the idea that interactivity is a two-way process, yet Crawford is not alone in his conflation of the terms.

In her chapter on "The Many Forms of Interactivity" in *Narrative as Virtual Reality 2* [150], Marie-Laure Ryan points to Murray's *agency* as potentially analogous to her own use of *interactivity* while acknowledging that Murray's intention is to create more precise language around "purposeful actions that alter a world meaningfully" than those offered by the term *interactivity.* In trying to enumerate a taxonomy of the various forms of interactivity, Ryan largely describes her classifications in terms of the actions users take and what kinds of control they have, suggesting that her use of the term also incorporates the kinds of purposeful actions with which Murray is concerned. Distinctions between Murray's use of the term and Ryan's become muddled, and the two seem to be using *interactivity* and *agency* to refer to similar concepts, supporting that she considers the two terms fairly interchangeable.

Ryan also entertains the premise that her use of the term "interactive" may be similar to Espen Aarseth's use of *ergodic* in his definition of *ergodic literature* as a class of works in which "non-trivial effort is required to allow the reader to traverse the text" [1]. Aarseth specifically rejects the term "interactive", arguing "the word interactive [...] connotes various vague ideas of computer screens, user freedom, and personalized media, while denoting nothing". And while Aarseth's definition of ergodic literature has largely been adopted by the electronic literature community as a means for understanding the kinds of works that community is interested in creating and studying, many such works do not center reader choice, or even necessarily reader participation, while still falling under the definition of requiring non-trivial effort on the part of a system or external force. Consequently, Ryan concludes that *ergodic* and *interactive* actually have distinct meanings:

> According to Aarseth's definition, ergodic literature is a class of works in which "non-trivial effort is required to allow the reader to traverse the text" (1). In this line of thought, turning the pages of a book is a trivial effort, but clicking on links while surfing the Web or reading hypertext fiction is not because these activities involve choice. But if we understand ergodic design as a built-in mechanism or feed-back loop that enables a textual display to modify itself, so that users will encounter different sequences of signs during different sessions, then the "non-trivial effort" does not necessarily come from the user. For instance, "The Speaking Clock," a digital poem by John Cayley, generates an ever-changing display of words selected from a fixed textual database, using the computer clock, as well as aleatory selection devices, to trigger transformations. Since ergodic texts may be closed systems that operate without human intervention, I believe that ergodism and interactivity are potentially distinct properties, even though they often fall together, and that interactivity as human intervention retains a legitimate place in the toolbox of media studies. There are noninteractive ergodic texts, such as "The Speaking Clock," or the output of an Eolian flute operated by the wind; nonergodic interactive texts such as conversation (which, as free-flowing exchange, lacks the global design of ergodism); and texts that are both ergodic and interactive, such as hypertext and computer games. [150]

Here, in her reference to both hypertext works and computer games, and throughout places in the rest of the chapter, Ryan's use of the term *interactive* seems to indicate that she is using the term as a superset of Murray's *agency*, as a phenomenon of back-and-forth action between the user and system—an implication which contradicts the previous observation that agency and interactivity are analogous. To add to the confusion, in the section on "Reconciling Immersion and Interactivity," Ryan does adopt the term "agency" in places, but usually only in referring to the user's ability to enact desires upon a system.

Crawford's notion of interactivity seems to parallel Ryan's, both in their conception of interactivity as a back-and-forth participatory cycle between two parties, and in their use of conversation as the prototypical case of interactivity for which other definitions should be able to account. Crawford, in fact, uses the metaphor of conversation as a central concept for his definition—one of the more clear and succinct definitions of interactivity and the one which my model adopts:

> I offer this definition of interactivity: A cyclic process between two or more active agents in which each agent alternately listens, thinks, and speaks. [...] The value of this definition lies in its reference to conversation, a well-known form of interaction. Our experiences with conversation offer useful guidance in software design. Obviously, the overall quality of a conversation depends on the particular quality with which each step (listening, thinking, and speaking) is carried out. Even more important is the way those three qualities combine. [35]

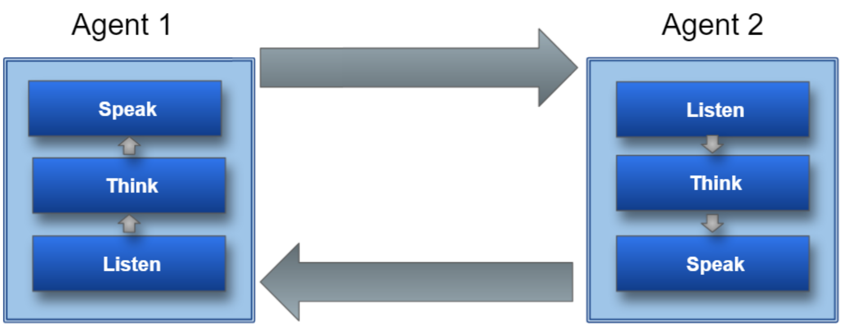

Figure 3.1: Depiction of Crawford's Listen-Think-Speak (LTS) model of interactivity [35]

Crawford's Listen-Think-Speak definition of interactivity (which I shorthand to LTS throughout this dissertation) emphasizes three points that become important underlying assumptions for us:

1. Interactivity is a cyclical process between two parties.

2. In order for interactivity to occur, both parties must recognize the output of the other, interpret it, and respond.

3. The quality of each component step, and the way those steps combine, influences the experience of the overall interaction.

We might illustrate the LTS model as depicted in Fig 7.1. The focus on two parties supports Murray's own suggestions that interactivity and agency are indeed two distinct phenomena. In her chapter of *Hamlet on the Holodeck* titled "Agency," Murray makes an effort to distinguish agency from interactivity, noting that interactivity does not necessarily require the player to have any specific intention. She writes:

> Because of the vague and pervasive use of the term *interactivity*, the pleasure of agency in electronic environments is often confused with the mere ability to move a joystick or click on a mouse. But activity alone is not agency. For instance, in a tabletop game of chance, players may be kept very busy spinning dials, moving game pieces, and exchanging money, but they may not have any true agency. The players' actions have effect, but the actions are not chosen and the effects are not related to the players' intentions. [119]

Though Murray seems to be using interactivity as a shallower form of action-taking and system-response than true agency, I would argue the examples of interactivity she provides instead offer the full interactive loop at a more granular interaction level, where her definition of agency is instead focused on a more abstract level of interaction intention. We will return to the idea of levels of interaction later in the chapter (section 3.3.6). Still, though she does not distinguish agency from interactivity as we are defining it here, her distinction reveals how we might.

Murray's implication that the player's will is a central component of agency centers her conception of agency on one side of LTS diagram: the player assesses her options (listens), forms an intention (thinks), and then takes an action (speaks) then waits for the system to respond to confirm whether the intention was recognized and incorporated into the game state—indeed we could even represent agency as one half of that diagram *from the player's point of view*. We might construct the diagram from the player's perspective as depicted in Figure 3.2.

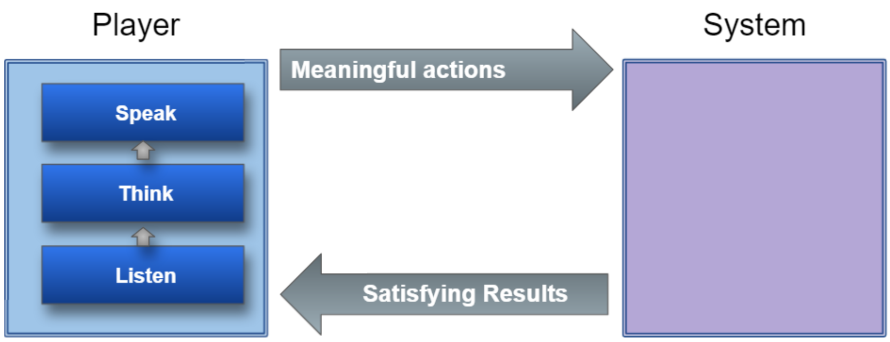

Figure 3.2: Reformulation of Murray's agency as enacting one side of the LTS model from the player's point of view

Thus, for our purposes, agency and interactivity are two distinct phenomena. Subscribing to Murray's definition, agency is "the satisfying power to take meaningful action and see the results of our decisions and choices", and is a subcomponent of interactivity. It is only the portion of the listen, think, speak model that is experienced by the player. Interactivity, by contrast, is the entire loop as experienced by both sides.

This distinction is important, but it is only a starting point. In the next section we will review current research conceptions of agency in an effort to add nuance to our understanding and further explore the conditions under which agency occurs.

3.3 Understanding the Experience of Agency

By taking Murray's definition of agency as "the satisfying power to take meaningful action and see the results of our decisions and choices" as a starting point, and incorporating other researchers' insights about agency, this section will explore agency's relationship to intention, consequence, and context in an effort to build a common understanding on which to create our theoretical framework. In doing so, my goal is to create a theory of agency that synthesizes current understandings from researchers and industry practitioners alike.

Based on our current understandings, we can glean the following information about agency:

1. Agency requires intention
2. Agency is bounded by a system
3. Agency requires buy-in and reframing within that system
4. Agency is proprioceptive and requires feedback in order to be felt
5. Agency is not dependent on eventual consequences
6. Agency may exist at different levels of granularity simultaneously

Many of these points are interrelated. I have tried to present them in a logical order that allows each point to build on others, but in some ways this presentation is arbitrary and the interdependence of these points means there will necessarily be lots of references to other sections and chapters. I appreciate the patience of the reader and invite you to jump around if digressive tangents better-suit your needs.

3.3.1 Agency Requires Intention

In building our understanding of agency, it makes sense to start with Murray's foundational definition and build from there. For Murray, agency as "the satisfying power to take meaningful action and see the results of our decisions and choices" carries player intention at its core. Murray's distinction between interactivity and agency, in fact, seems to rely on intention as the most important distinction between the two:

> But activity alone is not agency. For instance, in a tabletop game of chance, players may be kept very busy spinning dials, moving game pieces, and exchanging money, but they may not have any true agency. The players' actions have effect, **but the actions are not chosen and the effects are not related to the players' intention.** (emphasis added) [119]

The focus on intention here mirrors colloquial use of the term, and its use across other academic disciplines which concern themselves with the capacity of an actor to carry out their will or express power (cf agency in philosophy, cognitive science, politics, critical theory, etc.). Agency really is *about* intention.

More specifically, Murray links the idea of intention with the idea of choice. After all, if a player does not have more than one action available to them, their intention does not matter; they may only pursue a single course of action. Thus agency is dependent on the idea that at least one course of action is presented to the player, and that the player's intention to pursue a particular course—to choose one action over another—is the foundation for agency.

An often-cited Sid Meier quote that "games are a series of interesting decisions" [3] seems to underscore the importance of agency and intention. And a point that often arises among practitioner talks and guides is the idea that choices must be sufficiently framed so that players can formulate some idea of the consequences of a choice [93] Short). The reason this framing is so important is not only because a random choice is hardly an interesting one [3], but also because an ill-framed choice does not allow the player to form or understand intention. Intention requires the player to have some level of investment in the outcome. Choice without intention is not agency; it's a dice-roll.

Some researchers, however, have disputed the idea that agency and intention are related at all. Muriel and Crawford, for example, taking an approach grounded in actor-network theory, argue that the crux of agency is the ability of an actor to make changes or transformations. They write, "Agency exists because, in some way, it transforms reality. Agency, therefore, does not have to do with the intention, desire, or the will of an actor, but rather with the transformations that occur; which are effectively observable and traceable." [117]. This argument effectively implies that agency is more about counterfactuality[3] among the outcomes of an action as observed from the outside than investment in one outcome over another by the acting agent. While their framing is useful in describing the "paradox" of a player's agency as necessarily bounded by the system within which that agency occurs, to focus on the outcome of a choice to the exclusion of the intention of the agent is to miss the very core of what makes games

[3]Counterfactuality here is used in the sense of a counterfactual conditional: situations in which something else could be true in different circumstances.

pleasurable. After all, actor-network theory gives just as much agency to a dice-roll as to a choice, yet most players would argue that the two have very different aesthetic experiences.

This additionally raises the question of whether counterfactuality must exist for agency to occur. For Muriel and Crawford, and many players, the answer is an emphatic yes. The question has raised interesting thought experiments [102] and philosophical questions of free will in games [130] [185], but for our purposes, a philosophical digression can be avoided for reasons discussed later in the chapter.

As we will see in future sections, intention does not necessarily mean that the player must be able to "do anything", nor that the player must get exactly what they want. And certainly agency does not preclude unintended consequences. But rather, agency requires that the player is able to express intention, given the tools available, in a way which the system recognizes and validates.

3.3.2 Agency is Bounded by System Affordances

If we think about agency only as a phenomenon focused on the realization of intention, it is easy to frame it as the player "being able to do whatever they want" or to get lost in discussions about free will.

To focus *only* on intention ignores the fact that play is naturally bounded by players' expectations, and the constraints and affordances offered by the system. When we play a first-person shooter, for example, the fact that we cannot stop to write poetry for the enemies does not feel like an imposition on our agency. If we found we could not fire our gun, on the other hand, that would certainly feel like a limitation, despite the fact that firing guns in real life is a relatively rare occurrence [193].

In "A Preliminary Poetics for Interactive Drama and Games", Michael Mateas describes a poetics for a Neo-Aristotelan theory of interactive drama [99]. In the context of interactive drama, Mateas incorporates the idea of an *affordance* from interface design [125] in which an element offers an action to the user:

> In interface design, affordances are the opportunities for action made available by an object or interface. But affordance is even stronger than implied by the phrase "made available"; in order for an interface to be said to afford a certain action, the interface must in some sense "cry out" for the action to be taken. There should be a naturalness to the afforded action that makes it the obvious thing to do. For example, the handle on a teapot affords picking up the teapot with your hand. The handle cries out to be grasped. [99].

Within this context of affordance, Mateas introduces *formal affordances*, those offered by the narrative and representational elements of an interactive drama that motivate a player to take an action, and *material affordances*, or the affordances offered by the systemic components of an interactive work as supported by its inputs and computational model.

As players, we recognize formal affordances as a dramatic (narrative) possibility space, constructed from our understanding of the plot, genre, framing scenario, our character, the characters around us, and so forth. We know that the protagonist of a noir detective story usually performs certain actions, so if we find ourselves playing a character a noir detective in an interactive story, we would naturally expect to be able to perform those actions. We would probably not expect to be able to fly, regenerate limbs, or lift cars above our heads, though these would all be reasonable formal affordances for a story in which we play a superhero. As we progress through the story, certain narrative actions "cry out" to be taken.

Material affordances are the actions the system allows us to take. Each system recognizes certain inputs and maps those inputs onto actions in the game. The material affordances are those the system recognizes as valid and provides to the player for use. We will look more closely at the nuances of both formal and material affordances in Chapter 3.

Mateas argues that agency occurs when both the formal affordances (what the player is motivated to do by the representational layer) and the material affordances (what the player *can* do given the constraints of the system) are in balance.

> A player will experience agency when there is a balance between the material and formal constraints. When the actions motivated by the formal constraints (affordances) via dramatic probability in the plot are commensurate with the material constraints (affordances) made available from the levels of spectacle, pattern, language and thought, then the player will experience agency. An imbalance results in a decrease in agency. [99]

When formal and material constraints are in balance, agency is achieved, and an imbalance of the two decreases agency. If a game offers formal affordances that suggest and motivate a possible action that the material affordances do not support, the player feels cheated that the computational model is not robust enough to support her intentions. For example imagine a parser-based interactive fiction piece that describes a small jewel sitting on a table. The player tries to TAKE the jewel, but the author has not modeled the jewel in the room's model, so the parser prints the message "You can't see any such thing." Here the player experiences a break in agency; the system has not actually modeled the items described, and even though the room description has suggested the possibility of a certain action, the system fails to support that action and in the process informs the player that her narrative mental model of the world is wrong. Similarly, if several actions are available to the player but either she does not know she can take them or she doesn't understand why she should take one particular action over any other, the system is offering more material affordances than formal affordances; in this case, again, agency is interrupted. In order for agency to be maintained, the balance of formal and material affordances must remain intact.

Throughout the rest of this chapter, as I examine other dimensions of agency and its pleasures, I will ground these examples in this model and use the language of formal and material affordances to discuss them. And as we construct a model of responsiveness in game systems, we will return to this definition of agency as the balance of formal and material affordances continually.

The idea that agency is a phenomenon bounded by a system is further elaborated in "Agency Reconsidered," by Noah Wardrip-Fruin et al.:

> The concept of "agency" in games and other playable media (also referred to as "intention") has been discussed as a player experience and a structural property of works. We shift focus, considering agency, instead, as a phenomenon involving both player and game, one that occurs when the actions players desire are among those they can take (and vice versa) as supported by an underlying computational model.
>
> [...]
>
> Agency is not simply "free will" or "being able to do anything." It is interacting with a system that suggests possibilities through the representation of a fictional world and the presentation of a set of materials for action. Designing experiences toward the satisfactions of agency involves balancing the dramatic probabilities of the world with the actions it supports. In other words, the design task is to entice players to desires the game can satisfy—whether

> this is traveling across space, managing resources, engaging in battle, or making conversational moves. [...] Supporting agency requires employing or crafting a computational model of the play domain suggested by the work's dramatic probabilities. Depending on an inappropriate or overly-simplified model leaves the designer with two choices: extremely constrained input (so that players are effectively not able to play) or breakdown as the work is unable to continue coherently suggesting dramatic probabilities and players lose faith in the materials provided for action. [193]

One point implied by this discussion, is that players are constantly building and revising a mental model of the computational system as they play. A player comes to a work of digital media with preconceptions about what the work is, and thus what they will be able to do within it. As they play, the system's feedback about both its formal and material affordances changes the player's mental model, and ideally their mental model of those affordances ever-changes to more closely resemble the intended experience of the creator.

> Players come to games with assumptions about the domain of play. To play successfully they must transition from their initial assumptions about this domain (e.g.,, movement or conversation) to an understanding, often largely implicit, of how it is supported by the software model. Because we do not have a "Holodeck" this will inevitably be different from how it is supported in the everyday world, though it may be quite close to the support found in other games (which also contribute to many players' initial assumptions). [193].

This process explains why we do not feel a lack of agency in the example of not being able to write poetry in a first-person shooter: our mental model of what the game allows does not tell us that we *should* have the ability to write poetry.

3.3.3 Agency Requires Buy-In and Trust

In addition to framing agency as experienced within constraints, I want to explicitly highlight that in this process, the player must undergo a buy-in toward these constraints.

This buy-in reframes the player's expectations of their array of choices to align with actions available to them in the system. One way to read this is as a reiteration that formal and material affordances are in balance. After all, the idea of situating oneself within formal affordances implies reframing within formal constraints. However, the idea of buy-in additionally positions narrative agency as a phenomenon that occurs within a magic circle [72] and takes place with the willing participation of the player. As such, the player is engaging with a kind of *trust* in the system in which she expects that the system will uphold its fulfillment of her intentions if she takes the specified actions.

Karen Tanenbaum and Theresa J Tanenbaum explicitly move to distance the idea of agency as "freedom" with an emphasis on choice, and instead position agency as "the process by which participants in an interaction commit to meaning" [179]. They argue that while the freedom offered by agency is often an illusion, and that the discussion of whether this illusion is real is misguided. By framing agency as "commitment to meaning," the authors offer useful expansions to our understanding of agency; namely, by rearticulating agency in this way:

1. We are able to understand places where explicit lack of choice leads to deeper role

play and narrative immersion.

2. We are able to understand places where players feel they have been let down in the system's response to their commitment. [179]

To illustrate the first point, Tanenbaum and Tanenbaum cite an experience from *Metal Gear Solid IV* [131].

> What makes this sequence remarkable is how little control the player has over any aspect of the experience. In it, the player is told that he must maneuver Snake down the tunnel quickly or he will die. As the character moves through the gauntlet of microwaves, the top half of the screen shows the final moments of a climactic battle in another location. Suddenly, there is a flash of sparks from the wall of the tunnel and Snake collapses to the floor. A voice over the radio implores "Don't give up on me, Snake!" and the player is prompted to tap the Triangle button on the controller in order to keep Snake moving forward. Over the next moments of gameplay this escalates: Snake grows ever more crippled, the corridor becomes more and more hazardous, and the prompting from the system demands ever increasing button mashing in order to move the character forward.
>
> [...] This sequence leads the player to buy into the drama and the desperation of the situation. Through carefully arranged cues the player is invited to frantically mash a single button in order to advance the game. The beauty and irony of this is that even this limited set of interactive possibilities – to rapidly push the button or not – is not actually required by the system.

The authors use this example to note a place where the player experiences agency even in the absence of many of the things we might associate with agency. There is no counterfactuality; the scene will not end in a different way depending on how the player participates. There is also no real choice (nor does it occur to the player to make a different choice). But still, the formal affordances are very much balanced with the material affordances, and the player feels that she has elected to keep moving forward; she feels responsible for Snake's success.

Again, we can feel the counterfactuality question resurfacing: this example raises the question of whether a single choice gets to "count" as agency, even if that single choice represents a balance of formal and material affordances. It is easy to dismiss this example as enactment, to dismiss it as a different pleasure from agency. But if we approach this argument as a lens rather than a box, if instead of saying "does this count as agency?" we ask "if this is agency what does it teach us?" it yields interesting observations.

First, the act of engaging, of making a choice as the player's character even as the formal constraints tighten around her to the point that there is only one choice to make, fosters and reinforces roleplay and identification with that character. This idea is echoed in both Murray and Ryan, but the idea of this process feeling like a conversation between designer and player emphasizes a collaborative trust in a way I had not previously seen articulated.

Second, it deemphasizes the idea of "choice" and instead emphasizes intent. The authors elaborate on this point:

> This definition, that agency is the process by which participants in an interaction commit to meaning, is particularly well suited to interactions with narrative and story-based games. This understanding of agency proposes that game designers should strive to create game and narrative experiences in which the player can demonstrate commitment to the experience, and, crucially, where that meaningful commitment is reinforced by the game's behavior. Agency is not about selecting between options in this case, but is instead about expressing intent, and receiving a satisfying response to that intent. Commitment in this sense might be a purely cognitive process, or it might involve player actions.

Both of these points together become very important not only to roleplay, but to the narrative mental model we construct as we decide how to *play* as our character. When I play a game that gives me different options around how to roleplay a character—do I want to be the righteous hero? The deceitful rogue? The tyrant ruler?—I find myself constructing a backstory and taking actions that align with that intent. At Telltale, we called these *roleplay rails* and specifically authored choice-spaces to afford certain roleplay rails [93], which were decided upon before a single scene was written. Would a player choose to be an honorable Batman or a vengeful vigilante? Would she value friendships or self-interest? Rails were signalled to the player by early decisions, and often a player would commit to an identity that they intended to carry through their experience.

Intent colors the player's motivations; in other words, it shapes the mental model against which she judges any formal affordances presented to her. As a game's affordances confirm these desires as valid, the intent itself seems to be confirmed as valid by the system. This validation is usually the place being transgressed when we colloquially talk about a game "stepping on agency," a phrase used to describe places where a game has interjected something about the player's character that runs counter to the way she intended to roleplay that character.

As an example, I had an experience in *Detroit: Become Human* [40] that felt like it lazily stepped on agency through a mismatch of formal and material affordances with one of the characters I played, an android cop named Connor. Over the course of the game, Connor discovers that some of the androids have broken from their programming and are engaged in an uprising. The player can steer Connor toward continuing down a path of lawfulness that will betray his fellow robots or toward becoming a "deviant" himself and aiding the uprising.

In the first scene in which Connor encounters a "deviant" robot, the game has already set Connor up as a "good guy" who highlights the shades-of-grey morality of his grizzled partner. Connor, off by himself, encounters a deviant and rather than the game offering the choice or whether to reveal the information to his partner, Connor just calls out to him on his own. While playing the game I literally called out "Oh my god, that should have been a choice! No! No!" [96]. Even though I later came to understand that the designer was probably intentionally trying to limit the amount of moral agency it offered for Connor at this point in the game because of his eventual arc, with the knowledge I had at this point in the game and the signalling that the game had given me that this might be a turning point for my character, I felt deeply undermined in how I wanted to define Connor through the cumulative effect of my choices. My *intent* was to play Connor as a good character as the game wanted, but with his goodness being defined by his propensity to help his fellows rather than by his lawfulness and willingness to report his fellows to the law. My lack of agency did not come from my lack of choice per se; it came from the game injecting information than ran counter to the formal affordances it had presented in a way that interrupted my narrative immersion; the effect was that rather than feeling a lack of choice, it instead felt like the game had *made a choice for me* that ran counter to the affordances it had presented. In the process, it broke my trust in its ability to carry out my roleplay intent. Even using UI elements to highlight my lack of agency in this particular moment would have offered acknowledgement of my larger intentions and reassured me that my intended roleplay experience was not "wrong"; highlighting my lack of agency in this choice would have reassured my agency in a broader sense of how I wanted to roleplay my character. (see section 3.36: Agency Exists at Different Levels of Abstraction).

Similarly, Tanenbaum and Tanenbaum cite the occasional misalignment of representative text on choices with the actual line a player character says after the choice is made. They write

> For the most part, this two stage relationship between the communicative commitments of the player and of the character works seamlessly. The player commits to a general "flavor" of communication, which is then executed by the character within the predefined conversation possibilities. Occasionally, however, there will be a mismatch: the player will instruct Commander Shepard in what she believes will result in one type of communication, only to receive something unexpected in response. One player describes this on the Mass Effect forums
>
> > "For the most part, I loved the dialogue wheel... but I'd be lying if I said there weren't a few occasions where I selected a response intending for my character to behave or say something a certain way... only to have them either say or (in worse situations) do something I neither intended or wanted to happen. . . One case example was when I was tracking the guy who'd gone off to create a cult of biotics. When I confronted him, I highlighted a response that I thought would probe or provoke him a bit... but instead once I selected the response, I drew a weapon and shot his head off."[1]
>
> The communicative commitments of the player may not always align with those of the game designer, as this example shows.

Both the example cited by Tanenbaum and Tanenbaum and the example of my own experience with *Detroit: Become Human* demonstrate that the experience of roleplay is dependent upon the trust that formal affordances will not only be supported by material affordances, but also that they will be internally consistent and that the player will be the one to act out key moments, she will not have those moments acted out *for her*, especially in ways that run counter to her roleplay intentions.

When a decision is made *for* us that runs counter to our mental model, a system's material affordances have misaligned with its formal affordances; the system has suggested a mental space was viable that turns out not to be. From this, we can conclude that agency does not just apply to individual *choice moments* but it also includes the intention for roleplay as expressed in formal affordances.

3.3.4 Agency is Felt Through Feedback

In reviewing the modern origins of agency in our field, Wardrip-Fruin et al note that both Murray and Church begin their arguments toward agency by describing movement-oriented pleasures. They write:

> Murray's overview — of ways that the agency enabled by games might be combined with the meaningful narratives of fiction — begins with a section titled "The Pleasures of Navigation." Similarly, Church's influential essay "Formal Abstract Design Tools" [5] begins its search for conceptual game design tools with the movement-oriented pleasures of Super Mario 64. However Church's focus is not on movement itself, but rather on how the simple and consistent controls offered for movement, combined with predictable physics, make it easy for players to have intention. [179]

It is not a coincidence that spatial agency through movement is one of the first examples theorists reach for in describing the phenomenon of agency; movement is deeply tied to proprioceptive feedback [67] [174] [178], and this makes it very easy to immediately feel and recognize our agency over our spatial presence.

In describing how we understand our presence in a virtual world, Ryan argues that we experience the feeling of immersion in virtual space largely because of the proprioceptive feedback we receive from our own movement:

> The ultimate test of the material existence of things is our ability to perceive them under many angles, to manipulate them, and to feel their resistance. When my actual body cannot walk around an object or grab and lift it, it is the knowledge that my virtual body could do so that gives me a sense of the object's shape, volume, and materiality. Whether actual or virtual, objects are thus present to me because my actual or virtual body can interact with them. [...] Perspective creates a similar effect by suggesting that the depicted objects have a hidden side that could be inspected by a mobile body. [150]

We experience our own presence not just by taking in what our five senses detect, but rather we experience our own presence through the feedback loop between our own movement and the response we receive from the world in reaction to the movement. The kinds of things we might describe as spatial agency—the ability to move through an environment or the ability to change physical aspects of the world—are easy to recognize as physical manifestations of our will with clear, immediate feedback that can be deeply felt.[4] We experience spatial agency through that feedback we receive as we try to exert our own effort and observe the effects of our effort.

While proprioception is one of the easiest ways to understand our own agency—our ability to manipulate ourselves in space is one of the most basic and fundamental

[4]The idea of spatial agency and its phenomenological aspect was something I previously tried to capture in the idea of *affect* [95]. At the time the idea that agency operates at different levels of abstraction simultaneously was still forming, but I had not yet realized that spatial agency is not the only agency dependent on feedback observation and the updating of the player's mental model. To me, it was clear that affect (spatial agency) felt different from the kinds of agency we experience in narratives, but that was because I did not fully recognize the importance of feedback in players' understanding of their own narrative agency, nor did I make the connection that this feedback was analogous to proprioceptive feedback in space.

ways we understand the world—the process we use to understand our own movement in a space is analogous to how we understand our agency in a systemic sense more broadly. We experience our own agency within a system by observing how the game responds when we take actions. Just as spatial agency requires us to experience feedback in the form of audio/visual sensory changes to our environment, all other forms of agency in interactive narratives require feedback from the system to demonstrate that our actions have had an effect.

I mentioned Wardrip-Fruin et al's argument that players build and revise a mental model through which they constrain their expectation of intention and mediate that intention through affordances in Section 3.3.2. It is worth reiterating the point to specifically note the role that feedback plays in that process. As players, we come to a game with some conception of what the game is; a mental model has already begun to form before we even begin to play. We adjust our framing and our expectations of what actions we will be able to perform based on this mental model. Once we begin to play and we encounter formal affordances, our mental model may update further. We now form an intention to act, take an action, and note how the game responds to that action. We filter the feedback we see, noting how the game has responded to our intention, and our mental model of the game's underlying system is either reinforced or it changes to incorporate the new information. Based on this new information and potentially new mental model, we form a new intention. This pattern of forming intention, taking action, observing feedback, and updating mental model happens across all forms of agency, from learning how to control a character to wondering whether a complex strategy will lead to victory.

It also explains the thrill of discovering unexpected agency, the moment where we take an action that we do not expect to pay off, but it does. An example from "Agency Reconsidered" is the process of improvisation [193], which conforms to this model brilliantly: an actor has an idea for what the scene is they are acting in, they take an action by saying a line or performing a movement, another actor responds, and based on this response, the first actor either confirms their mental model of the scene or updates it. Agency remains intact as long as the second actor's feedback signals that the first actor's action was received and listened to before the second makes their own move, which might change the scene.

This mental model is something we will return to in different contexts. In the next chapter we will also discuss the mental model updating in the context of the thrill a player gets at the sensation that a system is "listening" to input that the player did not expect it to hear.

If we only know we have had agency based on the system's reaction to that agency, the system's response is clearly an important area of research.

3.3.5 Agency is Dependent on Feedback, Not Consequences

While agency is dependent on feedback in this proprioceptive loop, it does not preclude unexpected consequences to player actions. If a system confirms or denies a player's intention, it might be tempting to reframe this as a player must get whatever they want in order to experience agency. In this section, we will explore why the player does not need to "get whatever they want" for agency to be maintained.

For this argument, separating *choice* from *consequences* becomes a useful distinction. Creative practitioners often separate choice as the moment of decision-making that occurs in the course of a narrative and consequence as the eventual result or payoff of that choice [93] [35]. Additionally, I want to separate *feedback* from consequence; in other words, the message from the system that reveals a state change or recognition of an input is distinct from the eventual result of that choice.

In a lecture on best practices for designing branching narratives, narrative designer Molly Maloney describes her process for designing impactful narrative choices for Telltale Games [93]. She argues "choices and consequences are not the same thing. A choice is in the moment; it's about what's happening right now [...] consequences are about the future." Maloney explains that consequences can be unintended and they need not be obvious when the choice is made. Using the example of a game offering a choice of whether the player should sit next to Max or Amy in a classroom, she explains:

> Consequences are awesome in that they can recontextualize previous choices that I made where in the moment, I might have thought it was a great idea to sit next to Amy, but later when I found out that Amy was a real jerk, I felt bad about that retroactively. That's okay for consequences to do, and in fact consequences make choices even juicier and better.

To extend Maloney's example, consider a game that offers you a choice of whether to sit next to Amy or Max in a classroom. I choose to sit next to Amy, the user interface (UI) indicates that my choice has been understood by the system—perhaps by changing color once I press a button to confirm my selection, and the game shows my character walking over and taking a seat next to her.

Now, no matter what happens after that moment, my agency over that choice has been signalled to me as a player. In Maloney's example, maybe I immediately find out that my choice had consequences, particularly consequences I may not like. And these unsavory consequences do not indicate a lack of agency. The game acknowledged my choice as valid and allowed me to carry it out; it just provided unexpected outcomes after the fact. The fact that an unintended consequence resulted from my choice did not detract from my agency at the moment I made the choice. We can imagine, however, situations in which I might have felt my agency was undermined.

If, for example, I had no framing for why I might want to sit next to one or the other—a lack of actionable formal affordances—I would not be able to adequately form an intention, and might just be choosing one at random. Additionally, if I selected to sit next to Max, but the game forced me to sit next to Amy anyway, I may feel cheated, as if something were offered then taken away. The separation between consequence as the eventual payoff of a choice and feedback as the more immediate acknowledgement of a choice becomes useful in this case.

Here we can break down the choice, feedback, and consequence of this moment as illustrated in Table 3.1.

Though narrative designers do not always mention feedback explicitly as separate from consequence, the distinction is often implicit in their arguments, and feedback is so elementary of a design principle that it is often taken for granted. For example, Maloney also notes that consequence is not a substitute for feedback at the moment a choice is made even though she does not call out the term "feedback" explicitly:

Choice	Sit next to Max or Amy? (I choose Amy)
Feedback	*Input feedback:* the choice I selected changes color before the choice UI disappears from the screen. *Narrative feedback:* The game shows me walking over and sitting next to Amy.
Consequence	In the next scene, I see Amy bullying Max. She expects me to participate.

Table 3.1: The distinction between choice, feedback, and consequence in a narrative choice scenario.

> I'm sure you've played a game where you've made a choice and the game doesn't seem to react to it in the moment, but then later you actually find out that it had a huge difference that you weren't aware of. It can be frustrating if you don't feel like the game is listening to you at the moment you made the choice. If you put down that game because you're like "oh, this game isn't listening[5] to me", you'll never stay. The player won't stick around to find out what that amazing consequence was.

Thus for agency to be upheld, there must be a tight coupling between choice and feedback, but eventual consequence can be almost anything. While feedback is important to maintaining agency and the sense that the game is "listening," consequence does not have any bearing on whether agency is maintained. If a story had no surprises, players would soon lose interest, agency or not; we play interactive narratives in part *for* the surprising consequences of our decisions.

[5]Many industry practitioners speak about the game "listening" or "understanding" as a key component to pleasurable decision-making in games [5] [**?**], [160]. Since they do not directly reference Crawford's model, I have to presume that not all of them are citing him, and are arriving at this metaphor independently. Yet the prevalence of the system "listening" as a metaphor reinforces the idea that it is a key component of the player feeling their agency.

3.3.6 Agency Exists at Different Levels of Abstraction Simultaneously

Sometimes our conversations about agency appear to be referring to wildly different phenomena. The phenomena Tannenbaum and Tannenbaum discuss feels like a radical departure from the bounded systemic understanding proposed by Wardrip-Fruin et al. One way researchers have reconciled this seeming tension is to enumerate various types of agency; the exercise is sometimes a way to draw attention to applications of agency the research community tends to ignore [5] [95], and sometimes it is used as a way to help characterize the nature of agency [5] [150].These enumerations are useful shorthands to refer to agency across different aspects of games, but I do not believe they point to phenomena of agency that function in fundamentally different ways.

McGrenere & Ho [111] point out that affordances are often nested in software design. Following Gaver's idea of "hierarchical affordances" [66], McGrenere and Ho enumerate how affordances might be nested inside of other affordances:

> This case relies on the notion of nested affordances. The button has a clickability affordance, which is specified by a raised-looking push button. But users are not interested in clicking on a button for its own sake; they are interested in invoking some function. It is generally the icon or the label on the button that specifies the function to be invoked. Therefore, button clickability is nested within the affordance of function invokability. This is much the same as we would describe a piano as having an affordance of music playability. Nested within this affordance, the piano keys have the affordance of depressability [111]

Similarly, formal and material affordances might exist as nested within each other. In McGrenere & Ho's example, the formal affordance of button clickability is nested within the more abstract affordance of saving one's work. Thus the user might invoke agency

Level of Abstraction	**Agency Exerted**
Low	Player presses a button and sees her avatar throw a punch as a result.
Medium	Player inputs a complex series of button presses and her character performs a super combo as a result.
High	The player employs a "zoning" strategy, throwing fireballs to keep her opponent at a safe distance and hoping for her to jump so she can capitalize. Her AI opponent slides under the fireball as a result.

Table 3.2: Explanation of the player's experience of agency at multiple levels of abstraction simultaneously while playing *Street Fighter V*.

when she clicks that button to save her work, and in the process that agency is operating at multiple levels of abstraction simultaneously.

From this simple example, we can see that the phenomenon we call agency may exist at many different levels of abstraction.[6] Within a fighting game like *Street Fighter V*, we might talk about agency at the level of pressing a particular button and seeing her avatar throw a punch, or we might talk about it at the more abstract level of intending to execute a complex super combo which requires a series of correct button inputs at specific timing, or we might talk about it in terms of employing a specific more abstract strategy such as "zoning" or "punishing."

[6]Note that Ryan also discusses "levels of interactivity" but by this, she is referring to a something slightly different. For Ryan, choices in an interactive narrative exist in degrees that are "closer" or "further" from the text as distinguished by levels of textuality as understood by narratologists. For Ryan, interactivity at the outermost "layer" gives the player control over the presentation layer—how the story is told. Other layers include the story presentation layer, interactivity within a somewhat-authored story, and narrative created through "real-time story generation." [150]

In contrast, the nested levels of abstraction I present here encompass the multiple levels of abstraction that occur simultaneously when the player exerts agency. Rather than provide a taxonomy of types of agency or interactivity in particularly works with this idea, I instead want to present the idea that when we exert agency, we are operating at potentially multiple levels of abstraction at once

In his dissertation, Michael Mateas identifies that agency can exist at a global level or a local level, such as having agency over a narrative storyline or moment-to-moment spatial agency respectively. These point to the extremes, but I argue that the player actually experiences agency at a range of nested levels simultaneously, as the *Street Fighter V* example above illustrates. Indeed even low, medium, and high levels are a simplification; there are surely more than three levels, but I offer these three to illustrate that a player experiences agency at multiple levels simultaneously through a range of nested affordances.

Agency operating at different levels of abstraction accounts for some of the varied uses of the term that might, on the surface, seem to indicate contradictory meanings. While spatial agency, agency in narrative dilemmas, strategic agency, and protagonism might seem on their surface like fundamentally different phenomena, I argue that these are all actually encapsulated by the same model of agency, they are just operating at different levels of abstraction and employing different affordances to highlight different pleasure-experiences of play.

From this point, I would also argue that certain games afford narrative roleplay at a high level of abstraction. Formal affordances for identification with a character and the player's ability to shape her responses to align with a certain version of her character are a high-level formal affordance that the material realities of the game must support. When a game responds on our behalf in a way that undermines how we wanted to roleplay our character, we experience this as an imbalance between formal and material affordances and experience the frustration I described from *Detroit: Become Human.*

Low-Level Agency	**Mid-Level Agency**	**High-Level Agency**
Spatial Agency	Narrative Choices	Roleplay
Controls & Movement	Character Customizations	Long-Term Strategies
	Next-Move Strategies	

Table 3.3: Classifications of agency at different levels of abstraction.

While it is probably more correct to think of levels of agency as existing along a continuum, for the sake of illustration I have grouped them into broad categorizations of low-level, mid-level, and high-level agency in Table 5.1. In general, when we talk about "agency in games," I think we tend to be focused on agency at a mid-level. When I discuss "narrative-level agency" in later chapters, it is this level that I mean.

In the next section, we will situate different forms of agency within the insights brought forward here.

3.4 Reconciling Varieties of Agency

One goal of our theoretical model of responsiveness is that it should provide a lens that is broad enough to account for models of agency as we currently understand them while still being specific enough to add nuance to these understandings. In this section, I want to look at places where various "types" of agency have been enumerated, especially in places where this enumeration was motivated by an argument that areas of game studies discourse were overlooking specific types of agency. I want to reconcile these various "forms" of agency with the understanding of agency provided in the previous sections.

In reviewing these lists, it is sometimes tempting to dismiss one form of agency as not applicable to our task. Take, for example, the idea of protagonistic agency—the ability to imagine ourselves in the shoes of our protagonist and roleplay that character as we choose.

It seems at first glance that this may not be agency at all. It is not, for example, an act of taking decisions within the bounded material affordances a system offers us; it is instead a phenomenon that occurs entirely in our own heads: we imagine ourselves to be a certain character based on the information the game gives us. Still, I think the impulse to situate these different phenomena under the umbrella of agency suggests commonalities. We have not previously had language that could help us reconcile some of the more radical forms of agency, but the model and explorations here have been able to account for some of these seeming differences.

We call a lot of different things "agency". Many of them are operating at different levels of abstraction or on different types of affordances. This section will look at enumerations of agency in an academic context by reframing Ryan's work on categorizations of agency within our model. We will then explore common practitioner concerns by reviewing Ashwell's Bestiary of Player agency through this lens as well.

3.4.1 Categorizations of Narrative Agency: Story and Telling

In her chapter "The Many Forms Of Interactivity", Marie-Laure Ryan builds upon categorizes various forms of "interactivity" [150] which she separates using two dichotomies:

1. *Internal vs external:* whether choices are made from within the narrative world or from outside of it

2. *Ontological vs Exploratory:* power to take action that changes the narrative world vs power to watch/explore but not change

Though Ryan calls these forms of interactivity, as we have seen in Section 3.2, her language in describing these forms primarily focuses on categorizing these forms by the nature of the choices the player makes and the narrative elements these choices affect. All of the forms of interactivity she notes here seem to be expressing agency at the narrative level.

Ryan categorizes internal vs external interactivity as the textual level at which a player is making decisions: either from within the story world or from outside of it. She describes the distinction thusly:

> When the user of an interactive text plays the role of an individuated member of the storyworld, interactivity is internal. The actions of the user correspond to events in the history of the world. On the other hand, when she does not imagine herself as a particular member of the storyworld, or when she controls the storyworld from a godlike perspective, interactivity is external.

And she describes the distinction between exploratory and ontological actions thusly:

> In exploratory interactivity, the user looks at what exists in the storyworld but has no creative power. Her involvement with the storyworld has no lasting consequences. In the ontological variant, her actions create objects that become part of the storyworld or cause events that bring lasting changes. The storyworld evolves as a result of the interaction. While exploratory interactivity precludes world-changing actions, ontological interactivity is a broader domain that allows purely explorative actions.

Ryan then categorizes various examples of works through cross classifications of the dichotomies as depicted in Table 3.4.

	Internal	**External**
Exploratory	*The Manhole*	*afternoon, a story* *Victory Garden* *Patchwork Girl*
Ontological	*Second Life*	*The Sims* *Civilization*

Table 3.4: Marie-Laure Ryan's categorizations of types of interactivity.

Ryan's categorizations here share interesting overlaps with what I have previously called *diegetic agency* and *extra-diegetic agency*, which I described by means of the choices the player makes:

> Diegetic choices are those that a player makes as a character or presence within a story world that affect story, while extra-diegetic choices are those that a reader makes as a removed observer that affect discourse [95]

Where Ryan focuses on the player's relation to the world and whether the effects of her choices enact story changes, the split for me fell to whether the player is making choices

.

In internal interactivity, the player makes decisions that in turn change the narrative at the level of the fabula [20] or story [30]. This overlaps with my idea of diegetic agency completely. However, in external interactivity, the player acts from outside of the narrative world, but the affordances offered might still come from within the narrative world and decisions might still change the narrative's fabula. In contrast, extra-diegetic agency operates on affordances at the level of sjužet [20] or discourse [30].

The distinction reveals interesting insights: the split between diegetic and extra-diegetic agency still serves as a useful shorthand for what Ryan refers to as internal-ontological and external-exploratory interactivity respectively, but it raises a useful distinguishing question: if narrative agency can be diegetic or extra-diegetic, does that label refer to *the position of the player at the time of the choice* or does it refer to *the impacts of her choices*? Ryan's split is useful in separating the player's position within the narrative from what facet of narrative her choices can impact.

However, the split between ontological agency and exploratory agency feels too narrow a dichotomy for the kinds of impacts the player can have. The dichotomy between whether agency affects a narrative's story or its telling, however, opens up different avenues of inquiry. "Exploration" as Ryan uses it is focused on a fixedness of content, and exploration is only one possible way to affect the telling. We might also imagine other types such as agency over point of view or level of detail. The Jew's Daughter [115], for example, provides an interesting example of a work whose agency is clearly discursive, but is not exactly exploratory. Works like *Sleep No More* [133] also feel discursive while still being exploratory.

Thus combining the two approaches seems to provide the clearest and most specific vocabulary for discussing narrative agency by discussing:

1. Whether the player is making choices internal or external to the narrative

2. Whether those choices impact the story or its telling

Within our model, I would argue that all of these choices and the varieties of agency Ryan cites here are happening at narrative levels of abstraction. When deciding whether to classify a work as internal vs external, we ultimately interpret this through the formal and material affordances provided to us. We might ask ourselves if formal and material affordances at the narrative level offer identification with the protagonist. Examples of such formal affordances might include:

- Ability to take actions and move an avatar in a physical space
- Narrative point-of-view
- Emphasis on verbs that change the actions of the protagonist in a text-based game

And when deciding whether to classify a work as ontological or discursive, we would similarly want to examine feedback such as to whether the consequences of actions manifest as story changes or discourse changes. We might look for signals in the feedback as to whether characters, environments, or plot events change in response to our actions.

3.4.2 Bestiary of Player Agency

Sam Kabo Ashwell's article "Bestiary of Player Agency" [5] is often passed around design communities interested in agency in games [158], and it serves as a good starting point for an idea of the "types" of player agency I have heard mentioned among industry practitioners. It is also a useful place to test our model, since many of the forms of agency Ashwell mentions do not seem to be agency as previously understood through purely academic literature. Useful questions to ask while examining these forms of agency include:

1. To what formal and material affordances does the player respond?

2. Can the player form and express intentions through these affordances?
3. At what level of abstraction do these affordances exist?
4. How does the player understand feedback to confirm her agency?

Throughout this chapter I have tried to integrate some of the forms of agency listed in Ashwell's bestiary as part of our discussion. I have, for example, discussed *spatial agency* as a very local form of agency. Similarly, I have alluded to *tactical agency* or the player's ability to solve a problem in a game in her preferred way when discussing levels of abstraction, but it is worth noting that this form of agency is experienced as a balance of formal affordances—strategies the game suggests will work—and material affordances—those the game actually supports.

Big Decisions—the kinds of large, telegraphed choices that we might see in Telltale's *The Walking Dead* series [59]—are examples of the kinds of choices at play that we discussed while talking about feedback versus consequence, though Ashwell notes that these kinds of choices set up expectations for how those choices should be paid off, and they are often a shorthand for players presuming different endings:

> The other thing I'm wary of is that Big Choices are often about choosing endings – in some discussions it's taken for granted that multiple endings are the ultimate determinant of how much agency the player has. That's easy for the author, because they don't have to implement the consequences of choice; and it can be satisfying for the player, who gets control over a major element of the narrative. But increasingly I feel that games are about means more than ends, process rather than closure. This makes me a lot more interested in other kinds of agency. [5]

These other forms of agency are interesting for our discussion too. For example, *reflective agency*—the ability to make a choice that does not change the state of the system in any way but does help the player mentally fill in her own backstory—echoes Tannenbaum & Tannenbaum's conception of commitment to meaning. Here, roleplay is being offered to the player at a very high level; the choice itself is a formal affordance that suggests a player will be able to continue the story while keeping a mental model of her character that the game let her fill in. If the game later forces the player into a narrative situation that undermines that mental model, thus not offering the material affordances necessary to support the roleplay's formal affordances, the player will experience this as her agency being undermined. This form of agency is interestingly internal, but not *necessarily* ontological or discursive.

Negative agency is also an interesting type of agency for our model. This technique sets up the expectation of certain kinds of agency and then intentionally signals its loss. In *Depression Quest* [135], for example, certain choices are presented as stricken through as the player character becomes more depressed and unable to carry out basic tasks. Interestingly, negative agency is often signalled through formal affordances that specifically indicate that a choice is unavailable, which maintains a balance in formal and material affordances toward demonstrating the player's lack of choice to her. Unlike the examples of undermined agency I have provided above, this technique maintains consistent roleplay framing even as it takes choices away from the player—thus it paradoxically undermines agency less than inserting unwanted information that undermines a player's intended roleplay experience.

3.5 Situating Agency Within a Model of Interactivity

In Section 3.2, I framed agency and interactivity as two distinct phenomena with agency serving as one half of interactivity as experienced from the player's side. Over this chapter we have constructed an understanding of agency that brings nuance to the model. Now I want to return to the model to explore more deeply how our understanding of agency fits into the idea of the player listening, thinking, and speaking.

Models, by definition, distill complex phenomena into salient components, the workings of which are metaphorically illustrative of the whole. The LTS model of interactivity builds upon the metaphor of conversation that we have seen across several definitions of interactivity. The metaphor is abstract enough to be applicable to a number of phenomena, but specific enough to bound that endeavor to useful analogous ideas. It is a model that makes intuitive sense and resonates.

3.5.1 Player LTS

Throughout the chapter I have mentioned the process by which the player forms a mental model of the system as she plays. We can also reframe this process in terms of how the player listens, thinks, then speaks in conversation with the system.

1. **Listen:** The player takes in initial information from the game —the box art, the marketing materials, the title screen. Everything from the logo to menu sound effects conveys information to the player that begins to formulate the player's mental model. When she starts the game, she takes in the first formal affordances

offered by the gameplay.

2. **Think:** The player forms a rough mental model of the way the game will operate. She might take a few exploratory actions, but she quickly forms her first intentions.

3. **Speak:** Based on the intentions she has formed, the player tries to take an action. The action she is able to take is bounded by material affordances.

4. **The system issues feedback.**[7]

5. **Listen$_2$:** The player receives feedback from the system. Her action is either confirmed or denied as valid, revealing to her more information about the system's material affordances.

6. **Think$_2$:** Based on the system's reaction, she revises her mental model slightly. She might find other formal affordances. She forms a new strategy/intention.

7. **Speak$_2$:** She then takes another action as she is able to through the system's material affordances.

As with any model, this process is oversimplified. In practice, the forming of intention is a messy endeavor, and players do not always recognize their own intentions this neatly. There is no clean delineation between interpreting the system's feedback and recognizing the affordances it communicates. Sometimes recognition of affordances comes in a

[7] I leave this step intentionally vague for now to foreground the player's experience of agency as one side of the listen-think-speak loop. We will further explore what happens on the system's side in Chapter 7.

discrete "Aha!" moment, but usually a player's mental model of affordances is built up slowly over many iterations of this loop.

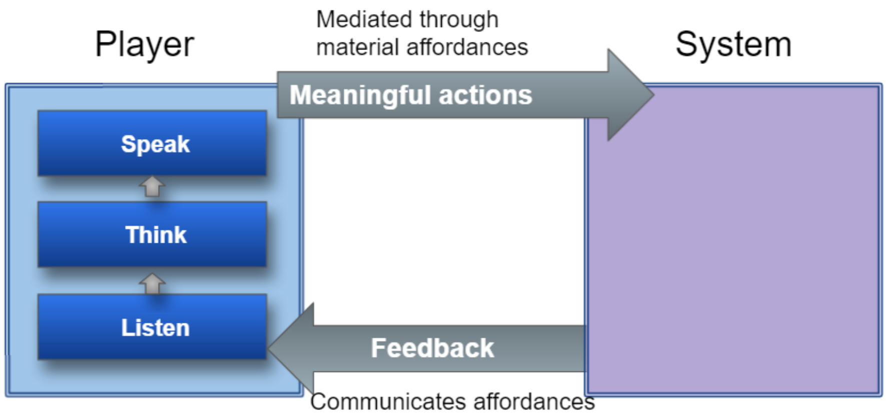

Figure 3.3: Revised model of agency as half of the LTS interactivity loop. Agency is bounded by system affordances, which color the player's expectations and understandings. Her actions are mediated through the material affordances availabe to her and she understands her agency is confirmed by the feedback she receives.

Still, despite the known flaws with metaphorically representing complex phenomena in a model, models like this one offer clarity and understanding in their simplicity. This model, in its simplicity, raises interesting and important questions for us:

1. If agency is dependent on feedback, but exists as the balance of affordances, how does feedback communicate affordance?

2. What do we know about affordances? Are they inherent properties of a system? Do they only matter if the player knows about them?

3. What needs to happen on the system's side for the feedback to communicate affordances appropriately?

In this model, the affordances and feedback serve an important communicative and intermediary role. They are the tools through which the system expresses its state and its workings to the player. Consequently, if a system is to respond to the player that response must come through feedback and affordances.

The next chapter offers a deeper dive into the workings and assumptions at play when we discuss affordances. Though some of the points in this chapter will be reiterated, the specific role of affordances in those processes will be explored in more depth.

Chapter 4

Models of Affordances

In the previous chapter, we saw that player agency arises from the balance of formal and material affordances. If affordances are the primary vectors for the system to communicate its workings to the player, and for the player to express herself to the system, then they are surely of great importance to our conception of responsiveness. I will argue in the next chapter that responsiveness is actually a product of a system's ability to change its affordances and feedback.

Before we interrogate how the system uses affordances to communicate that it has "heard" the player, it will help to have a better understanding of affordances and how they work, especially in relation to agency and player actions. If affordances are deeply intertwined with the player's mental model, how can a designer influence this model in order to both make a system feel responsive and also to make sure that the balance of agency is maintained? After all, if responsiveness is about changing something that is fundamentally striving for balance, doesn't that imply that a responsive system would upset agency? From our experience of responsive games, we know this not to be true; a game's responsiveness generally enhances and deepens a player's feelings of agency by making the player feel heard and acknowledged. The answer to this seeming contradiction lies in the assumptions we make when we model agency as the balance of formal and material affordances.

The concept of affordances exists at various theoretical levels across disciplines like design, HCI, and psychology. McGrenere et al point out that even among the HCI communities, certain assumptions are made about the nature of affordances, and these assumptions necessarily color how we discuss them [111]. If material affordances are properties that exist within the system, and formal affordances are in part mental constructions that the player makes through interpretation, should we think of affordances as inherent properties or as subjective constructions? After all, if affordances are a key tool for a system to communicate its workings to a player, and the system's responsiveness rests in the change of its affordances, the distinction is important.

In this chapter, I will tease apart these different assumptions with an eye toward their impacts on the process of designing for agency. We will examine how affordances foster agency in more detail and will move toward a model that shifts its focus from player action to system response by looking more deeply at how the player's mental model enters the picture. In doing so, this discussion will also serve as background toward the model of responsiveness that we will discuss in the next chapter

4.1 Gibson's Affordances

The term *affordance* as we use it in a design context was first coined by psychologist James J. Gibson [67], who argued for an ecological understanding of perception—that is, an understanding that an animal's perception of its surroundings (and itself) is constructed through interacting with stimuli in its environment rather than, for example, understanding spatial relationships by interpreting visual data in the way a camera would and making reasoned understandings of space from interpreting that data.

Gibson's affordances are framed as properties of environments that "offer" certain uses by an animal.

> The *affordances* of the environment are what it *offers* the animal, what it *provides* or *furnishes*, either for good or ill. The verb to *afford* is found in the dictionary, but the noun *affordance* is not. I have made it up. I mean by it something that refers to both the environment and the animal in a way that no existing term does. It implies the complementarity of the animal and the environment. (original emphasis) [67]

For example water might offer, or afford, drinking, splashing, bathing, and so on; these are the uses of water that are available to a particular animal. Gibson's thesis is that a creature perceives through the affordances provided in an environment. It is not that water is an empty vessel onto which we prescribe a meaning that we cognitively conjure, but rather that we understand and perceive what water is directly, without an intermediate layer of cognitive processing, based on the affordances it offers us.

Though Gibson frequently stresses the contextuality of affordances, in practice the way he talks about them does not always reflect the kind of deeply nuanced reliance on perception that later theorists presume. Instead, Gibson's argument is that we perceive affordances directly rather than interpreting them through a layer of interpreting "senses," and because of this causality—that object suggests affordance—affordances must be at least somewhat inherent. But this is not to say that affordances are *entirely* fixed. Rather, Gibson seems to argue that affordances are somewhat fixed and somewhat contextual: water has different affordances for a fish than for a dog, but barring special cases, water has the same affordances for all dogs, regardless of their previous experience with water or their perception of it:

> Different layouts afford different behaviors for different animals, and different mechanical encounters. The human species in some cultures has the habit of sitting as distinguished from kneeling or squatting. If a surface of support with the four properties is also knee-high above the ground, it affords sitting on. We call it a seat in general, or a stool, bench, chair, and so on, in particular. It may be natural like a ledge or artificial like a couch. It may have various shapes, as long as its functional layout is that of a set. The color and texture of the surface are irrelevant. Knee-high for a child is not the same as knee-high for an adult, so the affordance is relative to the size of the individual. But if a surface is horizontal, flat, extended, rigid, and knee-high relative to a perceiver, it can in fact be sat upon. If it can be discriminated as having just these properties, it should *look* sit-on-able. If it does, the affordance is perceived visually. If the surface properties are seen relative to the body surfaces, the self, they constitute a seat and have meaning. (original emphasis) [67]

Though Gibson specifically argues that affordances are not an inherent property, the description above suggests a universality that implies affordances are somewhat of an inherent property. They may be mediated through context, but they are not dependent on an agent's perception in order to exist. Water's affordances are not inherent properties, but it does not *not* have inherent properties either. And indeed, even Gibson recognizes the tension in his conception:

> An important fact about the affordances of the environment is that they are in a sense objective, real, and physical, unlike values and meanings, which

> are often supposed to be subjective, phenomenal, and mental. But, actually, an affordance is neither an objective property nor a subjective property; or it is both if you like. An affordance cuts across the dichotomy of subjective-objective and helps us to understand its inadequacy. It is equally a fact of the environment and a fact of behavior. It is both physical and psychical, yet neither. And affordance points both ways, to the environment and to the observer.

Gibson's intentions were to present an environmental approach to perception as distinct from other psychological approaches to perception at the time; consequently, his theory centers the environment and its objects more than our conception of affordances today. Even in framing affordances as what is "offered" makes the object doing the offering the active party, an idea which carries interesting implications into its framing as purveyor of agency from system to player in our use. Though an agent receives and interprets what is offered to them—or more precisely, they receive what they can do with a particular object in an environment—ultimately in Gibson's conception, offerings are universal for a cohort of agents such as a species or similar group of animals.

Crucially, for Gibson, the affordances of an object do not change based on the needs of the user. He writes:

> The affordance of something does *not change* as the need of the observer changes. The observer may or may not perceive or attend to the affordance, according to his needs, but the affordance, being invariant, is always there to be perceived. An affordance is not bestowed upon an object by a need of an observer and his act of perceiving it. The object offers what it does because it is *what it is*. To be sure, we define what it is in terms of ecological physics instead of physical physics, and it therefore possesses meaning and value to begin with. But this is meaning and value of a new sort. (original emphasis)

Gibson notes, however, that this does not mean that affordances cannot change at all. For example, events might change affordances, and affordances can change in response to changed context. Similarly, the affordances of an object might change as the observer's position around the environment changes. A cliff, for example, might afford falling off at the top or climbing at the bottom. And we might ascribe positivity or negativity (usefulness or danger, for example) to an object as its affordances change based on our position. Gibson argues that although these judgements are contextual based on the observer's current situation, such judgements are universal rather than subjective to an individual based on interpretive or emotional coloring.

> Note that all these benefits and injuries, these safeties and dangers, these positive and negative affordances are properties of things *taken with reference to an observer* but not properties of the *experiences of the observer*. They are not subjective values; they are not feelings of pleasure or pain added to neutral perceptions. (original emphasis)

Gibson's conception of affordances as inherent offerings of environments provide useful conceptual frameworks for game designers, particularly as we consider material affordances. If an affordance exists regardless of a user's perception of it, then we can design for functionality without worrying about the seemingly-impossible task of predicting what a user will think or perceive. However, designers cannot design products in a vacuum without considering users' perception of those products; they *need* to be able to predict what users will perceive to some degree. It is not enough to offer a functionality if the user never knows the functionality exists.

And finally, the idea of inherent affordances poses a problem for defining what makes a game feel responsive. If we want our system to be able to change what it offers players in response to their actions, the idea of a fixed, inherent affordance presents challenges. Gibson may also suggest inroads to an answer to the problem in his discussion of the affordances offered by people and other animals:

> The richest and most elaborate affordances of the environment are provided by other animals and, for us, other people. [...] They are so different from ordinary objects that infants learn almost immediately to distinguish them from plants and nonliving things. When touched they touch back, when struck they strike back, in short, they *interact* with the observer and with one another. Behavior affords behavior and the whole subject matter of psychology and of the social sciences can be thought of as an elaboration of this basic fact. (original emphasis)

While I do not mean to imply that a game system offers anything near the complexity of interaction offered by people or animals, the idea that interaction might unlock a different kind of recognition and response than other objects or invariants in an environment is intriguing. We recognize people and animals as special because of their capacity for interaction and response to our intentions, and the richness and elaborateness of the affordances offered through interaction certainly deserve closer examination.

4.2 Norman's Affordances

Following Gibson's introduction of the term to the psychology literature, the notion of affordances was largely popularized in the design community [111] through Donald Norman's foundational book *Design of Everyday Things* [126], the first edition of which was released under the title of *Psychology of Everyday Things* [125] (often abbreviated in Norman's own writing to POET) .

In POET, Norman writes:

> The term affordance refers to the perceived and actual properties of the thing, primarily those fundamental properties that determine just how the thing could possibly be used. [...] Affordances provide strong clues to the operations of things.

Like Gibson, Norman was a psychologist, but unlike Gibson, he was largely interested in research questions with more practical applications to the field of product design. Thus his original definitions were less focused on air-tight arguments that could survive academic scrutiny, and instead focused on bringing forth concepts that could serve product and software design fields and the general public. Consequently, he focuses on affordances as aspects of a product that "cry out" for a certain use. Within a book that frames the most important aspects of design as those that align the user's expectations with the reality of how a product functions, his use of *affordance* as the abilities a user *perceives* in a product added a perceptual component to the definition of affordance that was missing in the original definition. The budding HCI community latched on to this conception of affordance, however, and Norman later spent considerable effort to distinguish perception from use [111].

A decade after the book's first publication, Norman laments the misunderstanding of his original use of affordance:

> Alas, yes, the concept has caught on, but not always with complete understanding. My fault: I was really talking about perceived affordances, which are not at all the same as real ones. [...] POET was about "perceived affordance." When I get around to revising POET, I will make a global change, replacing all instances of the word "affordance" with the phrase "perceived affordance." The designer cares more about what actions the user perceives to be possible than what is true. Moreover, affordances, both real and perceived, play very different roles in physical products than they do in the world of screen-based products [124]

Norman argues that a user has a mental model of how a thing works, and that it is important that her mental model aligns with the actual capabilities of the system. This is certainly true, but only tells part of the story.

Norman's view of what real affordances changing means is ambiguous, in part because he is conflating affordances operating at different levels of abstraction. For example,

> In graphical, screen-based interfaces, the designer primarily can control only perceived affordances. The computer system already comes with built-in physical affordances. The computer, with its keyboard, display screen, pointing device, and selection buttons (e.g., mouse buttons) affords pointing, touching, looking, and clicking on every pixel of the screen. Most of this affordance is of little interest for the purpose of the application under design [124]

This passage is confusing because he seems to indicate that real affordances change only when the user is literally prevented from physically clicking when a click would not yield a fruitful result. This is likely a case of speaking about affordances at different levels of abstraction simultaneously and wrestling with the different natures of affordances at different levels.

> Now consider the traditional computer screen where the user can move the cursor to any location on the screen and click the mouse button at any time. In this circumstance, designers sometimes will say that when they put an icon, cursor, or other target on the screen, they have added an "affordance" to the system. This is a misuse of the concept. The affordance exists independently of what is visible on the screen. Those displays are not affordances; they are visual feedback that advertise the affordances: they are the perceived affordances. The difference is important because they are independent design concepts: the affordances, the feedback, and the perceived affordances can all be manipulated independently of one another. [124]

Here Norman reveals that targets on the screen do not qualify as affordances, in the way he uses the term later. Norman effectively sees *real affordances* as the functionality the system actually offers, something more akin to material affordances, while perceived affordances are the functionality the user believes they have. But absent from this model is the formal affordance, which the target on the screen represents. While Norman's conception of affordance here and his distinction between "real" and "perceived" is a departure from the way we have been using the idea of affordances, this split is very useful, as we will see in the next section.

4.3 Reconciling Gibson and Norman

The question of whether an affordance is a perceived property as Norman argues or an inherent one as Gibson suggests, reveals different approaches to the nature of affordances, and the implications of these assumptions have impacts on the discussion of design. The disparity in definitions has led to confusion among the HCI community [111], and deserves further exploration in games studies discourse. Gibson's book argued for a form of perception that was explicitly anti-cognitivist, in a tradition we would now call an ecological approach, where Norman's split between "real" and "perceived" affordances, though argued to be a closer understanding of Gibson's terminology, actually exemplifies the return to a form of cognitive perception from which Gibson was explicitly distancing his notions of perception. Where Gibson is arguing that a subject *perceives* and experiences their reality through their experience of opportunities within their environment, Norman argues that an object includes opportunities that are true and correct (e.g., that a teapot is graspable by the handle), and that there is an additional layer of comprehension for more abstract objects that involves whether such an affordance is obvious enough to be perceived.

Gibson's work, though grounded in what we would call an ecological approach to perception today and largely interpreted as a departure from cognitive modeling altogether, focuses on distancing from a cognitive approach to *visual* perception specifically. His examples involve very primitive types of perception and equally primitive actions within an environment—breathing, climbing, running, sitting. His efforts to debunk a cognitivist approach to perception are focused on how humans experience visual input; he is concerned with how we see objects and argues that our vision depends on our operation as agents in an environment. How we build mental models when operating computers or performing more abstract tasks involving language or culture is not at odds with how we see rocks and understand them as throwable.

Norman's examples, by contrast, largely consist of man-made objects with socially agreed upon functions in abstract situations. The two are largely dealing with different domains, and thus their insights are not mutually exclusive, and indeed games studies conversations around affordances benefit from recognizing both approaches as valid [26]. It is possible that the examples of the HCI community might gravitate toward Norman because many of their products are similarly abstract and environments involve agreed-upon constructs, but elements of Gibson's approach are still applicable.

Both models offer important lessons. Understanding the cognitivist approach to perception is important for understanding how the player will update her mental model, which is important to understanding agency. Similarly, understanding what sorts of things will inherently suggest certain actions—either through inherent affordance, mapping, or convention—provides a toolbox for designers, and moreover, the idea of a fixed affordance is useful for how we will be conceiving of responsiveness changing affordances.

4.4 Formal and Material Affordances

The idea of two different models of affordances—one in which perception and recognition play key roles, and one in which offerings are inherent regardless of perception—find echoes in Mateas's conception of formal and material affordances that we discussed in the last chapter.

Mateas argues that the experienced phenomenon of agency might be abstractly modeled as occurring at the point where formal and material affordances are in balance. Mateas's conception of material affordances echo Gibson's original conception, and Mateas acknowledges this influence [99]; material affordances are the systemic features available to the player, and they exist independent from the player's perception or understanding of them. As in Gibson, the player recognizes her "reality" within the game and begins to strategize through the lenses the material affordances provide.

Formal affordances, by contrast, correspond with Norman's conception: they are aspects that "cry out" for particular interaction and help the player form the intentions that they will express through the material affordances the system provides.

An important contextual note is that Mateas's first introduction of formal and material constraints largely discusses a very particular structure of interactive drama, one in which a player is acting as a character within an interactive drama. In this context, formal constraints influence and motivate what a player wants to do within a given interactive drama, and are largely a product of the constraints imposed by a narrative situation and our expectations of stories of the particular type into which our interactive drama is situated.

As the idea of agency and its relation to formal and material affordances broadens in "Agency, Reconsidered", however, the idea of agency as the balance of formal and material affordances is applied to a broader range of interactive experiences including the chatbot *Eliza* [195] and *Super Mario Bros.* [137], a game whose player motivations more arguably arise from mechanical motivators and feedback than from a player's narrative expectations of plumbers in mushroom worlds invaded by tyrant turtles. To echo this shift, in Mateas's later lectures [102], the definition of formal affordances moves away from representing the possibility space we would expect to have in a particular *narrative* situation and instead begins to represent the possibility space we expect to have and the actions we are motivated to take by the representative layer the game offers the player.

This definition allows for application of the agency-affordances model to a wider range of experiences, but it makes the idea of a formal affordance more tricky to prescriptively pin down. Formal affordance conventions vary by genre, by intended pleasure experience, and have changed substantially over time. Additionally, the kinds of things we might call formal affordances at a very low level of abstraction look very different from the kinds of things we would call formal affordances at a higher level of abstraction. Let us return to *Super Mario Bros.*, a game whose core pleasure experience is not primarily narrative or roleplay oriented and examine some of its formal affordances for a clearer idea.

The box art for Super Mario Bros.sets our first expectations by offering the framing narrative and asking "Do you have what it takes to save the mushroom princess?" (see Figure 4.1. The goal to save the princess is a formal affordance at a very high level and answers the question of "why am I doing this?". When the game opens, Mario is positioned all the way to the left of the screen, facing right, with much of the world to the right, the combination of which is a formal affordance inviting us to move to the right. The first thing the player encounters is a bright, flashing block with a question-mark on it (formal affordance advertising power-ups), but as the player approaches, a mushroom with angry eyebrows moves to block it (formal affordance advertising enemies). Having only two buttons (material constraint), the player quickly learns she must jump to avoid the first enemy—a mushroom with thick angry eyebrows—and when she does, positioning of the blocks makes her very likely to land on it. When she lands on it, a satisfying sound is made and numbers appear on the screen to indicate her score has increased (feedback). The player has now learned the two core mechanics (move to the right and jump on enemies to defeat them) [36]. She may learn that pits and lava kill her purely through a process of trial and error, but the game tries to advertise these aspects of its system to the player; those advertisements are formal affordances.

Though none of the elements being represented here are actually grounded in narrative, (e.g., the player is not motivated to move to the right because Mario's character is established to only move right) the representational elements of the game

Figure 4.1: Box art for *Super Mario Bros.*. The box art and other marketing materials offer formal affordances to start the process of building the player's mental model of the game before she even begins playing.

including Mario's position and facing, the placement of powerups, the visual distinction of the power-ups from other blocks, and other aspects of the level design function as affordances that motivate and inform the player in a way analogous to the narratively-grounded formal affordances of interactive drama.

In other words, formal constraints as a definition begin to shift away from a solely narrative constraint and encompass a broader range of constraints designed to signal to the player what she should *want* to do. From a practical standpoint, formal affordances encompass the tricks designers use to try to communicate systems and state. They are the tools we use to motivate players. Player motivation is a tricky thing. Motivation might arise from a myriad of sources, some of which can be difficult to predict and many of which are deeply personal. A designer's job is to try to shepherd those motivations through affordances.

I also want to emphasize that although our conception of formal affordance has broadened, the definition of agency as the balance between formal and material affordances does not change. If we recontextualize formal affordances away from narrative-specific focuses and frame them as the formal elements that advertise certain functionalities of the system to the player and motivate her to take various actions, agency is achieved when a balance exists between what a player wants to do based on the opportunities signaled to her (broadened formal affordance) and what a player systemically *can* do (material affordance).

4.5 Affordances and the Mental Model

If affordances arise from the mechanical and narrative expectations of the player in a given situation, both of which are naturally influenced by the player's past experience with narratives and games. As we saw in the previous section, formal affordances influence what a player expects to be able to do, which in turn gives rise to a motivation or intention. How then do players construct these intentions?

Many of these intentions come from the intersection of formal affordances and *conventions*[126], socially agreed-upon norms that suggest affordances when none are apparent. While our understanding of an object's affordances might arise from innate understanding of an object's properties (as in Gibson), often it is the result of our perception of conventions that offers an initial mental model. *Mapping*, the process of of assigning (often visually) analogous relationships between an object's control and its functions—such as turning a car steering wheel to the right to turn a car to the right—can also help us mentally model an object when no affordances immediately present themselves [125].

Under Norman's conceptions, a user will generally approach an object with a rough mental model of what it is "for" and how it is used. Using this model, she forms an intention and attempt to carry out our intention, and in the process we receive feedback from the object as to whether our intention was successful or not. If our action was successful, our mental model is confirmed, and if the action was not successful, we update our mental model, even if only to note that a certain action is not valid. We then begin looking for the next way to interact with the object, until our goal is reached. Over the course of interacting with similar objects or mappings, we build up a body of knowledge that we use to approach new objects, and this body of knowledge we adopt in that process is called "convention." Norman notes that convention is arbitrary and learned; it may intuitively make sense but it could be any other thing and still function and make sense; it is not inherent. [124]

Convention exists in game design too. Players approach any new game with the learnings from previous games they have played in that genre [193]. Conventions around common control schemes, common types of objectives, abilities, and so on color the experience a player expects to have. Some of these conventions are even explicitly codified by game platforms. For example, Microsoft, Sony, and Nintendo will not certify a game for release on their console platforms if the menu buttons for confirming a choice or going back on a menu screen do not conform to the conventional buttons on the controller that players will expect.

We might even argue that genre itself is largely a construct designed, in part, to corral convention of specific types of pleasure-experiences so that players know what mental models to bring to a game [1]. We know that a first-person shooter will likely involve certain elements conventional to that genre, thus if we start a game knowing it is a first-person shooter, we have already started to construct a mental model around how the game "works" which will color our expectations. [193]

In addition to relying on convention to begin building a mental model around a game's systems, players discover which affordances are available to them by taking action within a game space and observing how the system responds to that action [88]. According to Linderoth, game environments offer *exploratory* actions, which players perform in order to learn more about the affordances available to them, and *performatory* actions, which players perform in order to try to realize their intentions. Once the player has taken a performatory action, her intention is either confirmed or subverted, thus affirming or changing her mental model, respectively. Interesting to note in the discussion of conventions, intentions, and mapping is that mapping is only apparent after a user has taken an exploratory action upon the object to discover the map—e.g., a user might not realize pressing the left button on a directional pad will mover her character to the left until she blindly tries it.

Linderoth's perspective and integration of ecological understandings of affor-

[1] See also Bernstein & Greco on the economics of genre in new media [17]

dances are critical to adding nuance to our understanding of the mental model, but also to more correctly describe the LTS loop. Not *every* action the player takes is an informed, conscious decision that she decides upon and then enacts. The lines between listening, thinking, speaking, and acting are not perfectly clear and discrete. Rather, it is probably more helpful to think of the mental model as constantly in flux and updating. The system offers a conversational counterpart that provides input to the player's mental model via feedback and shapes the player's output via affordances is a useful one.

The process of updating the mental model is constant, not only in games, but in other narrative experiences as well. We want narrative experiences that surprise us, and we find stories where all aspects of the story and discourse contain no surprises to be tedious. Even in situations where we know the outcome of a story, we often experience it in the moment as if we do not [150] , so surprise and immersion remain intact.

For games, I would argue that a lot of pleasure lives in the gap between the player's mental model and the system's true nature. The constant loop of a player realizing the outcomes of her actions—exploratory or performatory—and the process of updating that loop is what keeps motivation and keeps the player motivated to continue moving forward. The gap is the carrot.[2] As soon as our mental model stops expanding, a game becomes "solved." When a game is solved, a player so deeply understands the remaining possibility space that the conclusion becomes foregone, and most people lose

[2]This idea is supported by Csikszentmihalyi's work on flow state [37] , which I have previously likened to mechanical immersion [95]. In Csikszentmihalyi's conception, we need a little bit of "reach" to maintain flow, and it is the fact that our model needs updating a little every time that keeps us interested.

interest at this point. Norman illustrates this point with the fact that few adults play *Tic Tac Toe* among themselves, because the game is solved immediately. As adults, our mental model of *Tic Tac Toe* overlaps so completely with the possibility space of *Tic Tac Toe* that the conclusion is foregone. We have mastered *Tic Tac Toe* too completely for it to be interesting [126].

Chess provides another interesting example for this idea: if a game ceases to be interesting, once a sufficient mental model is built up, why is *Chess* still interesting to *Chess* masters? Surely their mental model of the game is quite advanced? *Chess* is not actually a solved game; the large possibility space means that the number of optimal moves in any game is large enough that the conclusion cannot be drawn until a few moves from the end. Thus each player's mental model of the moves that will be available to them in the future, not to mention which of those moves will be strategically advantageous, is in constant flux until late in the game. But usually a conclusion is reached before the actual end of the game, and it is common to forfeit once that point becomes apparent. In other words, the affordances are always shifting, and the moment the mental model perfectly aligns with the remaining possibility space, the game ceases to be interesting.

As our mental model "solves" each level of abstraction, we seem to move to the next level up to try to master the game at that level. Sudnow's journey in *Pilgrim in the Microworld* [174] echoes his experience as a jazz pianist in *Ways of the Hand* [175]—he practices the lower-level affordances: movement, openings, etc until they become muscle memory–so internalized that the he does not need to think about them so that he can turn his attention to higher-level concerns. Similarly, a player learns affordances at a lower level of abstraction—awkwardly and consciously at first. As she improves, that lower level becomes internalized to the point that she performs those actions unconsciously and can turn her attention to higher-level strategies. These higher level strategies offer their own meta-affordances composed of the smaller affordances nested in the levels of abstractions below.

4.6 Perceived Affordances, Real Affordances, and Feedback

In an attempt to reconcile the shortcomings of both the ecological approach to affordances exemplified by Gibson and Linderoth and the cognitivist approach to understanding object design implicit in Norman, Rogelio Cardona-Rivera and Michael Young call for a cognitivist approach to affordances that extend and clarify Norman's distinction between real and perceived affordances while also providing some additional context for how game designers might approach affordances toward their aesthetic goals.

> The primary thesis that our cognitivist theory contends is that, in the context of game design, designers should primarily focus on what players perceive they can do, as opposed to what players can actually do in an interactive virtual environment. While this might seem trivially true, it has a subtle implication: if a virtual environment does not support an actual affordance, but never presents feedback to elicit that course of action, player agency may remain unaffected (Wardrip-Fruin, et al. 2009) [26]

In enumerating their cognitivist approach, the authors reintroduce concepts from Norman [126] and frame them in the context of game design. They define the components of their theory thus:

> 1. *Real Affordances* – what is actually possible in an interactive virtual environment; these affordances are actionable by the player. Borrowing the terminology presented by Wardrip-Fruin et al. (2009), these affordances lead to actions that are "supported by an underlying computational model." This is the first of two entities that are actually manipulable by the game's designers and developers.
> 2. *Perceived Affordances* – what players perceive to be possible. Perceived affordances do not necessarily correspond to real affordances. The perceived course of action must conform to what a player believes is possible (a reasonable action) and must be consistent for the player within the game's context. A player's beliefs can be informed by what a player has experienced in the game (her perception and attention), as well as be guided by what similar games have typically expected from her in analogous situations (her memory and analogical thinking skills). Similarly, a player's sense of consistency of actions within a specific game context can be informed by what actions have been available in similar games; by similar, we mean games that can be considered to be within the same genre (Miller 1984).
> 3. *Feedback* – this is perceptual information used in the game to advertise the real affordance in the hopes of eliciting an accurate perceived affordance. This is the second of two entities that are actually manipulable by the game's designers and developers. [26]

In their model, as in Norman, a designer is trying to bring a player's understanding of their available affordances in line with the "real affordances" of the system through the use of feedback. In their discussion of how their theories relate to existing theories of affordances, the authors describe the correspondence to formal and material affordances thus:

> Mateas (2001) argues that agency arises as a result of the balancing of material and formal affordance. Material affordances (opportunities for action presented by the game to the player) are delivered as feedback to communicate a game's real affordance. Formal affordances (motivation to pursue particular material affordances) are also delivered as feedback to communicate a game's real affordance. Note that, while both material and formal affordances take the form of feedback in our framework, their purpose is different. Material affordance feedback targets perception and attention (e.g.,

> through the use of lighting (El-Nasr, et al. 2009)), whereas formal affordance feedback targets problemsolving and decision-making (e.g., through the use of influence (Roberts, et al. 2009)). (pg 4)

Implicit in this argument is the idea that their use of the term affordance is operating at different philosophical levels than Mateas, particularly in the case of formal affordances. Where Mateas is concerned with an abstract idea of affordance, Cardona-Rivera and Young are more interested in tangible offerings of a system. In their model, both elements of Mateas's model are signalled through "feedback". Though they describe formal and material affordances as "taking the form" of feedback, it seems more correct to say that formal and material affordances are abstractions, constructed in the mind of the player and through the combination of complex systemic processes respectively, both of which are *communicated through* feedback. Feedback is the language the system has to communicate.

It is tempting to simplify "real affordances" to "system features" but we must remember that affordances operate at multiple levels of abstraction. At the lowest level of abstraction, yes we are trying to help players understand what actions the system will understand–what words a parser will recognize to what button combinations make an avatar run and jump. These are affordances at the most basic levels. However, this theory can be applied to higher levels of abstraction too (such as when a player anticipates certain enemy behaviors in a speed run, or higher still such as when a player recognizes the strategy an AI opponent is employing in *Starcraft II* [42]). And at slightly higher levels of abstraction, we understand that the player's ability to bring their understanding of the game's affordances in line with what the system offers extends beyond mere "features" into things like "successful narrative actions" or "viable complex tactical strategies."

As a designer, the use of feedback to bring perceived affordances (both formal and material) in line with real affordances aligns with how we try to signal available actions to players. Such a design strategy is evident in the quotes we have seen from practitioners so far. This conception of affordances, though operating at a different philosophical level from Mateas, gives us a way to talk about the changing of affordances as separate from *the communication of that change.*

In building our model for system responsiveness, the combination of formal and material affordances into "real" affordances is not specific enough for our purposes, however the separation between affordances and the feedback that communicates them provides a means to help explain not only how systems enact responsiveness but also how they communicate it.

4.7 Agency as a Loop

With these different conceptions of affordances, let us return to the LTS model. The idea that agency is achieved when what the player is *motivated to do* as motivated by formal affordance is balanced by what the player *can do* as mediated by material affordance makes sense as a general abstraction, but the extended vocabulary offered by other models of affordances, aided by a deeper understanding of the assumptions therein, allows us to fill in additional details and add information to the model. Let us look at what the LTS model might look like with additional details.

In the model, a player takes actions as available through the system's material affordances. Those actions are then interpreted by the system, the system updates, and outputs feedback to the player. This feedback, we previously argued, offers suggestions from which formal affordances arise. As long as the formal affordances suggested to the player by the system remain in balance with the material affordances the system can understand, agency is maintained.

But as we saw in the previous section, the system does not communicate purely in formal affordances; a translation process between the system's feedback and the player's understanding of available actions and motivations happens at the think stage. Agency under this model of affordances starts to look more like a loop that looks like this:

1. Player takes exploratory action based on perceived material affordances (speak)
2. Action is recognized by the system (listen)
3. System updates (think)
4. System outputs feedback (speak)
5. Player receives feedback, interprets formal affordances (listen)
6. Feedback confirms/disproves mental model, perceived material affordances update, player forms new intention (think)

7. Player takes next action based on new model (speak)

Recall from our previous chapter that the player understands her own agency through feedback. Feedback either confirms the player's intended action as valid, or denies it as invalid, expanding the player's mental model and updating her expectations. As long as the communication loop is maintained, and the system is not denying the player the ability to be heard and her intentions carried out—even if to her own demise—agency is maintained; formal affordances remain in balance with material affordances.

However, the agency loop fails if communication breaks down along the loop. The player will feel the disruptions at the point of feedback, or more specifically, in the lack of feedback appropriate to her intention. Perhaps the player presses a button and feedback does not come at all; her action was not acknowledged. Or perhaps the feedback explicitly indicates that the system has not understood the intention (e.g., the classic mechanical interactive fiction parser response of "I don't understand that verb."). Or perhaps the feedback does confirm the player's previous action as valid but does not deliver sufficient formal affordances for the loop to continue, such as the feedback offering a "blind choice" for which the player does not have sufficient framing). The disruption of the loop is a disruption in agency.

Supposing the feedback does confirm the player's intention as valid, this recognition might come in two forms:

1. The feedback might recognize and carry out the player's will, but add no new information or additional affordances or

2. The feedback might recognize the player's will, but add or change affordances, offering the player a next action to take or strategy to act upon.

This second case is interesting, and opens the door for an exploration into changing affordances, how systems respond, and how that response reinforces and maintains a player's agency.

Chapter 5

Responsive Systems

5.1 Introducing Responsiveness

One of the most exciting promises of our medium, is the feeling that a system has heard our intent and has responded. Agency is the ability to take meaningful action—form an intent, take an action, and have the system carry out that action—but in order to do that, the system must *understand* us. At its most basic, the system does what we request; but at its most brilliant, it collaborates with us, offers a response that demonstrates that it has recognized our input and is taking it into consideration as it offers us new possibilities.

Responsiveness, like agency, is an aesthetic experience that the player *feels*. Responsiveness is the property that makes us exclaim, "Wow, I didn't expect that to work!" or "Wow, I didn't realize how consequential my actions were!" It is the clever response from an interactive fiction parser to a verb the player didn't expect it to recognize; it's the thrilled unease a player feels the first time they encounter the promise that "Clementine will remember that" [59]; it's the excited realization that a narrator is reacting to your moves in *Portal* [187] or *Bastion* [58] or *The Stanley Parable* [22]; it's the delight an audience-member feels when their shouted word is chosen by an on-stage improv group for use in their skit. Responsiveness describes the degree to which a system can meaningfully respond to the input actions of a user. It is fostered by, supports, and reinforces our feelings of agency.

Among industry practitioners, the analogous idea of "reactivity" encompasses "the ability of the game to react to player input." [45]. Though the term is used informally among industry narrative designers, its meaning or implications have never been properly explored. Developers talk about how "reactive" a game is in a similar way to how early practitioners talked about "interactivity": the community presumes a general sense of shared understanding, but the exact meaning is nebulous, and few efforts have been made to lock down the nature of the term and the processes that foster it.

"Reactivity is incredibly important for agency; the more a player feels like their choices matter in a story, the more engaged they are, and the more agentic they'll feel." [45]. This idea that "choices matter" tends to colloquially mean that a system has responded to a player's choice in a way that indicates counterfactuality with another response. In other words, something the system models has changed because of player choice, and selection of a different choice would have resulted in a different outcome. This is a very specific and limited idea of what "matters" about choices (see [6], [47],

[78], [107]). One might argue that the prevalence of the term—such as its use as a categorizational tag applied by users on Steam and other popular game distribution platforms—indicates that the player's ability to affect a game's model is core to certain narrative pleasure-experiences [127] .[1] But is it actually the player's ability to change the model that matters in these cases, or her recognition or understanding of that ability [189]? Is it something else entirely? Understanding responsiveness (or reactivity) is key to understanding how our actions have affected the game's model.

When trying to apply a term to the phenomenon of the system responding to a player's action in a way that indicates it has "heard" the player, there are several we could choose from. "Reactive" and "reactivity" are currently in use within industry. "Responsive design" is also in use. For our purposes, I have chosen "responsiveness" because it feels more intentional than merely reacting, and feels like an appropriate term to keep with the metaphor of interactivity as a conversation that we have built upon throughout this text. Participants in a conversation might react, but that connotes that they do so with little thought or intention. *Responding* in a conversation, on the other hand, implies deliberation, or at the very least comprehension. If our systems are to indicate that they are listening, thinking, and speaking in reply to what a player has done, the term *responsiveness* feels more appropriate than *reactivity* to me.

[1]While the ability to counterfactually influence a narrative model might be core to certain narrative pleasure-experiences, the games tagged with "choices matter" show wild inconsistencies in how the term is applied. In general, games for which the community has decided "choices matter" do not always correspond to how much ability the player has to affect the game model, despite that being the commonly accepted colloquial meaning. Instead, it often corresponds to things like how often choices are made, how central choice is to the core mechanic, and how much the player is able to recognize and understand the impacts of their choice when making it. Though inconsistencies in the use of the term might make it easy to dismiss any commonality to such experiences, they instead point to a need for better language around the distinctions between choice, agency, and responsiveness.

I have suggested previously that responsiveness is the system's counterpart to player agency – it is how the system communicates that it is upholding its side of the listen-think-speak model. Responsiveness as a concept still centers the player; after all, the most important aspect of a system feeling responsive is the player's sense that she has been *responded to* — that she has been heard, understood, and acknowledged. In examining how a system comes to feel responsive to a player, we will examine the listen-think-speak loop from the system's side. We will examine how the process of "mental modeling" as done by a system shares components and deviates from those of the player, how the system may increase responsiveness through improvements in the listen-think-speak loop, and how responsiveness might exist at multiple levels of granularity and abstraction, just like agency.

Additionally, we will examine how one of the most important elements of a player feeling she has been heard is ensuring that the system listens/thinks/speaks in a way that communicates that it has done these things with the player, thus reinforcing rather than undermining agency. The tools that the system has to respond to the player –its affordances–then become central to the system communicating it has understood and is working with the player.

I argue that **a system's responsiveness is the degree to which the system alters its feedback and affordances as a result of the player's inputs.**

Capturing what is ultimately a felt experience in a theoretical model is difficult, but ultimately we want to capture the fact that players feel heard. In this chapter we will examine how affordances and feedback changing contributes to this feeling. We will look at how systems listen, think, and speak and how affordances and feedback function within this loop to create a system response.

5.2 A Model of Changing Affordances

As we saw in Chapter 4, affordances and feedback are the tools available for the system to communicate its state and processes to the player. If the system is to communicate that it has heard and understood the player, this must be done through changes to its affordances and feedback. This section will explore the implications of those changes.

Intuitively, in a conversation in which both parties are engaging in the listen-think-speak loop, it is not enough for Party A to issue commands to Party B; Party A must receive an answer that indicates understanding and that Party B is holding up their end of the conversation. The absence of such changes amounts to Party A shouting at a wall. In order for something that feels like a conversation, we need the system to not just understand and carry out our will, but to respond. As we have seen, the system's ability to communicate with players is mediated through its affordances and communicated through its feedback, so in order to communicate that a system has adapted in some way, its feedback and affordances *must* change in response to the player. Let me demonstrate the point by way of an example: let us consider a light switch.

First, let us consider a traditional light switch. The user enters a room and sees a switch on the wall that resembles other lightswitches she has seen in the past. Its form and shape offer formal affordances, which the user interprets to form the understanding that flicking the lightswitch up (material affordance) will turn the light on. She does so. The room lights up; feedback tells the user that her action was successful and her will was carried out. At a very local level, agency has been achieved. But the affordances do not change. No matter how many times one flicks the switch, the formal and material affordances of the light switch will remain constant. And similarly the feedback, though changing with the light's state, will not change in its character; the room will light up when the switch is on and will darken when the switch is off.

Alternatively, consider a smart home voice agent that controls the lights instead of a switch. The user says to the agent "Silli, turn on the office lights." The agent turns on the lights, and then the agent asks "You sent the same request yesterday at this time; would you like me to add this to a daily routine?" Here the agent is responding. The agent has heard the user's request, carried it out (as demonstrated through the feedback of the lights coming on), then additionally done some comparison processing to recognize a pattern, and offered additional feedback in the form of a question. The question indicates that the agent heard and understood the user's initial request and additionally has changed its formal affordances (outputting a new line) which reveals two new material affordances to the user: (1) the ability to establish routines and (2) the ability for the system itself to detect the user's patterns. The system's adaptation is understood as a change in affordances (either formal, material, or both), and communicated through feedback.

The second example is clearly more responsive than the first, and that responsiveness also reinforces and supports agency. In the case of the traditional light switch, the user has one key choice: whether she wants the light on or off. She can exert agency at this very local level, and the formal and material affordances of the switch support her agency in this choice. Feedback confirms that her choice was successful and agency was upheld. But her range of options are limited, and nothing about the lightswitch feels as if it is adapting to her. The second example, by contrast, the feedback the system offers both confirms the user's intended action and presents new formal and material affordances for her to consider. Presuming the system can deliver on its offer and can add the operation to a routine, the new formal and material affordances remain in balance, so when the user responds with, "Sure, Silli, that would be nice." her agency is maintained and confirmed as soon as the agent responds with "Got it. The routine is set."

But could we not argue that the simple lightswitch is still responsive? After all, the expected feedback does come and the user's agency, local as it is, remains upheld. And the more sophisticated agent is arguably only presenting a second choice after the first—is that really so different?

To answer these questions, let us explore changes in feedback, formal affordances, and material affordances in more depth.

5.2.1 Changing Feedback

As we saw over previous chapters, all changes to the system are communicated through some form of feedback. For agency to be maintained, a change of state must be communicated through feedback.

Feedback, as we saw in the lightbulb example, might communicate a change of state without changing affordances, or it might communicate both a change of state and a change in affordances. To say that feedback has changed (but affordances have not) implies that the game is communicating state changes without necessarily communicating changes to its computational model or processes. Feedback does not need to communicate changed affordances in order to uphold agency, as the lightswitch example shows, but without changes to affordances, an interaction will soon feel shallow. After all, tools and objects change feedback without changing affordances; those interactions feel much more like *uses* than *conversations*, and our lightbulb example echoes this.

till, changed feedback alone can offer its own pleasures, and indeed many of the things we might call responsive are actually tight, well-tuned feedback. For example, when we talk about responsive controls, we normally mean that the feedback we feel in the player avatar's movement is tightly-coupled with the player's inputs [178].

Bastion [58] and *The Stanley Parable* [22] are both games that rely on responsive narrative feedback to player actions in the form of voice over. In both cases, a narrator recounts the player's actions through voice-over in third person. As the player makes a move in the world—through hack-and-slash combat and exploration of a top-down fantasy world in that case of Bastion and through first-person exploration of a surrealist office space in *The Stanley Parable*—a voiced over narrator makes comments like "The kid fired into the crowd" or "Stanley took a detour into the office closet" respectively. In both cases, the effect is the feeling that the narration is responding to the player's moves, even though the narration effectively serves as feedback that parrots the player's own actions back to them. We will return to these games to compare them in more depth in Chapter 6. For now, I want to highlight that feedback alone can be extremely effective in fostering agency and the sense that a game is listening, especially when that feedback comes in unexpected depth or quantity.

At the outset of the chapter we noted that responsiveness is the ability to make the player feel heard. Feedback can make the player feel *heard* without necessarily changing the computational model in response to her, but as we saw in the light switch example, feedback that offers a change in affordances creates a sense of responsiveness that is different in character from feedback alone.

5.2.2 Changing Formal Affordances

Having seen that the system's ability to change its feedback leads to the sense that the player has been heard, let us more closely examine the implications of formal affordances changing.

In the context of interactive drama, a change in formal affordances might come in the form of a new affordance being introduced. For example, a new character might enter the scene with a problem. Or an NPC might mention a new destination the player should visit, or reveal new information about the murder the player is investigating. The introduction of new formal affordances, used in this way, are often a way to move the story along and motivate the player toward her next goal.

The key to a formal affordance feeling responsive is that it must feel that it has come *as a result of the player's action.*

To return to our light switch example, imagine if instead of the agent asking if you would like to add the lights to a routine as the result of a user's action, it asked this unprompted while the user was sitting at her desk working. Here the agent would still be offering a new formal affordance, but instead of it feeling like a responsive reply to the user's request, it would feel more like a random advertisement of its own workings. This approach is often how games currently handle formal affordance changes, which is also why many game narratives feel like transactional exchanges designed to dress up a set of rules rather than conversations.

Consider ways we might alter each of the formal affordance changes mentioned above to happen as a result of a player action. Instead of a character entering the room with a problem unprompted, suppose she entered a few seconds after the player turned on the light or after she fired a shot. Suppose the NPC only mentions the neighboring town's wolf problem after seeing the player has a wolf-skin sarong equipped. The context surrounding a system's presentation of a new formal affordance matters. It might be easy to dismiss such examples as indicative of an intractable amount of branching, but I would argue that expensive branches of content is not necessary to introduce formal affordances as a result of player actions; often clever writing and design can circumvent expensive branching—in addition to the technical solutions we will explore in later chapters.

Changes to formal affordances are one of the key ways we recognize narrative progression as players. In a story, we expect to have different actions to us at the end of the story when we are the powerful, changed protagonist confronting the villain than when we are the wide-eyed novice at the beginning. What makes sense to do next, thus what the game should draw us to do next, should change over the course of a well-formed narrative; after all, for many definitions, change of character, circumstance, and/or understanding is the basis of well-formed narrative arcs [112].

Games offering a more cinematic experience, such as Telltale games, do this well. Though many of the player's actions in each scene or level often amount to having conversations or collecting and using items, at a higher level of abstraction, player actions cause relationship changes and depend on subtle nuances signaled by particular dialogue choices. The player is actually strategizing about the impacts of her choices—how certain conversational moves might be received, for example—almost all of which are necessarily different choices at the beginning and end of the game once relationships are established, narrative context changes the impact of certain decisions, and the nature of the things being decided are totally different.

Outside of adventure games and hypertexts, changes to formal affordances feel under-leveraged in many genres of interactive story. We have constructed games to give the protagonist more and more abilities, but very rarely do townspeople react to the player differently after she has defeated powerful enemies, or behave with more and more wonder or treat the player as a celebrity on subsequent returns to a town. In fact, it is a common joke of MMO RPGs that the world does not ever seem to narratively react to the player's accomplishments, asking her to slay some menial pest only moments after declaring her a legendary hero for defeating a formidable enemy in the previous zone.

One of the most recognizable changes to formal affordances that occur as a result of player actions are story "branches" or narrative counterfactuality that occurs as a result of a player action. This might take the form of different story events, particularly those occurring after major dilemmas or highlighted player choices. It might additionally take the form of different character relationships, especially those occurring as the result of a major story choice or as the result of "faction" objectives being completed. Or finally it might take the form of different roles for the player character.

In the context of games that do not center narrative pleasure-experiences, formal affordance changes as a result of player action take on a slightly different form. In these cases, responsiveness through formal affordances might take the form of levels or terrain adapting to players' actions, enemy AI changing as a result of a player's strategy, items spawning as a result of player's actions, and so forth.

One of the greatest difficulties with formal affordances is that they offer the promise to the player we mentioned in the agency chapter and will return to in Chapter 7. In order for that promise to be upheld, they generally need to change a model or process in the system. Formal affordances suggest to a player that she can take an action, so in order for agency to be maintained, they must always be balanced by material affordances to actualize those promises.

5.2.3 Changing Material Affordances

For the computationally minded, the idea of a responsive system changing its affordances conjures images of programs that can rewrite themselves. This is the domain of artificial intelligence (AI) as depicted in science fiction—the ever-learning program that can rewrite itself—as well as the kinds of techniques we see as the focus of Game AI research today.[2]

Still, the promise of a system's greatest ability to really respond to something seems to rest in its ability to alter its affordances due to the player's actions, and a computational system with the ability to change its own material affordances in response to a player's actions is arguably the pursuit of many AI research efforts even if they do not overtly frame their contributions in this way.

In an interactive narrative context, material affordances are deeply-felt by the player. They might encompass things like:

1. What actions/verbs the system offers the player

2. What actions/moves the system itself can perform

3. System changes Its rules for evaluating what to do next

[2]When we think of popular speculative depictions of AI in works like *Her* [76], *Ex Machina* [65], *iRobot* [132], or *Detroit: Become Human* [40], the recognizable facet that makes the work about AI is the fact that computational agents can change their own programming—they can change their own material affordances. In reality, AI as a field is wide-reaching and varied in its pursuits. Mateas has convincingly positioned interactive drama into the AI tradition, not just in terms of believable autonomous agents, but also in the techniques of classical AI used to create interactive drama [100].

Additionally, each of the examples mentioned as potential changes to formal affordances must have material affordances available to support them in order for agency to be maintained. And just as Gibson suggests that our perception of our physical existence is shaped by the affordances of our environment, the player reframes her reality—her observations, curiosities, and abilities—through the material affordances offered by a game. She strategizes through the moves she understands she can make.

Changes to material affordances at the local and narrative levels are a rich area of exploration for the narrative research community. Common examples of material affordances changing in response to players might include:

- Conditional links or choice moments that are only available to players who have made certain decisions
- An AI opponent changing its strategy in response to player's moves
- Different approaches to a conflict becoming available as a result of player choice
- Dynamic difficulty adjustments
- Changes to the verbs available to the player

Many popular genres rely on material affordances changing to signal progression. RPGs, for example, typically offer new abilities to players as they progress. Players often expect to be more powerful at the end of the game than at the beginning, and that "power" is largely conveyed by a larger palette of abilities. Games with more simulation-focused gameplay such as strategy-simulation games offer players higher-level material affordance changes through the emergent possibility spaces created through the combination of various rules and states. In *Crusader Kings 3* [171], for example, the strategies available to the player at any given time depend on the complex interplay between the player's stats, traits, health, relationships, relative power, whether she has a viable heir, and so on. The viable actions available to the player in any given situation—whether a seduction attempt, assassination attempt, or war is likely to succeed in eliminating an enemy, for example—change dramatically in response to her actions. A player might spend several in-game turns executing a strategy. The game's response to this strategy (e.g., how NPCs respond to the player's actions) sets off a cascade of rules and state changes that changes how the player must approach the next obstacle. Here, the simulation creates affordances through the combinatorial interplay of the rules and states. (We will discuss changes to combinatorial affordances in Chapter 6).

Other genres, such as idle games (also called clicker games), feature changes to material affordances by virtue of changes to the scale of player actions. In *A Dark Room* [52], for example, the player begins by collecting sticks to stoke a fire in a dark room. But soon other people arrive, attracted by the fire, and form a village. As the game progresses, the player is able to automate the process of collecting sticks in favor of building buildings to attract villagers (who collect more sticks), and soon the player is focused on town management instead of stick collecting. Once the player manages the town well enough for it to become successful, she is able to leave the town she encounters a world map—common in other games, but surprising when first encountered in *A Dark Room*, since the game's affordances up to this point have been communicated through a purely textual interface. Much of the delight of the game rests in its affordances continuing to change in surprising ways as the game progresses.

It is worth remembering, however, that changes to material affordances must come as a result of player action for responsiveness to be upheld. *Frog Fractions* [173] , for example, is a game famous for surprising changes to its material affordances. The game starts like a typical math-learning game for children in the style of *Missile Command* [7], then turns into a typing game, then moves to space, then takes the form of a visual novel, and so on. The humor in the game comes from the nonsequitur changes to the games affordances and style. But the player does not do anything to cause these changes, nor are they signalled before they happen; if they were, the effect would be lost and the humor killed.

5.2.4 Changing Perceived Affordances

In the previous sections, we returned to a focus on formal and material affordance changes as we formulated our conception of responsiveness. In part, this is because formal and material affordances are the basis by which we judge the player's agency, so it is a natural starting point for assessing the system's response to that agency. But as an interesting aside, let us ask: what about the cognitivist theory of affordances—if responsiveness is about changing affordances what happens if we focus on changes to perceived affordances independently of whether formal or material affordances change?

In Chapter 4, we noted that much feedback is designed to bring a player's perceived affordances more in line with the real affordances a system offers. For example, Cardon-Rivera and Young discuss an example from *The Elder Scrolls V: Skyrim* [170] in which players receive UI feedback designed to advertise an affordance.

> A player, while exploring the surrounding area, has the chance to find her way to the top of the waterfall shown in Figure 2. A player likely knows – due to either from memory of this game, or other platform/adventure games, or even real life scenarios – that falling down a chasm typically results in them getting hurt and/or dying. Beyond jumping to certain death, a player perceives no other affordance. In reality, however, acting upon the cliff's affordance leads to the discovery of another affordance: that of the lake below which breaks the player's fall. Note that the cliff always provides the real affordance of jumping off the cliff. The player, however, perceives no affordance beyond jumping off with the consequence of dying.
>
> Unfortunately, there is narrative game content that requires that a player to act upon what Gaver (1991) would consider a "Hidden Affordance," (see Figure 1) which is problematic if the game's designers intended for the player to pursue the content; we assume so, otherwise, why insert it in the game? To counter the possibility of the player relying on her perception and problem solving skills, and consequently ignoring the content, the game's designers insert feedback to elicit a perceived affordance; Figure 3 illustrates what is presented. The discovery of the landmark is not triggered by the game until the player is precariously perched at the edge of the rock formation. When triggered, the words "Bard's Leap Summit Discovered" appear on screen. If we frame the interaction between a player and a game as a dialog between them (Young 2002), we can take the perspective of the player and ask: "why would the game have presented that landmark discovery at the precise moment that we approach the cliff?" Assuming that the game's designers are being cooperative (Young 2002) vis-à-vis Grice's Maxim of Relevance (Grice 1975), both the textual overlay's content and timing signify more than just a landmark discovery; this reasoning is borrowed from the cognitive faculties we use during every day discourse processing. Of course, while the textual overlay is feedback to elicit a perceived affordance in the player, it is up to the player herself to construct the correct mental representation that will enable action. [26]

As the authors note, feedback also communicates changes to formal and material affordances. In this example, we might reframe it as an example where the player's perceived affordances changed and the material affordances did not. The player's recognition of a material affordance she previously did not understand leads us to an interesting point: changing perceived affordances without a change to material affordances is actually mastery.

In discussing responsiveness as a system's ability to change its affordances in response to player actions, it is also worthwhile to examine what happens when a system does not change its affordances as a means of comparison. One interesting place where we see such an approach[3] is in games that rely on epiphany and recontextualization as a means of progression, such as in *The Witness* [181] or *Her Story* [10]. These are examples of games in which progression is signalled primarily through the changing of perceived affordances—that is, much of the pleasure of the game comes specifically from the lack of changing material affordances. Instead, the player gradually understands more and more of the material affordances at play as her mental model of perceived affordances expands. These cases are useful for us to isolate the aesthetic impacts of changing perceived affordances only without changing material affordances.

In *The Witness* [181], the player explores an empty island solving maze puzzles of increasing elaborateness and difficulty. The player's progression in *The Witness* is dependent on her ability to recognize shapes and structures within the world as she solves perspective-puzzles and recognizes more and more of the terrain of the island to be "solvable" in the game's core maze-solving mechanic. The final puzzle is solvable from the very beginning, but the player only knows how to look for it after she has played

[3]With the caveat that all interactive systems will be responsive at a low enough level of abstraction. For these works, I am looking past the building blocks of interaction—logics like navigation, collison, and on-screen display—and am looking at narrative-level affordances. To examine narrative, we must necessarily be looking at affordances that are abstract enough to combine these lower-level logics into a high enough level that the player can begin to form narrative understanding.

through enough of the game to change her own understanding of perceived affordances.

With each moment of realization, the player's perceived affordances broaden to accommodate new solutions to puzzles, each new solution offering a new technique to take forward that allows the player to recognize more and more mazes all around her. The player feels a sense of epiphany, thinking to herself, "oh wow, that has been here the whole time; I just didn't know how to see it!" It is not just that the player's skill increases, but much of the pleasure comes from the player's ability to recontextualize what is already there, her ability to see new formal affordances that have been in front of her all along, eventually culminating in a final puzzle that has been visible since the start of the game. After solving this puzzle, the player enters a final area before credits roll, signalling the completion of the game.

The power of the game's final puzzle comes from the epiphany that the final puzzle had been visible the whole time. The effect would be greatly diminished if the final puzzle were behind a door that remained locked until the player gained enough skill to unlock the door rather than being accessible from the outset. The pleasure-experience entirely depends on the material affordances remaining constant while the player's perception expands.

Her Story [10] offers an interesting comparison to *The Witness.* In *Her Story*, the player searches a database of video clips of a police interrogation. As in The Witness, the information is all there from the beginning, and it is through the player's increased understanding that a picture of narrative events begins to coalesce. She may type anything into the search bar, but from all of the possibilities in the english language, she only knows which keywords will yield fruitful searches as new narrative information is delivered. The player's epiphanies, their recontextualization of the information presented so far, and the process of making connections and discoveries encompass the core pleasure of the game.

Her Story takes the pattern a step further even than *The Witness* does by offering no formal ending sequence in the way *The Witness* does; the game ends and credits roll when the player decides to log off from the database. Effectively, the player just stops playing when she has finished collecting information, when she has formed an understanding of the story that satisfies her, having gained enough information about the narrative and having resolved enough of her open questions that she feels closure. This kind of design is bold; it really depends on always having at least one question at all times that the player is trying to resolve. Even *The Witness*, as focused as it is on progression through the player's understanding, makes sure the player can always see objectives ahead of her; she knows which areas of the island she has not explored, and areas of the island visually change after completion. *Her Story* is so confident in its ability to raise questions, that it does need the player to see remaining obstacles—the game is simply over when the player decides she has no more questions.

In both cases, one could argue that with each new revelation, formal affordances emerge in the recontextualization process. In *The Witness*, new formal affordances take the form of recognizing objects in the environment as solvable portions of puzzles whereas in *Her Story*, they take the form of new words or bits of information "crying out" for the player to search the database for particular keywords to uncover the next bits of a story. But an interesting distinction is that neither game is changing these formal affordances for the player as a result of her actions; rather, the player is merely understanding formal affordances that are already present through a new context and a clearer idea of how to utilize them. The key change is the player's perception.

I highlight this technique in particular to offer a contrast in later sections when we talk about the changes to "real" affordances—to highlight differences between perceived affordances changing versus formal and material affordances changing. Changes to perceived affordances alone feel more like *mastery* than the system offering *responsiveness*, though this is not to argue that such a technique does not offer aesthetic merits because of a lack of affordance change. On the contrary, it is precisely the staticness of the system's affordances, their lack of change, that makes this technique so powerful. The lack of system response allows the system to serve as a foil to the player's mastery. The lack of response is the very thing that creates the impact of the player's broadening understanding.

5.3 Nested Affordances and Levels of Abstraction

Recall that we justified nested levels of agency with the fact that nested affordances offered the ability to intend multiple nested actions at once. And because responsiveness offers the system's counterpart to the players agency within this model, we can similarly examine responsiveness at different levels of nesting as well by examining the changing of affordances nested at various levels of abstraction within a player

Level of Abstraction	**Agency Exerted**
Low	Player presses a button and sees **her avatar throw a punch as a result**.
Medium	Player inputs a complex series of button presses **and her character performs a super combo as a result**.
High	The player employs a "zoning" strategy, throwing fireballs to keep her opponent at a safe distance and hoping for her to jump so she can capitalize. **Her AI opponent slides under the fireball as a result**.

Table 5.1: Explanation of system responsiveness at multiple levels of abstraction simultaneously while playing *Street Fighter V*.

action.

In Chapter 3 we illustrated agency at multiple levels of abstraction simultaneously with an example from the fighting game *Street Fighter V*. Just as the player exerts agency at multiple levels of abstraction simultaneously, the system *responds* to that agency at multiple levels as well. The player presses a button, **and her character throws a punch as a result**. The player inputs the correct series of button presses, **and her character performs a super combo as a result**. The player employs a "zoning" strategy, throwing fireballs to keep her opponent at a safe distance and hoping for her to jump so she can capitalize, **and her opponent slides under the fireball as a result**. The bolded portions here highlight the points at which the system is issuing feedback that validates and reinforces the player's agency.

Each of the responses I have highlighted here come in the form of feedback, but each also changes the formal affordances of the opponent's position in relation to the player's, which open up different opportunities for attack and different optimal moves for the situation at hand.

How responsive a system "feels" can often be related to the level of abstraction at which the system responds. An opponent who blocks when a punch is thrown is changing formal affordances, but players expect some of their moves to be blocked by an opponent AI based on the conventions of the genre. The player probably does not expect an AI opponent that builds a complex player model of how often she tends to employ a zoning strategy and for how long, and thus comes prepared with the appropriate counter to that strategy. Though such a model would be functionally similar to a system that blocks when a punch is thrown, the higher level of abstraction makes the advanced player model feel sufficiently more responsive.

One key difference between the lower and higher level of abstractions, however, is that the higher level of abstraction allows more room for the system to change its affordances. At a lower level, the logics of the game tend to be fairly fixed. The way the game detects collision, for example, does not normally change. The higher-level logics, on the other hand, might emerge from the combination of the lower-level logics. There is potential for a larger possibility space of abstract affordances to result from those combinations, and thus the possibility that the system will have a greater variety of responses to wield. It is easier to change affordances in response to players when affordances are more abstract and thus more malleable.

5.4 Responsiveness and Emergence

If a larger possibility space of combinatorial logics makes it easier for a system to change its affordances, does that suggest that a system's degree of responsiveness is actually just its potential for emergence? Or to put it another way: are responsiveness and emergence synonymous? I would argue that they are not.

Let us consider a game of *Chess* against an AI opponent. *Chess* is a game with a large possibility and much room for varied play and strategies. *Chess* is normally played against a human opponent, who will naturally be responsive up to the level of her skill. A more novice player will respond to the current immediate threats on the board, but she may not recognize common strategies or be able to strategize about the game or her opponent at a very high level of abstraction. Different players may be able to adapt at different levels of abstraction or "meta game". *Chess* against an AI opponent might feel more or less responsive depending on the "skill" of the AI, and such factors might include how it models the player, how it plans moves, and how it adapts. The level of the emergence possible through *Chess's* systems has not changed, but the degree of responsiveness has.

As we will see in section 5.6, while games that depend on emergent storytelling have a lot to teach us about responsive narratives, critics might mean one of any number of things when they call a game "responsive." Those definitions, I suspect, really amount to a player's taste. It is for this reason that I have cautioned against calling games "more" or "less" responsive than others unless we are comparing things in very broad strokes—my goal is to introduce a lens through which we can discuss why games feel responsive in different ways.

However, just as a postmodern lens can be applied to many works of literature and some will yield more fruit, so too with emergent systems and an analysis of responsiveness. The lightswitch example demonstrates this: the basic lightswitch has very limited affordances, a single choice, and feedback at a very low level, so the total space of possible affordance changes is very small, and we can say just about all there is to say about them pretty quickly. The example of the lights with the voice-activated agent, however, gives us more to examine: we have more levels of abstraction to look at: the affordances and feedback of turning the light on, the affordances and feedback of the agent noticing a pattern, the affordances and feedback of setting a routine, the agent's introduction of theses changes, and so forth.

I will argue that it is not that emergence and responsiveness are the same thing, but rather that emergent possibility spaces in game narratives tend to exist within deeply-modeled simulations, which tend to be responsive in interesting ways.

5.5 Separating Responsiveness and Agency

I have argued that responsiveness is the system's counterpart to agency—that it fosters, supports, and reinforces a player's feelings of agency. But if agency is a "a phenomenon involving both player and game, one that occurs when the actions players desire are among those they can take (and vice versa) as supported by an underlying computational model" [193] , then is responsiveness actually separate from agency? Are we really talking about a *distinct* phenomenon, or are we actually talking about a *component* phenomenon? The answer, I would argue, is both.

We have seen how responsiveness can enhance and foster feelings of agency. A system's feedback functions as a component aspect of agency; the player needs confirmative feedback to fulfill the part of that definition of agency that where the player's desired actions are "supported by the computational model." Additionally, that feedback might provide new formal affordances to propel the player through the next iteration of the agency loop. And the system's ability to change its affordances provides the gap in the player's mental model that provides the next goal and keeps her engaged.

Systems can respond, however, in ways that undermine agency; they can change their affordances in response to player's actions in ways that betray the player's intentions or her trust. One example is dynamic difficulty adjustment systems causing unwanted intrusions. I had such an experience with a system in *Dark Souls 2* [50]. One of the pleasure-experiences of the *Dark Souls* series involves extremely difficult boss battles that feel immensely satisfying when you finally defeat them after many failed attempts. I had encountered a boss that I couldn't defeat, but each time I died, I got closer and closer to victory. Part of my trouble stemmed from the fact that I had to battle two huge knights between the spawn area and the boss lair, so I was often low on resources by the time I reached the boss. I had finally figured out the boss's attack patterns and almost defeated it, but it narrowly killed me. I was sure I would get it this time, as long as I did well on the knights on the way. Except *Dark Souls 2* implemented a system where enemies are removed from the game after you defeat them a certain number of times; the knights on the way to the boss were gone. Whether this feature was intended to be either a dynamic difficulty adjustment or a measure to keep players from grinding experience before the boss, it had robbed me of my victory. I arrived at the boss with full resources and killed it easily. The game had changed its affordances in response to my actions, but it had deeply undermined my intentions.

This experience raises the question of whether responsiveness as I have been using it is a term used to describe a pleasure or whether it is a property of a system that exists separately from the player's pleasure. Again, I would argue that it is both. Just as agency can be used to describe a balance in formal and material affordances that results in the player being able to enact her intentions regardless of whether the she particularly notes the pleasure she gains from that, responsiveness functions in a similar way: we might describe a system as responsive even if the player does not gain pleasure from it (as in my *Dark Souls 2* experience), but it can also be used to great effect as a pleasure unto itself, as we will discuss in some of the examples in chapter 6.

5.6 Responsiveness as a Lens

I have mentioned throughout this document that responsiveness should be used on a lens rather than a box. It is worth exploring this notion in more depth. We might start by asking, what kinds of conversations does this approach allow us to have?

For one, responsiveness as a lens shows the kinds of nuance in affordance we talked about above, and allows some of the kinds of analyses we see in Chapter 6 around particular aesthetic experiences. It allows us to look at how systems' abilities to listen, think, and speak can create or deepen methods for making the player feel heard. It allows us to understand how various design decisions can enhance players' agency.

The ability to separate the player experience of agency with the system's response that reinforces and fosters that agency, we are better able to understand how systems function within that loop. Many researchers have already provided very useful frameworks for understanding agency as we saw in Chapter 3. Extending these with some of the vocabulary introduced by Norman helps designers understand how we can use our palette of design techniques and allows us to better focus the aesthetic experiences we want to offer players.

Much design wisdom looks at how feedback interfaces with perceived affordances, but by examining the interplay between feedback and affordances—as well as how levels of abstraction factor in—we can look at which tools are appropriate to express responsiveness at different levels of abstraction.

Additionally, this model allows us to reconcile seemingly disparate examples of responsiveness, and the next chapter will explore the nuances of such examples in more detail. The responsiveness of *Bastion* is very different from the responsiveness of *Crusader Kings 3*, and subtly different from the responsiveness of *The Stanley Parable*. All of these are also very different from responsive controls in a platformer. Yet with this framework, we can understand that where Bastion is focused on responsive narrated feedback, *Crusader Kings 3* is focused on responsive emergent narrative moments, and both of these are operating at a different level of abstraction from the moment-to-moment responsiveness of “tight” controls or the macro-responsiveness of an adaptive AI opponent.

Allows us to talk about works that are obviously more or less responsive or are responsive at different levels of abstraction, but I would caution against trying to quantify felt experiences too granularly. These kinds of experiences are easy to compare at extremes and in broad strokes, but lenses are always more interesting than boxes. It is always more interesting to ask “what if we view X through the lens of which affordances change in response to the player?” than to ask “is X responsive?” All interactive media, after all, will be responsive at some level of granularity, otherwise it would not be interactive.

Additionally, the question of “is X responsive?” allows for a very imprecise answer. The question might mean one of several things ranging from “Is the player’s felt experience subjectively responsive?” to “Does the system offer opportunities for affordances to change?” When we talk about how responsive a system *feels*, this is a subjective claim. We might mean one of any number of things (each of which will be explored in more depth in the following chapter):

1. That input and feedback are tightly-coupled (as in responsive controls)
2. That the system changes its own rules, content, narrative scenarios, or input mechanisms as a result of player input—in other words, a game changes its formal and material affordances in response to player actions to a high degree (as in

games with a skilled Dungeon Master)

3. That a game offers a high degree of counterfactuality (as in simulation games)

4. That a game changes its affordances at a particularly high level of abstraction (as in a responsive AI opponent)

In practice, which flavor of "maximized" responsiveness most appeals to us likely depends on our tastes[4], but the ability of a system to change its affordances and feedback is the core that underlies all of them. We will explore how particular changes to affordances affect the aesthetics of a piece in more depth in the next chapter.

[4]As a researcher interested in how systems might respond to players in single-player narrative scenarios with strongly coherent narratives, my own personal biases lean toward option 2 as feeling the "most interesting" or "most responsive." Critics who find beauty in the specific interplay between rulesets, number patterns, and the symbolic representations offered by models might find they gravitate toward other answers. In practice, I have found that my approach has actually tended to take a bit from each of these approaches. We might think of each of these approaches as a spice that can flavor a soup: different critics will have different tastes, some might naturally complement each other, and some might find that large quantities of one adequately substitutes for others.

Chapter 6

Aesthetics of Responsive Systems

As I noted in the first chapter, a successful theory of responsiveness is one that we can apply to existing games to reveal new insights. In this chapter I will look at the aesthetic impacts of various affordance changes on existing games that are broadly considered to be "responsive." By using the vocabulary we have built up to this point, we are able to articulate differences among games that seemingly offer very different pleasures but all have been described as responsive, and can account for differences in how the game presents responsiveness to the player.

When a player says that a game "feels responsive", this is an aesthetic evaluation that can mean one of a number of things:

1. That input and feedback are tightly-coupled (as in responsive controls)
2. That a game changes its formal and material affordances in response to player actions to a high degree (as in games with a skilled DM)

3. That a game offers a high degree of counterfactuality (as in simulation games)

4. That a game changes its affordances at a particularly high level of abstraction (as in a responsive AI opponent)

In framing the contribution of responsiveness as a lens through which we can examine these properties, I will specifically avoid claiming that one of these is true responsiveness, and I suspect such a claim actually comes down to a matter of taste. Instead, I offer a discussion of some examples of each of these varieties of responsiveness. Recall from Chapter 1 that a successful theory of responsiveness should be able to articulate how games that feel responsive in different ways achieve their different aesthetic experiences. Over this chapter, I will attempt to do exactly that.

In offering examples of different games that demonstrate responsiveness, I cannot hope to be exhaustive. Even if I could, new design innovations would likely render such an endeavor obsolete the moment it was completed. Rather, I hope to demonstrate that the model of responsiveness offered in this dissertation opens useful avenues of critique and offers designers, critics, and researchers a vocabulary with which to interrogate how a game is constructing feelings of agency and a sense that the player is "being heard."

6.1 Tightly-Coupled Feedback

As we saw in Chapter 3, feedback is necessary for agency. Consequently, much of a game's development effort goes into designing feedback such that the player understands the impacts of her actions.

At the lowest levels, we recognize feedback about space, physics, and navigation through on-screen movement. At the narrative level, feedback often comes in the form of UI indications, dialogue, narration, or diegetic enactment that confirms our agency as we saw in the example of the choice of which classmate we want to sit beside.

To focus only on feedback as a means of properly delivering information misses feedback's potential for aesthetic impact. After all, some of the most memorable moments of game narrative are clever bits of feedback. These include things like:

- "You have died of dysentery." [136]
- "Thank you, Mario! But our princess is in another castle." [137]
- "You euthanized your faithful Companion Cube more quickly than any test subject on record. Congratulations.". [187]
- "You've met with a terrible fate, haven't you?" [39]
- "Clementine will remember that." [59]

Rather than solely relying on feedback for communication of state, several games have used feedback for aesthetic ends unto itself. In this section we will review a few examples of games whose narrative feedback contributes to a feeling of responsiveness.

6.1.1 "Clementine will remember that."

When Telltale's *The Walking Dead* released in 2012 to much acclaim [59] , one of its pieces of UI feedback quickly became iconic among certain adventure gaming communities. In Season One, the player controls Lee Everett, an ex-convict who encounters a little girl named Clementine while navigating the zombie apocalypse and takes her under his care. Matthew Byrd describes the phrase:

> The phrase “Clementine Will Remember That” references an on-screen message that appears early in the game shortly after the player decides how to address Clementine, a little girl whose parents have gone missing right at the start of the zombie apocalypse. Dialogue choices in gaming are nothing new, but it’s the ambiguity of that message that makes it so iconic. “What will she remember?” “What did I say again?” “Is that going to matter later on?” [21]

Ultimately though, the phrase offers a promise that the branching structure and underlying computational model cannot fulfill. Fans felt betrayed that the computational model did not actually change in accordance with the feedback they were given:

> You want to forgive [the developers], but then you remember that phrase “Clementine Will Remember That,” and you realize that there was deception behind that suggestion that still cuts deep. It sometimes feels wrong that The Walking Dead‘s legacy has been boiled down to a phrase that is ultimately meaningless within the context of the game itself. You can curse at Clementine, be the father she needs, or vary between both at your leisure. None of it really matters. In the end, she is going to end up kneeling beside Lee as his wounds slowly consume him.

The phrase eventually became a meme used to dismiss the lack of systemic modeling that Telltale Games actually offered, but such criticism also ignores the aesthetic impact that the feedback actually had. If Clementine wouldn’t actually remember our choices, why tell the player that she would?

The notification serves several purposes:

1. It creates a sense of anxiety. The player wonders whether she made the right choice or whether another choice might have been better. She reflects on her

actions and thinks about the paths not taken.

2. In reflecting on her choices, the player deepens her identification with the protagonist and roleplay by constantly asking herself "are my choices reflecting the way I intend to play my character?"

3. It draws attention to the fact that choices might have big impacts. As long as that possibility exists, it creates dramatic tension in big dilemma choice moments.

4. Once the convention is set up that the message will be displayed at key moments, it becomes an expectation that designers can subvert.

Early in the development of design conventions that later became the "Telltale Formula" [82] , designers focused their attention on maximizing the aesthetic experience of making choices–the focus was much more on how players felt while in the moment of considering which choice to make than how the eventual story would pay those choices off. Production realities around Telltale's episodic content development certainly contributed to the issue. However, the later catalogue of Telltale games focused much more on not only increasing the amount of branched content, but it also focused much more heavily on the ability to recontextualize content that needed to be reused such that it could be interpreted in multiple ways depending on what the player had seen. Season 2 of Telltale's Batman series, *Batman: The Enemy Within* [60] , for example, features two completely different takes on The Joker, and offers 5 different ways Batman and The Joker's relationship resolves. Their approach to branching and narrative payoffs changed substantially from 2012 to 2018.

While it is easy to focus on the fact that Telltale's *The Walking Dead: Season One* did not pair its feedback with an underlying systemic component, the way that its use of "Clementine will remember that" changed over time is itself an interesting case study in response. Even as player trust was broken that the message indicated a literal signalling of state change, it still functioned as a way to communicate to the player that she should be on her guard and approach choice with care. It still made her pause to reflect on whether she had made the right choice. After all, even if you were *mostly* sure the message did not mean anything, it still felt as if the game was threatening you with how much it *could* be hearing and remembering.

The most interesting use of the convention came in *Batman: The Enemy Within*, as Batman and Joker, both wounded from a battle the player has just fought, reflect on their history in what is clearly falling action as the game comes to a close. The Joker, in a moment of vulnerability, asks Batman if they were ever really friends, a key tension which the player has strategized around throughout the game. This is a big question, and one which the player can correctly surmise is one of the last choices she will have with The Joker. No matter how the player responds, the familiar feedback appears, but this time it subverts expectations and reads "The Joker will never forget that." No matter which version of the Joker the player has encountered, what the relationship between the two characters, or what the player has just answered, the message is impactful, either as a melancholic acknowledgement of a lost relationship, a vengeful threat to repay the player's betrayal, or a twisted start of something that resembles an abusive relationship, to name a few. Indeed, it is the message's ambiguity that allows its mutability. In any case, the feedback is both impactful in its subversion of expectation and in its confirmation that even if a choice doesn't change a state, the player has still been heard.

6.1.2 "The Kid just raged for a while."

Supergiant's *Bastion* [58] is another game that leans into feedback as an aesthetic force for making the player feel heard through constant moment-to-moment feedback. In Bastion, the player takes the role of the Kid, a silent boy who awakens to a world broken by an event called The Calamity. The player explores this world, while shooting and smashing things, and almost every move is narrated in third-person by a melancholy narrator. Jason Schreirer, reviewing the game for *Wired* writes of it: [153]

> Bastion's gruff narrator *is* the game. His play-by-play turns what could have been just another hack-and-slash adventure into a haunting, poignant experience. [...] Bastion 's narrator talks about what you're doing, what you're seeing and sometimes even what you're feeling, offering up tidbits and stories that add an authentically human emotion to the game's cartoony, abstract world. (original emphasis)

The narrator recounts most actions the player takes. The third-person, past-tense narration in response to player actions causes each move the player makes to feel narrativized, as if she is making history as she plays. Though many other games use voice-over responses to signal systemic changes or provide ambient narrative while the player keeps control, Bastion's feedback is notable in its quantity and the fact that almost every move the player makes is met with a quip or a bit of narrative explanation.

If the narrator is offering backstory, we might ask whether the feedback in

Bastion is offering new formal affordances. When I first played *Bastion*, the fact that the narrator commentated my moves was an interesting and novel (at the time) occurrence that led me to explore the bounds of that system. I expected that the narrator would only narrate progression-oriented exposition, so I started smashing pots in an area near the beginning of the game, off to the side. These particular pots clearly did not suggest that the game would be advanced by smashing them; I expected no response. Instead, to my delight, the narrator came back with "The Kid just raged for a while." It was a moment of acknowledgement that was absolutely delightful. As I played, however, it became clear that I was not *really* making choices that would impact the plot or change the game's direction through my actions. As many of the Telltale players were, I was disappointed that the feedback offered led me to expect a counterfactuality that was not actually supported by the game's narrative model.

Most of the narrator's commentary is explanatory. While it might occasionally suggest hints or offer bits that move the narrative along, in general it either clarifies action that is already occurring, or simply narrates what the player is doing. Even when the feedback does provide formal affordances by suggesting what the player should do in a particular situation, none of that feels triggered by a particular action the player has taken, nor does it signal that a particular narrative choice has been heard. The narrator's feedback does, however, offer constant, unrelenting reassurance that the game hears the player's actions—every single one. *Bastion's* narration offers a narrative equivalent to responsive controls achieved through constant, tightly-coupled feedback.

6.1.3 "Stanley stepped into the broom closet."

The Stanley Parable [22] is a game about choice and agency in games. Originally created as a fan modification (mod) of *Half-Life 2* [186], the game has the player take the role of silent protagonist Stanley, whose job is to push whichever button he is told to push at the time and for the duration he is told to push it, until one day nobody shows up to give him instructions and he finds that all of his coworkers are gone. The game features a narrator, as in Bastion that narrativizes the players actions in third-person past tense as she makes them. However, unlike *Bastion*, *The Stanley Parable* takes a pretty direct shot at agency and free will in games through the narrator's feedback.

AAA games at the time of *The Stanley Parable's* release utilized voice-over directives very heavily as a way to steer players toward objectives. *Bioshock* [57] famously commented on this trope by offering orders for the player to follow over a walkie-talkie, then admonishing the player for following these directives. In a reveal at the end of the game, the player learns that her character Jack was programmed to respond to the phrase "would you kindly. . . ?" through hypnotic mind control, and it is through this phrase that Jack (and the player) have been manipulated to carry out whatever directive was given to him. Because the convention is that the voice-over should steer the player, should drive her, *The Stanley Parable* feels genuinely surprising in the degree to which its narrator not only directs the player, but listens in response.

The Stanley Parable's narration blurs the line between narration and direction, but it is most impressive in its feedback. Steve Mullis writes of the game for NPR:

> It is when you get to your first choice between two doors that the true charm of this game comes alive. The narrator, brilliantly voiced by British actor Kevan Brighting, tells you what Stanley (or you) did in this part of the story. "Stanley went left," for instance. But here's the thing — you don't have to go left. You can go right, or back the way you came, or not do anything at all, and the narrator responds, often frustrated with you for not following the "rules" of the story. At times he even restarts the game to try and force you to play the story correctly.
>
> It's in this way that the narrator somewhat becomes the antagonist of the game; guiding you, or telling you, what to do, what you did, what Stanley is thinking or is not thinking and what the narrator himself thinks of all of the choices you make as you explore the office and try to find the narrative thread. [116]

The game is absolutely masterful in its interrogation of choice, but as this quote points out, much of the impact of that comes from the fact that the narrator responds. He has opinions about the player's actions and choices, and unlike *Bioshock*, she is free to make choices that run counter to the narrator's directives. And if she does, the narrator notices. He directs the player, but he also observes her. He is listening.

For example, if the player finds a broom closet and enters it, the narrator responds:

> "Stanley stepped into the broom closet, but there was nothing here, so he turned around and got back on track."

If the player then idles in the closet, the narrator continues with a new response every few seconds:

> "There was nothing here. No choice to make, no path to follow, just an empty broom closet. No reason to still be here."
>
> "It was baffling that Stanley was still just sitting in the broom closet. He wasn't even doing anything, at least if there were something to interact with

> he'd be justified in some way. As it is, he's literally just standing there doing sweet FA."
>
> "Are you... are you really still in the broom closet? Standing around doing nothing? Why? Please offer me some explanation here; I'm genuinely confused."
>
> "You do realize there's no choice or anything in here, right? If i had said, 'Stanly walked past the broom closet' at least you would have had a reason for exploring it to find out. But it didn't even occur to me because literally this closet is of absolutely no significance to the story whatsoever. I never would have thought to mention it."
>
> "Maybe to you, this is somehow its own branching path. Maybe, when you go talk about this with your friends, you'll say: 'OH DID U GET THE BROOM CLOSET ENDING? THEB ROOM CLOSET ENDING WAS MY FAVRITE!1 XD' ... I hope your friends find this concerning." (sic)
>
> "Stanley was fat and ugly and really, really stupid. He probably only got the job because of a family connection; that's how stupid he is. That, or with drug money. Also, Stanley is addicted to drugs and hookers."

And so on... the narrator continues this for several more individually-triggered lines. The broom closet here is just one example of not just the breadth of coverage of responses throughout the game, but the depth of responses written for each case. The feedback offers unexpected recognition of various actions and digressions; the game recognizes and provides feedback for so much more than the player expects it to, and then continues that feedback for much longer than convention would dictate is reasonable. And all of this amounts to a persistent feeling that the game is really listening to what the player is doing and what choices she is making.

6.1.4 Feedback and Counterfactuality

Due to the convention of leaning on feedback to convey system state change, the use of feedback at the narrative level generally implies counterfactuality. If feedback comes, convention tells us it generally means that something else would have happened (including nothing) if the player had not gotten this particular feedback or had instead gotten other feedback. If the player makes a choice and receives feedback that reacts to that choice, she presumes her choice caused that *particular* feedback.

This convention explains why there was so much frustration to "Clementine will remember that."; players expected that feedback to signal counterfactuality that did not exist. It also explains why the *Bastion* narration fades into window dressing; at some point the player realizes that most of the narration is not signalling narrative counterfactuality, nor is it signaling recognition of one choice over another. Still, the sheer quantity of feedback in *Bastion* creates an experience that feels like the game is listening; it is effectively the narrative equivalent of controls that feel responsive because the feedback is tightly-coupled with the player's inputs. Meanwhile, much of the feedback in *The Stanley Parable* is so powerful because it is directly referencing and commenting upon counterfactuality as it offers it. The Stanley Parable offers a high degree of counterfactuality in its The game wants me to take a certain action, and when I take a different one, it notices, and tells me so.[1]

And on the opposite side of the spectrum, some games use feedback to explicitly signal a lack of state change and counterfactuality. For example, Porpentine's *With Those We Love Alive* [71] offers feedback that explicitly signals that certain choices

[1]There is also an argument for differences in the feel of feedback among these games because of the formal affordances the feedback brings. "Clementine will remember that" arguably suggests a formal affordance that is not always paid off in a material affordance. Bastion's narrator only very occasionally raises new formal affordances as linear progression, but they do not offer any narrative counterfactuality. *The Stanley Parable's* narration constantly offers new and surprising formal affordances, each of which is supported by material affordances and counterfactuality.

will not change state, leaving the player free to express herself through choices she quickly learns are purely decorative (or reflective) [94] . Though the choices do offer counterfactuality inasmuch as they confirm decorative elements that change based on the player's choice, they explicitly do not affect the narrative state. The lack of state change leaves the player free to use these moments for roleplay without fear of them impacting the optimality of her play. In this case, it is the specific, signalled lack of responsiveness that gives the player this freedom.

Using feedback to make the player feel heard, as with many of the design techniques we will discuss, can be analyzed in the language of our model, but there is still an art to crafting aesthetically interesting feedback that makes players feel heard. Just as in writing, the existence of a story does not guarantee the story will be an aesthetically interesting experience; our enjoyment comes down to craft execution, surprise, play on expectations, and so forth. While feedback alone can make a game feel responsive, a certain level of feedback–even very surprising and clever feedback–is quickly internalized and normalized. In *Bastion*, for example, what begins as a delightful surprise quickly becomes internalized as simply how the game delivers its narrative. "The Joker will never forget that" is only so impactful because of its surprise and subversion of the trope. The fact that it is only used once is key.

6.2 Responsiveness Through Affordance Changes

Observant readers will note that our discussion of *The Stanley Parable* highlights not only the fact that its feedback is tightly coupled with player action, but also that this feedback provides the player with new formal affordances as well. For example, when the player enters the broom closet and realizes that the game has recognized the action as meaningful, the feedback is also presenting new formal and material affordances that signal to the player that entering the broom closet is possible and rewarding, and the dialogue provides formal affordances to steer the player back to the previous path. *The Stanley Parable* is full of affordances constantly changing as a result of the player's actions, and these changes make the narrator feel responsive in a way that is categorically different from *Bastion's* narrator.

Different games and genres employ different changes to their affordances. Here, I present a few examples of different approaches.

6.2.1 Directly-Authored Changes to Affordances

The simplest approach to changing affordances at the narrative level comes in the form of carefully crafted narrative branches and hand-authored flags. In this approach, an author specifies precisely which bits of content the player should experience if she makes a particular choice or triggers a particular event. Many of the games we have discussed in this chapter so far—*The Stanley Parable*, *Bastion*, and games that use the Telltale Formula—all use this approach.

When Stanley enters the closet, an event is triggered that tells the game to execute a particular bit of feedback—in this case, a particular voice-over. The trigger and the feedback were both directly specified by a designer, who hand-authored the line (which was then recorded by a voice actor). The line was carefully written to acknowledge the action the player would take to trigger the line, and to convey further hints (formal affordances) about the room—namely, that there was nothing special about it.

This authorial approach creates well-formed through-lines. Because the author dictates how affordances change as the player progresses, the author has a lot of control over the experience, making it easier for the author to ensure a coherent, well-formed narrative experience. Unfortunately, the authorial burden of this approach is very high, and many industry talks on narrative design tend to focus on mitigating the problems inherent to this burden.

Still, the majority of commercial games utilize this approach. Because of the associated authorial burden, many of the most narratively-responsive games to take this approach tend to be tightly scoped. The games that are not as tightly scoped tend to be remarkable for the sheer amount of content created in pursuit of this approach. *Disco Elysium* [196] for example, another game that utilizes this approach to great effect, contains over a million words [134].

6.2.2 Combinatorial Affordances

On the opposite end of the embedded - emergent narrative spectrum, games which heavily rely on the behaviors that emerge from the combination of rules and states offer changing affordances through the interplay of various rules and logics. Simulation games, for example, rely on a combination of processes to create changes to material affordances.

In *Crusader Kings 3* [171] , for example, the player takes the role of a monarch in a fictionalized version of Europe in the middle ages. Her character is generated, and has a variety of stats including skills (e.g., diplomacy, martial, intrigue, etc.), dread (i.e., how much her vassals fear her), health, fertility, stress, education, weight, and so on. Additionally, her character also has personality traits (e.g., arrogant, cynical, greedy, shy, stubborn, intelligent, etc.), temporary attributes (e.g., malnourished, imprisoned, pregnant), and additional modifiers she might gain from her actions (e.g., adulterer, kinslayer, maimed, pilgrim, etc). The number of possible traits is huge, and each of these stats and traits interact with each other and the rules of taking various actions in the game to create a complex interplay in which affordances are constantly changing.

At the start of *Crusader Kings 3*, the player's monarch controls a small area of land. Through strategic marriages, breeding, bequeathals, wars, negotiations, plots of intrigue, and so on, the player controls a dynasty over hundreds of years, taking the role of her heir as soon as her character dies. Because characters inherit physical traits as well as lands, the player must strategize about both the physical and political implications of marriages. She can change laws, religions, and customs, but doing so directly affects how her vassals view her and how cooperative they will be. Her goal is ostensibly to take over as much of Europe as possible, but in practice many players wander from this goal in pursuit of interesting narrative situations, roleplay opportunities, and self-directed challenges.

While many of the actions available to a player each turn are always available, the likelihood of success of any of those options varies wildly depending on the complex interplay between the player's stats, traits, health, relationships, relative power, whether she has a viable heir, and so on. As the player begins to understand that interplay, she finds that the viable actions available to her in any given situation change dramatically in response to her actions. Thus while lower-level affordances do not change very much, higher-level affordances change drastically depending on the player's actions.

For each new ruler in the player's dynasty, a common goal is to find a suitable spouse so that she can ensure a suitable heir. To accomplish this goal, she opens a list of possible spouses, and must decide whom she will send a proposal. The game helpfully allows the player to sort prospective spouses by alliance power, prestige gain, rank, age, and so forth. When making this decision, the player is trying to reason about which of those stats is currently most important—alliance power, for example, can be crucial for creating alliances to aid in military victories, plots of intrigue, and so on; however, a spouse who is next in line for a large inheritance might be more advantageous in this moment. Additionally, congenital traits are inherited, so if the potential spouse is *melancholic*, for example, it lowers the spouse's fertility and could be passed on to future heirs. Depending on her priorities, different options will appeal to her, and whomever she chooses to be her spouse will influence the traits of her heirs, which will in turn change which actions are available to her in the future. Each decision has a cascading effect on future available actions.

Unlike the approach in which a designer authors exactly how and when affordances should change, instead the designer authors a rules system out of which affordances arise and change through the combination of effects, statuses, available moves, etc. The combinatorial effect in which small rules combine together for larger stories to emerge allows for a greater space of affordances changing as a result of the player's actions, but predicting the outcome can be difficult. Consequently, games that rely on simulation for their underlying models might get interesting and unique stories, but also need more design overhead to ensure coherence.

6.2.3 Contextual Affordance Changes

Crusader Kings 3 offers an interesting example of a game in which contextual changes make different moves optimal or available depending on state. Though a player might attempt to murder an enemy when the game is in a variety of states, the plot is likely to be successful in a relatively few scenarios.

However, aside from just the possibilities of success vs failure—even when failure provides interesting narrative conflict as it does in Crusader Kings 3—the idea that a system might offer different affordances in response to the same action depending on context is an interesting way to make the player feel heard. *Façade*, as an example—

offers different formal affordances through conversation over the course of the game. In *Façade*, the player presumes the role of a guest visiting the home of college friends Grace and Trip, who descend into an awkward argument over the course of the evening. The game consists primarily of conversing with Grace and Trip, answering questions, interjecting during disputes, and helping them resolve their conflicts as the night unfolds. Not only do such formal affordances arise and change in the natural flow of conversation, but also various player utterances might be interpreted differently depending on the surrounding context, and thus the player will receive different responses to the same input.

For example, complimenting Grace on the decorating early in the experience might elicit a positive reaction, but later, looking at or commenting on the sofa might trigger a beat in which Grace complains about the decorating and tries to pull Trip into an argument. As conversation moves, new topics are introduced and change the context of both the overall narrative and the current discourse, both of which influence the particular way the system will interpret player inputs and respond to them. Certain phrases might have radically different results depending on when the player expresses them, and affinity between the player, Grace, and Trip changes accordingly as a result.

Unlike the emergent narrative of *Crusader Kings 3*, which relies on the player's work to interpret stat and state changes as narratively meaningful, *Façade* recognizes context changes and offers logical affordance changes in response. The result is an experience deeply grounded in embedded narrative. Unlike *Crusader Kings 3*, the player can remain narratively immersed, reasoning and making decisions from within the fiction rather than surfacing from the perfluent dream [64] to make strategic decisions and interpretations of numbers.

6.2.4 Community-Driven Affordance Changes

While the changing formal and material affordances of *Crusader Kings 3* emerge from the interplay of rules and state, esports offer an interesting example of affordances that the developers change periodically in response to player behaviors.

Overwatch [43], for example, is a competitive multiplayer first-person-shooter game in which teams of 6 players try to either take control of "objective points", a designated area of the map which they must stand on without the other team standing on it, or move an objective along a path by standing on it without the other team standing on it ("payload"). Players on the team have individual abilities and designated roles on the team: tanks, which protect other players and absorb damage; DPS characters, who deal the most damage to the other team; and support characters who heal other characters or provide them with bonuses and other helpful abilities. The interplay among the characters is very carefully tuned such that even small changes to a character — for example, decreasing a shield ability by a small amount — has cascading implications for potentially optimal strategies.

Competitive play is organized into seasons, which last for a few weeks. Each season, developers make such small changes with the explicit goal of changing the current optimal playstyle in a process known as "changing the meta(game)". Similarly to *Crusader Kings 3*, this changes the viable strategies, and thus changes high-level affordances, but unlike Crusader Kings 3, this is a deliberate move periodically undertaken by designers to respond to players' current strategies and playstyles. [2]

Blaseball [8] is a fascinating example of a game that has taken this approach to community-driven affordance changes and repurposed it for a narrative context. *Blaseball's* community-run wiki explains the game thusly:

> Blaseball is a game that simulates an absurdist version of baseball and encourages participants to bet (no real world money involved!) on the outcome. Participants use those winnings to buy voting and raffle tickets that influence new game mechanics, player and team statistics, player trades, developer-written storylines, and more. A blaseball season occurs over the course of 99 "days" that correspond to one real world hour. Blaseball seasons last from Monday to Friday (United States Pacific Time), then a post-season with four teams from each of the two leagues occurs on a Saturday. Election results are posted on Sunday and the game's plot progresses in accordance with the voting. [180]

Blaseball is still ongoing as of the writing of this dissertation. The events that have occurred in *blaseball* so far as a result of fan voting have been absurdist (to the point of making the game very opaque to outsiders). For example, at the end of season 1, players voted to open the forbidden book, after which several events occurred that changed the nature of the game. From the wiki:

> In Season 1, the book first made itself publicly known in the form of a decree available to participants of blaseball. Following the inaugural Internet Series, the decree called Open the Forbidden Book was successfully passed with 556 votes, or 61% of all decree votes. The immediate results of this decree were as follows:
>
> - The Book Opens.
> - Solar Eclipse.
> - Umpiress' eyes turn white.

[2] It is worth noting that this process happens in traditional sports as well, though these changes tend to be small and often remain the same for years or decades, where *Overwatch* changes its meta every few weeks.

- Star player Jaylen Hotdogfingers is incinerated.
- Hellmouth swallows the Moab Desert.
- THE DISCIPLINE ERA BEGINS

Subsequently, the end of each season often brings huge changes to the game that add to lore, change the nature of the mechanics, and completely reframe what is possible within the game in response to players' actions.

6.2.5 Improvisation as Changes to Formal Affordances

Improv theater is an interesting example of the ability for formal affordances to change, and I suspect why so many games research efforts are influenced by it.

In improv the 'yes, and" ethos means actors operate from a goal of taking any utterance by a counterpart actor as valid, and they work to fit their actions into a narrative that is being collaboratively spun. In practice, this means that everything can be recontextualized in a minute by a new utterance.

For example, in a sketch from the improv television show *Improvaganza* [27], Colon Machary and Ryan Styles act out a scene in which they are a couple having an argument. Throughout the sketch, they must try to work in lines written by the audience on small scraps of paper, which they unfold and read for the first time as the scene transpires. As they act out the scene, one of them might say a line (either from the audience or their own imagination) that changes the context of the scene:

Colin (sobbing): I've hurt you!
Ryan: You have hurt me.
Colin: I'm sorry!

Ryan: I'm sorry too. It's just...
[Ryan unfolds paper.]
Ryan: The super bowl. I'm just so excited about the Superbowl!
Colin: I understand. (suddenly excited) So am I!

The nonsequitur changes provide the basis for the comedy, but they also allow the scene to radically change formal affordances in an instant. Not only can one actor change the context of the scene for all actors in the scene, crucially she can quickly change *what makes sense to do next.* Something may be a formal affordance in one moment, crying out for an actor to take it in one direction, and another actor can subvert that affordance, take it in another direction, and suddenly it ceases to be an affordance. The narrative moves ever-forward. [28]

Games that depend on improvisation have potential for flexible changes to formal affordances. *Dungeons and Dragons* (*D&D*) [68] , for instance, is a game in which the players and a dungeon master (DM) collaboratively construct a story together as channeled through a rule system. The ability of the DM to change formal affordances on the fly is the soul of what makes *D&D* one of the most responsive narrative games.

The responsiveness of *D&D* arises from the fact that a human player can be very good at changing formal affordances on the fly. Of course players can change the rules, modify the systems of *D&D*, fudge dice rolls, etc., but even with no such changes, the game is highly responsive because the DM chooses which aspects of the systems to highlight at any given time. She sets the narrative world for the players, and changes that world in response to the players' actions. Without changing any of the underlying system that governs *D&D*, a narratively-responsive human DM can adapt which narrative elements she offers, what kinds of locations the players will encounter, how characters respond to players, and what strategy seems most optimal. If players are not "getting" a piece of information, a DM can choose how strongly to suggest a certain path. If players want to explore a new area, the DM makes up everything about that area and what the players encounter there. She can easily change formal affordances at will.

Each formal affordance the DM offers is a promise. If the DM suggests that there is a lush forest to the east of town (a formal affordance), but then tells a player she cannot go explore it, the DM has created an artificial barrier—she has restricted the material affordances of the scenario and created an imbalance in affordances that undermine the player's agency. Nothing in the rules prevents such an exploration, and this leads to the phenomenon known as "railroading".

But as long as the DM is not railroading players, the fact that the DM will adapt her planned narrative to the players' desire to head in an unexpected direction, to take an unexpected action, or to focusing on an unexpected detail creates a deep sense that the players' intentions are being heard and recognized.

6.2.6 Toward the Holodeck

One of the difficulties in constructing narrative systems with the ability to change formal affordances in response to player actions is that each formal affordance offers a promise to the player, and paying off that promise is expensive. Additional art assets, animations, processes, voice recordings, etc. can balloon development costs very quickly, making such a payoff prohibitively expensive with current tools and limitations. Much of the procedural narrative research currently underway in the game industry focuses on alleviating this cost as one of the most pressing hurdles to overcome. From a practical standpoint, the more formal affordances a game offers, the more robust a computational model must be to support the material affordances necessary to maintain agency.

Popular fiction examples of responsive narrative systems, such as the Holodeck from the *Star Trek* series [143] or the titular amusement park from *Westworld* [123], are so impressive because they offer a range of affordance changes that the computational systems of science fiction are able to seamlessly support. The promise of the holodeck, and I think interactive drama as a medium, is the promise of a piece that changes its formal affordances in response to player actions throughout the piece at a variety of levels of abstraction. It performs these formal affordances changes in an improvisational manner and has a robust enough computational model that its material affordances can support such changes.

Façade is the first fully-realized interactive drama, and much of its strength comes from the fact that it changes its formal affordances throughout. While the kinds of contextual changes to affordances we saw in section "Contextual Affordance Changes" (Section 6.2.3) offer robust material affordance changes that make sense with the narrative moment at hand, additionally the game offers a variety of responsive formal affordances throughout. The work is radically responsive in how its narrative design signals formal affordances at the narrative level (e.g., different kinds of appropriate choices or conversation topics), and different points in the story. Robust computational modeling accounts for most reasonable player inputs around those affordances, and player responses change relationships with characters accordingly. The whole story feels like it changes because key plot and relationship moments change.[3]

6.2.7 The Chicken or the Egg: Formal or Material Affordances First?

In discussions of formal and material affordances changing, it is important to note that we cannot change one without changing the other lest we break agency. Changes to material affordances must be signalled through formal affordances and changes to formal affordances must be supported through a computational model with material affordances robust enough to support the changes to formal affordances.

So if both must change to support each other, it is fair for designers to wonder, should creators design responsive formal affordances first and craft a computational system to support them, or design responsive material affordances first and craft formal affordances to advertise them?

The answer depends on the intended pleasure-experience of the game.

Many games will naturally fall more toward a pleasure-experience in which players reason and strategize about systems; others will tend toward an experience in

[3]Still, Mateas recognizes that *Façade* is successful in part because of its constraints [100]: it has a very limited cast, its setting is limited to a single location, it takes place in a social situation that offers natural behavioral constraints to the player due to social norms. As impressive as *Façade* is, there is still a huge leap between it and the promise of *Westworld* or the holodeck. And it is worth noting that even though *Façade* released in 2005, there have not been any full-scale AI-driven interactive dramas of its type released since, though the popularity of VR dramas with live actors, in-person immersive theater experiences, and procedural narrative systems has increased dramatically in that time. I wonder if the reason for the dearth of AI-driven interactive dramas is, in part, due to our lack of ability to fully understand what exactly made *Façade* such a compelling experience, and whether we consequently presumed replicating it was about solving technological problems rather than design problems.

which players strategize and reason about roleplay and narrative concerns. Some games really do fall directly between the two. For these games, we might think of individual features as being more narratively-led or more systemically-led, but often the formal and material affordance design will iteratively influence each other.

6.3 Responsiveness Through Signalled Counterfactuality

The existence of categories like "choices matter" and the wealth of developer design wisdom on how to manage branching narrative patterns suggests that for many players, the idea of lots of branches or multiple endings remains at the heart of what players expect from responsive narrative. For these players, the idea that their choices have been heard rests primarily on whether their choice actually changes an underlying computational model. This idea feeds the naive intuition that larger possibility spaces lead to more responsive narratives.

In practice, we know that the size of possibilities spaces is not actually what makes a game feel responsive. The "oatmeal problem" [32] is a well known problem in procedural generation communities: we can generate vast amounts of something, but humans are experts at pattern-matching; we are quick to recognized when differences among generated output are superficial. It is very easy to generate one billion variations on fundamentally the same thing; a the difficulty of ensuring that a large possibility space actually leads to meaningful difference is a key design challenge of procedural narrative.

Counterfactuality supports responsiveness, though most players will only know upon replay, and whether it is necessary for responsiveness is a rabbit-hole. Just as with agency, there are certainly compelling experiences to be had through the illusion of counterfactuality, but they will fall down upon replay. Whether that is important to a creator will depend upon her goals; very few players actually replay games, and cost ultimately governs the amount of counterfactuality a creator can provide. Personally, particularly in the age of streaming, collaborative play, and participatory culture, I am most interested in systems that provide counterfactual responsiveness.

Still, invisible counterfactuality could plausibly reach a point at which the player is unaware that her choices have actually had an impact. Consequently, for many narrative games, a lot of development effort goes into ensuring that counterfactuality is visible, signalled to the player by various means. By signalled counterfactuality I mean that either:

1. The road not taken is clearly highlighted
2. Or the possibility space is presumed to be so large that the player would be surprised that counterfactuality was not at play

Sometimes the solution is just to indicate systemic elements through UI feedback. The latter strategy is the one we associate with interactive drama and highly simulated narratives: of *course* the holodeck offers counterfactual responsiveness; it could not support the range of player actions it offers if it did not. However, in the absence of such deeply robust simulations, many narrative games opt for the former strategy and highlight the road not taken.

One way games signal their counterfactuality is by highlighting their systemic elements *as systems* rather than attempting to keep the player immersed in dramatic action and roleplay. Though much of game narrative still treats games and systems as forces in conflict (see section 2.3), this need not necessarily be so [191]. Narrative simulation games tend to have an advantage in the integration of systems and because their systems are already designed to convey narrative through the relationship between rules and the patterns of actions, the changing of traits and stat movement, and so on.

However games need not necessarily offer pleasure-experiences that center poking at simulations for games to highlight their systemic elements. Games like *Fallen London* [53] and *Long Live the Queen* [56] do this particularly well by offering interesting interplay between the narrative and systemic aspects of the game; though the player is constantly aware of the systemic elements and strategizing about them, this complements the narrative experience rather than detracting from it. In *Long Live the Queen*, the player takes on the role of a teenage girl, Elodie, whose mother dies a year before Elodie comes of age to become queen. Each turn, the player decides which skills Elodie should learn to prepare her to successfully rule, then she experiences a short choice-based narrative section during which her skills are tested as challenges, social puzzles, and assassination attempts assault her from all directions. The game signals each skill check very explicitly, so even during sections where the game is apparently more focused on the player's narrative immersion, she is aware of places where her skills are deficient, and the story might have changed. While the game does an excellent job of integrating the skill system with the narrative and roleplay opportunities available to the player, the player's stats and skills, and the systemic elements of the game are never far from the player's mind.

Figure 6.1: Screenshot from *Long Live the Queen*. The player has just made a choice and now a string of skill checks suggest that the choice had counterfactual outcomes the player is not seeing.

Where some games use systemic elements to signal counterfactuality, others signal counterfactuality diegetically through clues in the narrative. The goal of this approach is often to obfuscate the systemic elements of the experience (even if temporarily), to maintain narrative immersion and avoid the kind of choices that shift players from narrative roleplay to mechanical strategizing. *Life is Strange* [44] took both approaches—signalling counterfactuality with an on-screen butterfly icon, but also offering dialogue that hints toward other paths the player didn't explore. Signalling counterfactuality diegetically is difficult: if the signalling is too subtle, it is invisible, but if it is too hamfisted, dialogue can become stilted and awkward. When it works well, the effect can be very rewarding, but approaches to diegetic signalling must be varied. Lines offering very direct signalling such as “I could’ve taken a picture of her covered in paint, but I didn’t and we had a genuine moment.” might be fine if used sparingly, but it does not take very many of such lines for a scene to feel inundated with unnatural reports of previous actions.

Detroit: Become Human [40] took a different approach still, reminiscent of early hypertext literature [75] [18], by offering a map between scenes to showcase the explicit narrative structure and highlight places where the narrative might have branched.

Figure 6.2: Screenshot of level choice map from *Detroit: Become Human*. Choice maps at the end of each level make narrative counterfactuality very explicit.

While we know that how counterfactual a game is—that is, how large its possibility space is beyond a certain point—does not correspond directly to the feeling of responsiveness [32], the efforts to signal that counterfactuality might indirectly factor into feelings of responsiveness. In an effort to signal the path not taken, the examples above all offer a higher degree of feedback to achieve signalling goals. Similarly, as we saw in section 6.1.4 "Feedback and Counterfactuality," games that depend on UI feedback to offer acknowledgement of choices might imply counterfactuality through convention. Similarly games that tend more toward emergence and implicitly signal their counterfactuality by means of narratives being more simulation-focused might naturally tend more toward responsiveness by offering combinatorial affordance changes (see section 6.2.2 "Combinatorial Affordances"). In each of these cases, it may not be counterfactuality that makes "choices actually matter" on a first playthrough, but rather the efforts to signal counterfactuality make choices matter through increased use of other techniques toward responsiveness. Additionally, counterfactuality supports aspects of narrative games like replay, streaming, and participatory or comparative play.

6.4 Changing affordances at a high level of abstraction

When discussing responsiveness at different levels of abstraction, it arises that a more responsive system—a more responsive AI *Chess* opponent, for example—might correspond with the idea of a system that can respond at a high level of abstraction. After all, among human players, a player who can adapt to higher-level strategies is surely more responsive than one who can only respond at the level of individual moves.

The narrative equivalent of this idea—the ability for a narrative to respond to players at a higher level of abstraction—is an interesting area of inquiry. Currently, at the highest levels of narrative abstraction, a player might decide what hero archetype she will roleplay: will she be a heroic paladin? A rogue with a heart of gold? A priest with a darker side? But what if as she started playing, she realized she detested her party mates and instead she could decide this wasn't a hero's journey at all, but a tragedy about her rise and fall as a villainous fallen cleric.

So far the most successful responsive embedded narratives have offered very constrained narrative settings. Even ignoring the sophisticated, believable androids we see in *Westworld* [123], we currently do not have the technical capabilities to carry out something with narrative design at that level of complexity in large part because we do not yet have a system that can carry out responsive narrative affordance changes at various levels of abstraction, particularly the highest levels. These high level narrative affordance changes are also an area where *Dungeons & Dragons* excels.

But hopefully by better-formulating the problem, we might begin to take steps toward new approaches to solutions.

6.5 A Matter of Taste

Over this chapter I have articulated four different ways systems might convey responsiveness:

1. Input and feedback are tightly-coupled (as in responsive controls)

2. A game changes its formal and material affordances in response to player actions to a high degree (as in games with a skilled DM)

3. A game offers a high degree of signalled counterfactuality (as in simulation games)

4. A game changes its affordances at a particularly high level of abstraction (as in a responsive AI opponent)

Each of these offers a version of responsiveness and appeals to different tastes. They are not mutually exclusive, and I have tried to use some bridging examples to demonstrate that responsive games can offer one or many of these. Additionally these are not cleanly-delineated categories; naturally for a system to change its affordances at a high level of abstraction (4), it must be able to change its affordances at all (2). Instead, I mean to highlight that when we say a game is responsive, we should be clear about whether we mean one of these in particular or a combination of them. Additionally some of these flavors of responsiveness might naturally pair well with specific mechanics or pleasure-experiences. Some might naturally mix in certain quantities or ratios better than others.

This chapter demonstrates that the model of responsiveness offered in this dissertation opens useful avenues of critique and allows us to have new conversations around agency. The vocabulary I have introduced here offers designers, critics, and researchers a new perspective with which to interrogate how a game is constructing feelings of agency and a sense that the player is being heard.

Chapter 7

Interventions Toward System Responsiveness

This chapter examines how the aesthetics of responsiveness intersect with the system's technical capabilities in more detail. Following in the tradition of Expressive AI proposed by Mateas [100], I believe that the design of complex systems is inextricably linked with their technical offerings. This chapter explores places where technical design intervention creates responsiveness, situate current research efforts into some of these improvements, and highlight design considerations apparent in each intervention.

In the introductory chapter, I argued that a successful theoretical model might not only help us understand why works that feel responsive do, but it might also guide research efforts toward better integration into narrative experiences.

We previously framed a system in the context of the Listen-Think-Speak model (the LTS loop), and in Chapter 6 we explored how changes to affordances signals to the player that the system is responding. In this chapter, I take a closer look at each step of the LTS process to describe how the system's ability to listen, think, and speak factors into its ability to change its affordances and thus how changes in the system's architecture are in turn translated to different aesthetic experiences of responsiveness.

Over this chapter we will first examine *how* systems listen, think, and speak from a technical standpoint. Then we look at technical interventions toward responsiveness at each stage in the loop. Finally, I also explore the design considerations that translate each technical intervention into a responsive experience for the player.

Interactive narrative is a robust field of technical research, and I cannot hope to be exhaustive in covering various approaches. As with all enumerations in this work, it is meant to offer a skeleton, not a cage. Instead, my goal is to demonstrate that research efforts in the field *can* be analyzed in this way, and that doing so allows us to understand which of our research efforts might lead to responsive experiences, and which ones need more connective work—ideally this model might also serve as a starting point toward articulating research problems and opportunities toward bridging that gap. Still, I believe that demonstrating the utility of this kind of analysis, even if not demonstrated exhaustively, achieves our goal of applying our theoretical lens to the domain at hand to gain new insights into design and technical approaches to responsiveness.

7.1 How Systems Respond: The System-Side LTS Loop

Having now explored the phenomenological experience of responsiveness—how it feels to players, and examined how changing affordances and feedback affect player experience, let us now examine how systems respond at a technical level.

We have detailed how players understand a computational system through its affordances, using them to build a mental model of the system and form an intention, then taking actions based on those intentions and observing the results of their actions. But how does that process look from the system's side, and what aspects are key to responsiveness? A representation of responsiveness as the system's side of the LTS loop is depicted in Figure 7.1. Returning to our model, let us take a closer look at how systems listen, think, and speak, especially in the context of interactive narratives.

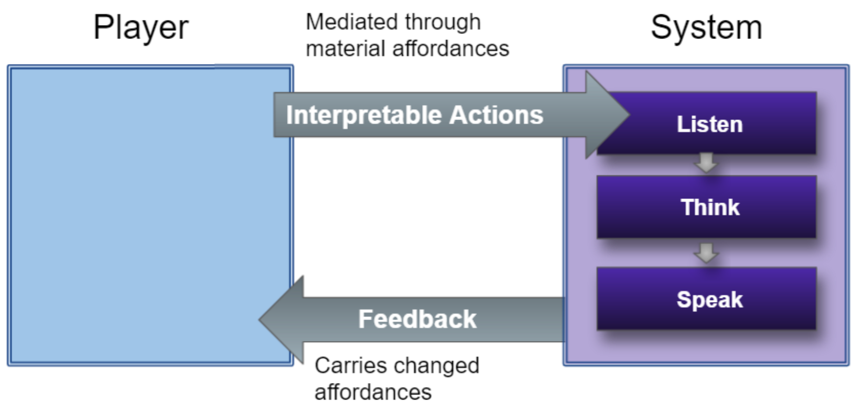

Figure 7.1: Responsiveness as a component phenomenon in the LTS interactivity loop on the system's side.

7.1.1 How Systems Listen

What systems hear is filtered through their material affordances, and ultimately, the set of inputs the system will recognize is predetermined. Modern interactive systems might recognize mouse inputs, keyboard inputs, controller inputs, gestural inputs, natural language inputs in the form of speech recognition or typed inputs, and so on. The types of inputs the system recognizes necessarily function as material affordances through which the player understands her possibility space and through which the system will understand her.

At the lowest level of abstraction, listening literally takes the form of event listeners—processes that check for inputs like mouse clicks, key strokes, button presses, and so forth. At a level above that, a system might listen for a certain link to be activated, a certain word to be input, a certain gesture to be made, etc. At each level of abstraction, the system is designed to recognize particular inputs as meaningful. As we reach narrative levels of abstraction, a key component of designing a responsive narrative system is deciding what the system should listen for and recognize as meaningful. Interesting questions in this design space include things like:

- If the player selects a dialogue choice, what kinds of things should the system interpret from that choice?
- If the player enters text or voice input, what kinds of things should the system listen for and interpret from word choice, tone of voice, speed of input, etc.?
- Should the system collect information about how long the player takes to make a

choice or take an action? What should the system glean from such information?

- If the player takes an action, what about that action should the system consider meaningful? And to what end?

At higher levels of abstraction, we can also envision systems that try to detect complex strategies to counter them. Systems might also analyze behaviors to look for cheating or *smurfing* (playing below one's skill level in a multiplayer game in order to have an unfair competitive advantage). And many games are turning to automated solutions to help with harassment detection between players. Such systems might look for patterns of play, players for whom particular metrics (a kills-to-death ratio, for example) are extreme outliers, or patterns of word usage and sentiment analysis respectively. Many of the most interesting problems for data scientists in the games industry are actually problems of figuring out how to listen for certain behaviors and intentions.

As players, without the ability to see a system's code, we primarily understand the listening capabilities of systems through the feedback they give us to reveal these affordances. For example, from our lightswitch example in section 5.2, we understand that the simple light switch has "listened" because the lights go off. The light switch assistant's question "You sent the same request yesterday at this time; would you like me to add this to a daily routine?" reveals that the assistant not only listened for us to speak and listened (and recognized) our request to turn the lights on, but also that it listened for a pattern to our actions with the lights.

Places where the system takes special care to ensure that listening is done well are particularly interesting. For example, parser-IF games often feature a disambiguation response: if the player types "TAKE" and there are multiple objects which can be taken, the parser might ask "Which do you mean, the apple or the orange?" Versu's [46] relationship system offered similar disambiguation for relationship intentions by tracking friendly, flirty, or hostile player moves towards a particular character. The system might then offer a menu option asking whether the player's intentions were in fact as they seemed (e.g., letting the player specify that she is indeed trying to make an enemy of this character). That input would consequently shape which further affordances were offered, and how the system understood later input. [161]

While the ability for systems to recognize all manner of inputs has increased dramatically over the last couple of decades, one of the primary difficulties in offering a wide range of possible inputs is ensuring that the range is sufficiently modeled; it does not matter if the system can perfectly hear and distinguish a huge range of natural language speech inputs if none of them actually do anything at the systemic level—i.e., if there is either no process that recognizes certain words as conversational triggers (material affordance) or no clear reason of what to say at any given time (formal affordance). How the system listens is deeply intertwined with how it thinks.

7.1.2 How Systems Think

When we hear of systems "thinking", particularly within the realm of artificial intelligence, it is easy to imagine various science fiction interpretations of exactly what that means. But in order to avoid a descent into the philosophical underpinnings of our field, let me clarify that I mean this in a metaphorical sense. Systems “think” via a suite of (1) data, and (2) processes which act upon that data [189]. Such a broad definition is a useful framework for considering our theoretical model as a lens that can apply to system actions across a range of architectures, player actions, and levels of abstraction. So if the player’s think step involves interpreting affordances to build and update a mental model of the system, what does the system counterpart look like?

First, the system undergoes an analogous process to the player’s wherein it interprets the player’s inputs through the material affordances available to it. Only inputs from the interpretive step that correspond to existing functionalities in the system can be recognized. Using these inputs, the system forms or updates a model of the player. Systems might model the player either explicitly (as mobile games do) [9], saving data about her position, stats, play style, spending habits, etc., or a system might model the player implicitly. For example, even though many hypertext works do not explicitly encode a player object to which variables are attached, the player’s position in the work, flags that capture whether various lexia have been visited, etc. model the player implicitly by capturing her progress and history in the world state. Various systems use player models to different ends as we will see in section 7.2.3, but a key split is whether the system is primarily modeling the player’s character (her position, character stats and traits, her relationships to other characters, her history in the world) or whether the system is attempting to model the human player herself (e.g., her playstyle, skill level, strategy, spending potential, etc).

In addition to modeling the player, the system must also keep track of the state of the game and game world as a result of the new inputs received. A game might track data for NPCs (position, relationships, abilities, state), levels (terrain, layout, event triggers, NPC spawns), items or collectibles, and so forth. How to update the state is determined by systemic or narrative rules that indicate which data should change and how. These processes govern the system's "behavior" and serve as the game's *logics.* As Noah Wardrip-Fruin argues in his foundation book, *Expressive Processing* [189], creating a model is by definition lossy process in which a complex phenomenon is distilled into its most salient features. The particular data and rules a designer chooses to include in the model make arguments about what processes or procedures are most important to the task being modeled. What a system models and how its processes create that representation is a key part of analyzing computational media.

Once the system updates its models via the enactment of its processes, it decides what to output to the player as feedback. Figure 7.2 offers an updated diagram of responsiveness.

7.1.3 How Systems Speak

Systems speak by issuing feedback to the player, which may potentially communicate new affordances. At the lowest level of abstraction, feedback takes the form of text output, visual changes to UI elements, feedback sounds, and so on. At a narrative level, feedback takes the form of dialogue, animations, cutscenes, etc.. Speaking is the most visible part of the LTS loop, designed to communicate relevant information to the player. Feedback has to convey the information the system wants or needs players to

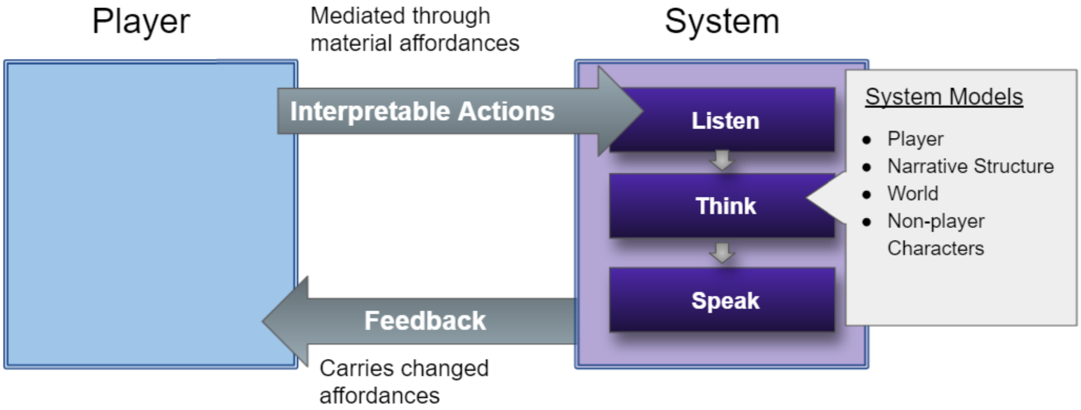

Figure 7.2: Diagram of responsiveness that depicts various narrative models at the think stage.

know. It is the part of the loop that lets players know that the system has heard their input, that it has (or has not) updated accordingly, and additionally informs the player of her next available options or actions.

One way feedback might communicate that a system has listened is simply by repeating key words or phrases back to the player. This technique is commonly used for cosmetic options such—such as a player selecting her character's name or gender—which usually do not offer changes to the computational model but do make players feel heard and can enhance their roleplay opportunities. *With Those We Love Alive* [71] uses reflective choices in this way to great effect, as we discussed in Chapter 6.

Similarly, narrative callbacks—moments where narrative elements refer to a player's previous choice—are very effective at surfacing places where the system has heard the player. As we discuss in Chapter 6, reminding the player of choices she has made can be an effective way to signal counterfactuality and let the player know that the game has heard her. However, diegetically reminding the player of her choices with dialogue can be awkward if not handled with care and variety.

In addition to letting the player know that the system is listening, feedback must also communicate state. Common forms of state communication might include visual effects to highlight resource or stat changes, pop-ups or toasts (floating/moving UI elements that signal change) to signal skill or relationship changes, visual elements to signal low health, etc. It might also take the form of characters behaving in certain ways toward the player to indicate their current feelings toward her.

And finally, feedback must also communicate affordances. The most direct form of this communication might come in the form of tooltips—a common form of formal affordance—that introduce or tutorialize material affordances to new players. Sound effects are often used to communicate confirmation or failure. An iconic sound effect from *The Legend of Zelda* [138] series is used as a formal affordance to indicate that a new material affordance has become available. Diegetically, feedback might come in the form of characters reacting to relationship changes, cutscenes playing after key choice moments, changes to the virtual environment, and so on.

7.2 Technical Interventions in the LTS Loop

Armed with a better understanding of how systems listen, think, and ppeak, let us examine specific technical interventions that designers can make at various points in the LTS loop to foster the affordances changes and feedback that signal responsiveness.

As we discuss technical interventions, we will specifically be looking at how the system's actions at different points in the LTS loop create the opportunity for responsiveness by hearing and interpreting player actions, changing affordances, and signalling these changes through feedback. Additionally, we will review some strategies toward technical interventions offered by different systems that have aimed to improve listening, thinking, or speaking, and how those systems have engaged with responsiveness. Finally, we will look at design considerations for each point in the LTS loop, which I hope will offer tangible takeaways for designers looking to design responsive systems.

7.2.1 System Listening and Responsiveness

Responsiveness is all about being heard, so naturally how well a system can listen is a huge part of that. If responsiveness is the ability for the system to change its affordances and feedback as a result of player actions, how it hears those actions and interprets them is critical.

One natural strategy toward increasing system responsiveness is recognizing a wider range of inputs from the player. After all, as we saw in our lightswitch example, when a system can listen for a limited range of inputs, the possible space over which it can change its affordances is usually very small. By contrast, a system whose material affordances support listening for a wider range of inputs—such as the assistant who listened for a pattern in lightswitch use—has a larger palette of possible affordances to work with. It can then use what it hears to make calculations and state changes in the think step that might ultimately lead to responsive feedback.

Simply offering more inputs to the player does not automatically create responsiveness. Consider an example of an interactive fiction parser (or any command line terminal, for that matter) that recognizes a huge range of commands, but gives the player no indication of what to type. A novice player has no formal affordances to advertise the material affordances to her and guide what she types. On the contrary, the formal affordance of a flashing cursor may suggest to her that she can type anything—a promise that the parser cannot uphold. As she tries command after command, and frustratingly gets the same output of "That's not a verb I recognize.", it is clear that the increase of recognized verbs does not automatically equate to a responsive experience. For responsiveness to occur, the input must change affordances and feedback *as a result of player actions.* Therefore if the system does not receive a player action that it recognizes, it cannot respond. If the parser additionally listens for repeated errors, and after several failed verbs suggests a hint to the player—thus offering a new formal affordance—a responsive interaction occurs. But that action is dependent on the system's ability to listen for the repeated failures.

The number and kinds of input the system recognizes is itself an affordance that colors the player's expectations. As soon as the player recognizes that the system is listening for a particular kind of action, she begins to expect that said action will have an effect proportional to the action. For example, if she recognizes that the parser understands the word "kick", she will then expect to be able to kick things, and for kicking things to elicit feedback that a thing has been kicked. She might also expect kicking to be an important action that might tangibly change stories, but such expectations will be influenced by her familiarity with conventions and her ability to map her mental model of this game's affordances onto other things she has encountered. As soon as we realize our smart assistant can detect patterns in when we turn on lights in the office, we might also expect it to detect patterns for the lights in the bedroom, and we might even expect it to detect patterns in the temperature controls.

The aesthetic experience of a larger number of recognized inputs is not automatically responsive, but it does act as a promise that the player has greater freedom. The flipside is that designers must deliver on that promise. When there is more the player can "do", there is more for the system to model in response.

7.2.2 Toward Listening Improvements

Advancements in natural language processing (NLP) and gestural input detection [152] over the last decade mean that systems can correctly identify speech utterances and gestures—two fundamental forms of input for embodied interactive experiences—better than ever before. One difficulty that arises, however, as this technology advances is how designers can leverage these technologies and integrate them with narrative and game-play models. A real difficulty arises in how to translate detection into meaningful understanding [176]. James Ryan's reductionist *Reductionist* framework [128], is an exciting step that is remarkable in its ability to generate semantic tags as it creates procedural text. A two-way approach as is described by Summerville et al [176] offers the potential for a huge range of inputs that are also semantically recognized with a system. Systems like Spirit AI leverage these emerging ML approaches toward a narrative experience [2]. The result is a parser that not only recognizes a huge range of inputs, but is also supported by a conversational model that contextualizes those inputs as meaningful.

7.2.3 Considerations When Designing Systems That Listen

As we move toward advances in recognition, detection, and input modeling—advances which will certainly allow for new kinds of experiences unlike those we have experienced to date—experience designers have many considerations to balance in designing systems that listen.

How much of the system to expose to the player is an evergreen concern, applicable to many parts of a system's design. As evidenced in the example of the interactive fiction parser that detects successive failures or the automated assistant who helps us with the lights, invisible triggers and verbs can be powerful tools toward responsiveness when applied well. In fact, a stated design goal for *Façade* was the inclusion of invisible conversational triggers feel entirely natural and promote narrative immersion, as in the example of Grace's comment on the decor in response to the player looking at the sofe [100]. Many moments of delight arise from unexpected invisible triggers, as we saw with *The Stanley Parable*. One benefit to unexpected invisible triggers is that they do not necessarily map from one instance to another; in other words, even though Grace comments on the sofa when the player looks at it, the player does not expect Grace to comment on everything she looks at.

As systems' abilities to recognize greater numbers of inputs increases, designers must be aware of how these inputs translate to formal affordances for the player. In deciding how much listening to advertise to the player as a formal affordance, cognitive load is something to keep in mind. A naive solution to a player not knowing which verb to type into a parser might be to advertise all possible verbs, however this creates an immense cognitive load for the player. Indeed, deliberately overwhelming the player with too many visible choices is a strategy used to great effect to deliberately provoke anxiety [97] in Dietrich Squinkifer's *Imposter Syndrome* [166]. By making each choice visible to the player, Squinkifer provides a huge number of conflicting formal affordances that "cry out" to be taken—the effect is that the player is pulled in many directions simultaneously. If a designer makes choices less visible, she potentially circumvents this anxiety by reducing the number of conflicting formal affordances calling out in different directions, however reducing formal affordances also obfuscates the workings of the system for novice players. A common way to reconcile this tension is to offer few formal affordances at the start of a game, and increase the offerings as the player achieves mastery.

In our attempts to listen and to shape the kinds of things the system can listen for, we must always keep the player's agency at the forefront of our designs; formal and material affordances must always be in balance. We must not overwhelm the player with too many switches and no reason to pull one over the others, but we also must be careful not to give her switches that do nothing. As the ability to hear more and more forms of input increases, designers must be aware of how inputs are advertised to players through formal affordances and how they are modeled in the system as material affordances.

7.2.4 System Thinking and Responsiveness

As we saw in the previous section, the system's ability to listen effectively is dependent on its ability to translate the inputs it recognizes into meaningful understanding. Whether its understanding is meaningful is dependent on its underlying model and what the system does with recognized inputs. In other words, what changes as a result of the recognized input? What data and processes should the designer include in the model?

As designers, we ask the question slightly differently in the process of creating the system, actually breaking it into component questions:

1. What should change as a result of player actions?

2. How should it change?

3. Under what circumstances should it change?

We saw in section 7.1.2 that narrative systems think in terms of models of both the player and the story. We might break "story" down further into component models of: the narrative world, the narrative structures, and the characters in the world. Taken together, the system is engaging at least all of the following models:

1. **Models of the player.** This might include a model of the player's character, a model of the player herself, or both as discussed above.

2. **Models of story structure.** This includes things like experience managers, planners, and other systems that decide which content to display next

3. **Models of non-player characters.** These govern the state and behaviors of characters in the story.

4. **Models of the world itself.** These might include things like resources necessary for gameplay, the physical state of the environment, etc.

As with all of the models presented in this dissertation, I present these here as a clean division, but real-world models overlap and intermingle these. For example, characters need data and processes to govern their behavior, but they also exist in the world environment; they are not cleanly separated from it.

Still, each of these is a component for which a designer needs to answer the three questions at the beginning of this section. Additionally, she must also answer these questions over the range of possible nested levels of interaction.

7.2.5 Toward Thinking Improvements

I mentioned at the outset of this dissertation that some of the research efforts of the narrative AI community that claim local improvements can lead to more narratives that are more tailored to players. Though they offer solutions to problems that are firmly situated in the literature as improvements to one of the models listed above, often they fail to recognize the gap between such improvements and the dependent design questions needed to bridge the gap between a technical system and the concerns most important to practicing designers. The strongest research efforts are those that can make the connection between their own work and how that work enables a new kind of experience, even if that experience is years away from being realized in practice.

7.2.5.1 Improved Player Modeling

For example, some researchers offer that more accurate models of player taste could lead to the ability to adapt stories to be more amenable to player preferences [183] [144]. While player modeling is the foundation of data science in the games industry, and some very impressive modeling [9] [194] happens in the context of ad monetization, adaptive difficulty tuning, and matchmaking using statistical classifiers and machine learning algorithms, efforts to integrate similar techniques into player modeling for narrative often feel disconnected from the realities of authoring the systems that would take advantage of such advancements. Even if we could perfectly predict what a player wants or what they will do—how does this help a designer answer the questions above? What kinds of experiences does this open? One of the reasons I think other applications of player modeling are so successful (for better or worse), is because their models are clearly optimizing toward gaining particular insights with a clear understanding of how those insights will influence game design. I would love to see more conversation in the academic technical research community around how new approaches might be used by designers.

These areas of research are an interesting first step—a piece toward solving larger problems—but they would be stronger if they understood the problem they are trying to solve not as a financial one [183] but as a component problem toward a particular experience, even if we do not yet understand what that experience might be.

Conversely, more research into new techniques for modeling a player's character rather than the player herself could be an interesting area of opportunity for AI researchers.

7.2.5.2 Improved Models of Narrative Structure

The research space around generative interactive narrative systems is vast and varied. Many of the approaches to the generation of procedural narratives have arisen from different traditions and are interested the the pursuit of vastly different goals. [142]. As Cardon-Rivera notes, "there is the potential to conflate procedural *interactive* narrative generation with procedural narrative generation" [25] (emphasis added), particularly when discussing specific design considerations for modeling narrative structure. However, while interactive narratives can build upon static narrative generation techniques, the generation of interactive narratives carries fundamentally different system design considerations if it is to generate a well-designed player experience. Specifically, such an experience requires the centering of player actions, affordances, and agency in the system's conception.

An interesting dimension of narrative events is that they can serve as both input and output for the player; they carry formal affordances that set the premise and motivation for the player's actions [26], but they also serve as feedback and reward for the action the player has just taken. Narrative structure must take into account how intertwined with player motivation and affordances it really is.

Some of the best research into models of narrative structure have approached the structure of interactive narratives specifically as a computational phenomenon. [62] [83] *Façade* side-stepped issues of agency as a force in conflict with well-formed narrative by reevaluating its working model of dramatic Aristotelean agency through a computational lens [99]. This approach demonstrates the level of narratological depth necessary to such an approach, and the results were clearly successful.

Any responsive narrative will need to understand how the structure of its narrative changes with player input and how narrative events interface with changes to affordances.

7.2.5.3 Improved Character Modeling

Characters are the soul of storytelling. We derive meaning and connection to narrative through characters, and it is no surprise that they have been an active area of research in narrative technology [142]. Thus, it is no surprise that improvements to character modeling—whether in the form of agent believability, character autonomy, or social simulation. As with other types of narrative model discussed above, character models, in order to contribute to responsive narrative, must understand their role in affordance changes.

From a narrative design standpoint, characters in games serve two simultaneous functions. First, they convey story, giving the game depth and relatable meaning.

Well-written characters raise the narrative's stakes and increase emotional investment in the narrative. Second, characters play a key role in delivering formal affordances. In the worst versions of this, NPCs primarily serve as diegetic dispensers of player objectives, strategic hints, and resources such as items or currency, a convention that leads many players to treat NPCs more as objects or resources than as full-fledged characters. When critics claim that narrative and gameplay are at odds, I believe the dual role NPCs play in game narratives, coupled with the difficulty of successfully marrying the two roles harmoniously is one of the key sources of that tension.

A key of the reason I suspect works like *Façade* and *The Ice-Bound Concordance* [139] have been so successful both as (1) an engaging story and (2) as playable experiences that facilitate agency is that they remain deeply aware of the need for characters to serve simultaneous functions narratively and systemically [100]. ABL, the bespoke character behavior language on which *Façade* is built, is specifically designed to perform both tasks: it understands that characters must introduce new avenues of drama into a scene as well as offer signals toward what the player should do next [103]. Indeed, the seamless overlap of the roles of characters in games and in storytelling is a primary goal for narrative designers.

Bad News [148] is a live-action piece that generates characters and social dynamics, from which a live actor selects interesting narrative bits and uses them to seed ideas for roleplaying situations with players. Though the the generated histories are not themselves embued with the warmth we might expect from a more character-centric embedded narrative, the presence of a live actor injects the characters with depth and personality. The actor bridges the narrative gap and serves both functions: he facilitates offers a human interest and identification to the story and also offers formal affordances to players to help guide them in their task.

7.2.5.4 Improved World Models

World models account for things like terrain, resources around which the player strategizes, and systems external to players and characters.

Simulation games have gotten very good at simulating environments, terrain, and weather through intricate world models. Minecraft is one of the most popular videogames of all time, and much of that popularity has arisen due to how it allows players to explore and strategize around the components of an environment model. Games which feature simulation-manipulation as a key pleasure-experience often feature intricate world models which provide fodder for the combinatorial affordances they construct.

As such, I am aware of less research in narrative technology that focuses on improvements to world models outside of social simulation. Thus how environment models interface with narrative models feels extremely under-explored, particularly when you consider how often environments are used as a narrative shorthand in games. In games where the pleasure-experience focuses on fast-paced action gameplay, terrain is often the primary means of conveying narrative. Yet we do not see very many systems integrating environment models with responsive narrative. Even fewer games utilize such models to change narrative affordances or feedback. Games that have radically changed the affordances of their narrative world such as *Blaseball* [8] or *World of Warcraft's* Cataclysm event [122] have been noteworthy, and even still World of Warcraft's affordance changes did not happen as a result of player actions and thus did not happen in the context of responsive narrative.

7.2.6 Considerations When Designing Narrative Systems That Think

7.2.6.1 How much to model?

When designing any model, a key question a designer asks herself is how much to model. By definition models are a distillation of a complex phenomenon into its most salient and necessary components. Figuring out how much to model is a deceptively difficult first step toward responsiveness.

Each component of the model that can change as a result of player input increases authorial burden, potentially to intractable degrees. Additionally each aspect of the model that the player understands as an affordance becomes a promise upon which the player will depend. Conversely, the more components exist and move in interconnected ways, the more possibilities for emergence and combinatorial affordances present themselves. If affordances are going to change, movement in the model is what facilitates that change.

7.2.6.2 Pay off the player model

If a player model is going to reinforce narrative responsiveness, the designer should understand the role the model plays in how affordances change. For example, if a game offers skills as part of her stats sheet (a dimension of a model of the player character), the game must also understand how those skills can interplay with other models to change affordances.

Any dimensions of a player model that are visible to the player are opportunities for roleplay. When designing for responsiveness, designers need to understand that the formal affordances of the player model should be a skeleton, not a cage. Designers need to provide numerous enough formal affordances (or roleplay hooks) that players understand the character they are playing and can form intentions about what that character might do, but while still leaving room for the player to inject herself into the role.

7.2.7 Keep NPC Models Focused

Humans are complex, interesting creatures, and as such many designers' first impulse is to capture various aspects of that complexity and richness in NPC models. However, as anybody who has ever failed to beat an AI opponent can attest: the goal with NPC modeling is not to create "more skilled" models of NPCs or even "more accurate" ones. In real-world design, NPC models almost always interface with feedback and exist to communicate something to the player.

Distilling NPC models into their salient features is a must for NPC modeling. Interactions with characters are a key way that narrative games offer formal affordances to players, and generative solutions need to keep this in mind when proposing new approaches to character behavior. Overly-modeled NPCs have the potential to offer terrible signal-to-noise ratios if players are e.g., trying to engage with NPC dialogue to discover where to find the next quest. Additionally, certain types of modeling can get expensive, such as reasoning over larger or complex social graphs. Designers need to understand how the features of their NPC models interface with other narrative systems to provide changing affordances and salient feedback.

7.2.8 System Speaking and Responsiveness

In many ways we have been answering the question of how the system's output conveys responsiveness over much of Chapter 6. However, there is a slight distinction between the question of what the player recognizes as responsiveness and the processes the system must undertake to convey it. In other words, we have been discussing *what* the system outputs, but now I would like to discuss the *how*.

As mentioned in Section 7.1.3, when systems speak, they output feedback that potentially communicates one of the following:

- Recognition of input
- Communication of state (change)
- Communication of affordances

At the lowest level, this is easy to recognize as recognition of action (e.g., a mouse clicking a button changes the button color to indicate the action has been recognized as meaningful), communication of success or failure of the action, and an optional communication of what to do next. But at high levels of abstraction, recognition of input becomes a harder problem.

Agency is felt through feedback, and responsiveness cannot happen if there is no response or if the response is too underdeveloped to provide feedback at the level of abstraction at which the player is concerned. The system need not tell the player "I'm employing the Fisher Defense", but the player needs to be able to recognize it in order to appreciate that the system has responded to her King's Gambit. While recognizing how to give feedback at a lower level of abstraction is fairly well-understood, how to give feedback toward more abstract recognitions is riskier [161].

To communicate responsiveness, a system must be able to communicate the principles of the aesthetics of responsiveness outlined in Chapter 6. It must offer tightly-coupled feedback and signal that affordances have changed, ideally while also signalling counterfactuality and doing these at the highest level of abstraction it can.

7.2.9 Toward Speaking Improvements

How a system communicates to a user has long been an area of interest among designers [125]. The HCI and commercial game industry has decades of experience communicating affordances through feedback, even if they were not necessarily using that terminology [66]. In many ways, this component of the LTS loop might be the farthest along when speaking about lower levels of abstraction.

Still, how narrative systems communicate feedback to players is an area ripe for study, particularly at higher levels of abstraction. While systems like *Curveship* [114] affect the telling of narrative events, they do not, for example, affect the way in which the system communicates its state changes. In other words, they affect *how* something is told, but they are not changing *what* is told—the level of informational depth or complexity, for example—as a result of player action.

7.2.10 Considerations When Designing Systems That Speak

In addition to how much to model and how those models should function, designers must also decide how much of the system's workings to surface to players and how strongly to signal formal affordances. This is a difficult problem for reasons similar to those discussed in Section 7.2.3 on designing for listening. Designers will generally want to avoid the Eliza Effect and the Talespin Effect in their models [189], and the balance of affordances to maintain agency remains forefront in their minds. Often the intended genre and pleasure-experience of a work might help guide the answer to this problem—with more mechanic-focused pleasure-experiences tending to signal their systemic workings more strongly while narrative immersion-focused experiences tend to obfuscate their systemic workings—as we discussed in section 6.2.7.

Chapter 8

The Lume System

This chapter details the technical specifications of Lume, an interactive narrative engine that Ceri Stagg and I developed in collaboration. Our goal with the design of the system was to facilitate a responsive narrative, and we approached this usign some of the design considerations that I laid out in the previous chapter. In this chapter, I discuss how the system's technical design tries to capture some of these technical design strategies.[1]

8.1 Motivations

8.1.1 Initial Motivations

Initially when we started building the prototype system that later grew into Lume, Ceri and I were trying to build a choice-based narrative system that captured some of the dynamism of combinatorial narrative approaches while maintaining the

[1]Much of this chapter is adapted from Mason & Stagg et al. [98]

coherence of well-authored embedded narrative experiences to the quality and emotional resonance of Telltale games. We were inspired by sculptural fiction [140] [15], emerging storylet models, and were particularly inspired by the mobile game *Reigns*. In *Reigns* [121], the player takes the role of a ruler who must make yes or no decisions to balance various stats to keep the health of her kingdom in balance. If any stat goes too high or too low, the player dies a comical death and the reign of the next ruler starts. We saw a lot of potential for something with the quick-cadence choices and stat-balancing of *Reigns* but with a much richer embedded narrative, more character development, difficult emotional choices, and more robust political relationships among the cast.

8.1.2 Related Work

Much work on procedural narrative generation has approached the problem by applying various models of narrative from narratology research to generate completed stories [85]. While this research has yielded interesting results, many of the concerns and innovations in those systems are aimed at trying to generate completed (static) narratives with salient plots. Though plot is a useful and important element of rich storytelling, Lume's goal is not to generate a plot arc autonomously, but rather to surface content strategically in response to player interaction. Thus the work done specifically in the space of interactive storylet approaches and procedural systems, and those that surface salient narrative based on player input, prove to be apt starting points for comparison. One of the foundational academic descriptions of a procedural storylet system is Mark Bernstein's Card Shark [15], a system that "begins with a set of lexia, all of which are connected to each other, and builds structure by removing unwanted connections" through a process that Bernstein and Greco term sculptural hypertext. The promise of sculptural hypertext—which Aaron A. Reed examines and folds into his broader definition of "sculptural fiction" [140]—is one of assembling a narrative through selection from a database. Yet Bernstein acknowledges,

> Where a sculptural strategy has been employed in the past—most notably, perhaps, in Malloy's (1993) Its name was Penelope and in Malloy and Marshall's (1996) Forward Anywhere—it has been chosen in part to deemphasize temporal sequence and narrative struc- ture (Golovchinsky and Marshall 2000).

Bernstein and Greco cite earlier attempts at sculptural fiction whose selection process was largely random. Unlike these, however, Card Shark selects eligible lexia based on the state of the world—that is, it selects relevant content—and presents all of those relevant options to the reader, who then decides which lexia to visit next. This approach forces readers to strategize about their own traversal through the work, making a decision as to which content they would like to experience next. In contrast, our system presents only one of the (potentially many) viable lexia to players in an effort to keep players invested in the role-playing opportunities the narrative provides (through diegetic agency). It never asks them to break that role by operating as their own narrator and providing extra-diegetic agency [95].

Like the works Bernstein and Greco cite, many popular games today employing a sculptural approach—most notably, the very card-like *Reigns* [121]—provide something of a loose emergent narrative. Yet this narrative is typically an after—thought to the ludic pleasures of stat-balancing. Little attention is paid to character arcs and development, scene length and pacing, rising and falling dramatic tension curves, and the causal relationship between events. The last of these is perhaps one of the most important considerations for interactive storytelling that seeks to create and maintain a sense of players' narrative agency.

Mateas and Stern [104] note that most interactive narratives fall into two general approaches: a more hand-crafted structure of nodes—such as the plot structures common to action/adventure games, hypertext fiction and hypermedia, and choose-your-own-adventure books—and a more procedural simulation approach that encourages emergent narrative, such as those found in sim games, virtual worlds, or the narrative created in the recounting of a sport or eSport match. Their technical approach to *Façade* sought to find a middle ground between these extremes through a drama management system that selects fine-grained narrative beats appropriate to particular narrative contexts. While *Façade* is widely recognized as a successful experiment in very reactive narrative, the cast was very contained, thus beats were not parameterized to be applicable to multiple characters in multiple narrative contexts. In contrast, parameterized beats applicable to a larger cast was one of the major design goals of *Prom Week* [109]. Yet while it was successful in creating dynamic emergent narrative events, these events lacked a clear, recognizable narrative structure. Our system aims to take from the best of both approaches, using parameterization that makes content more applicable and reusable across different points in the narrative, but also provides the narrative structure to offer stronger signaling of causal relationships.

The idea of building a parameterized storylet system is promising, and research into storylet systems is being pursued on multiple fronts. StoryAssembler [62] is an exciting approach that, like Lume, selects parameterized storylets from a pool of content. However, StoryAssembler uses a hybrid-planning approach (in contrast to our bottom-up logic-programming approach) that is organized around explicit authorial goals, using the concept of goal as a global mechanism for authorial control over narrative progression.

8.1.3 Motivations Toward Responsiveness

At the start of the design process we did not yet have a clear formulation of the theoretical model of responsiveness, but rather we were trying to design a system that would improve the player's sense of agency. We initially took a naive approach to this, figuring that a combinatorial card-like structure with some light narrative scaffolding would naturally lead to a greater sense of agency through the interplay of choices, stats, and preconditions for our storylets. Development took a very bottom-up approach, first designing a series of storylets and then trying to add coherence. Our earliest storylets were, like *Reigns*, largely stats-driven. When played together, they did not feel like much of a story . To increase narrative coherence, we started by largely looking to make interventions in the "Think" step, though we did not yet have a sense for the model of responsiveness presented in this dissertation. As I have suggested earlier, the iterative process of developing a system toward a certain aesthetic goal and then assessing the theoretical underpinnings of the goal itself led to the simultaneous development of both the theoretical model and the system.

We knew that our ability to stitch our storylets into a narrative would largely depend on how we decided to model the story; our approach was to start with the minimum narrative scaffolding necessary and add additional story structure onto our storylets as necessary. Our instinct was to find a narrative approach that took the narrative dynamism of the storylets and enforced narrative cohesiveness through external constraints. As we worked toward making out storylets feel narratively coherent, we realized that like *Reigns*, the system largely produced games that felt very reactive—which is to say the player was not actually exerting much agency, she was mostly just reacting to situations at hand as they arose.

In the process of prototyping various designs to surface consequences, we realized that we were able to surface event histories and real magic happened when we did. Very simple callbacks allowed players to understand the impacts of their decisions and draw causal relationships between storylets. This led to a realization: we knew surfacing system elements was important but did not realize how impactful just repeating a player's actions back to them could be. Experiments ensued. We realized we also needed to improve communication of various aspects of the system to players—that is we needed to improve how both feedback and affordances were communicated in the "speak" step. Around this time, I realized that we were no longer actually optimizing for agency per se, we were instead trying to maximize the system's ability to respond to the agency the player already had. From these realizations, the beginnings of the formation of a theoretical model of responsiveness started to take shape. And as more and more rule-systems to enforce coherence were added—including systems to enforce

NPC believability and narrative causality—the system presented here started to take shape.

8.2 System Overview

8.2.1 Narrative Model

Our concept of narrative is derived from a popular theory of narrative practitioners, conceiving of narrative as a series of interconnected moments [112]. Whether creators organize these moments into scenes, beats, film shots, or pages, we experience narrative linearly in discrete moments. We expect these moments to feature consistent characters, and to be related causally. These expectations are so strong that consumers of emergent narrative will often read causality into sequential moments, even where none is explicitly shown or modeled [163]. In general, storylet systems promise an alternative to handcrafted branching plots and emergent simulated plots. The narrative vignettes themselves may be as well-formed as the author likes, while the procedural combination of these vignettes can lead to the emergence of overarching narrative arcs. The creation of a well-formed narrative arises from selecting appropriate moments into appropriate slots, and the more specific the selected narrative scenario is to the current world-state, the more responsive the narrative should feel to the player's input.

Our hope for the Lume system was always to strike a balance between what we call *narrative dynamism*, the feeling that what is happening is one of many paths that could have been taken, and *narrative coherence*, the feeling that narrative events are developing causally from player actions or from logical NPC reactions. The Lume system offers a hybrid approach in which individual scenes constitute node trees, but the selection of the scene to be displayed next is left up to the system. By building on the capabilities of logic programming approaches to content selection, we are able to construct a coherent narrative through a series of moments appropriate to the given context. In addition, we parameterize several elements of the selected scenes themselves in an effort to allow both content reuse in more places, and (more importantly) a greater sense of narrative responsiveness, as more content is able to be catered to the player's past decisions. Thus the narrative experience is shaped by logical rules authored by creators, but the actual narrative emerges from the player's decisions.

Initial ideas toward stitching together disparate content units into a coherent emergent story were first inspired by *Façade's* beat system [100]. However, by starting from a bottom-up approach, we did not actually encode a top-down drama manager to manage individual drama metadata to enforce something like a smooth Aristotelian drama curve [99]. Instead we settled on a much looser fires-in-the-desert [54] narrative structure, which I will discuss more in the case study of *Rumina Woods*, a prototype game we have authored with this system.

One goal of the system was to modularize components as much as possible. Having worked at companies where an over-optimization in the tool pipeline eventually led to a stagnation of innovation, we favored an approach that ensured the core logic of Lume was as lightweight as possible, and things like higher-level narrative structures, or various game-specific functionality (e.g., items, while supported at the necessary lower levels, were created as modular components or left to individual games and authors. As such, many of the features of the core Lume functionality are fairly abstract, and more specific functionality is left to authors to determine.

8.2.2 Scene Node Trees

Scenes are the units of narrative moment we use as our storylets. The system presents a scene, and based upon the outcome of that scene, it updates the state of the world before determining which scene to present next.

Scenes are composed of a tree of nodes: a mandatory base node, which typically presents the premise of the scene and introduces initial characters, and optional child nodes for player choices or variable beats within the scene (discussed below). These nodes, taken together, form a tree that varies the outcome of any particular narrative moment (see Figure 8.1). Our procedural scene design provides a way to hybridize the branching node structure commonly used for well-formed interactive narratives—such as those found in interactive narratives with authored choice-paths—with the huge possibility spaces of more proceduralized approaches such as those found in simulation games.

8.2.2.1 Node Types

Base Node Each scene has a Base Node as its root. The constraints of the base node must be satisfied in order for the scene to be eligible for selection (and child nodes may impose additional constraints). All base nodes contain the following components:

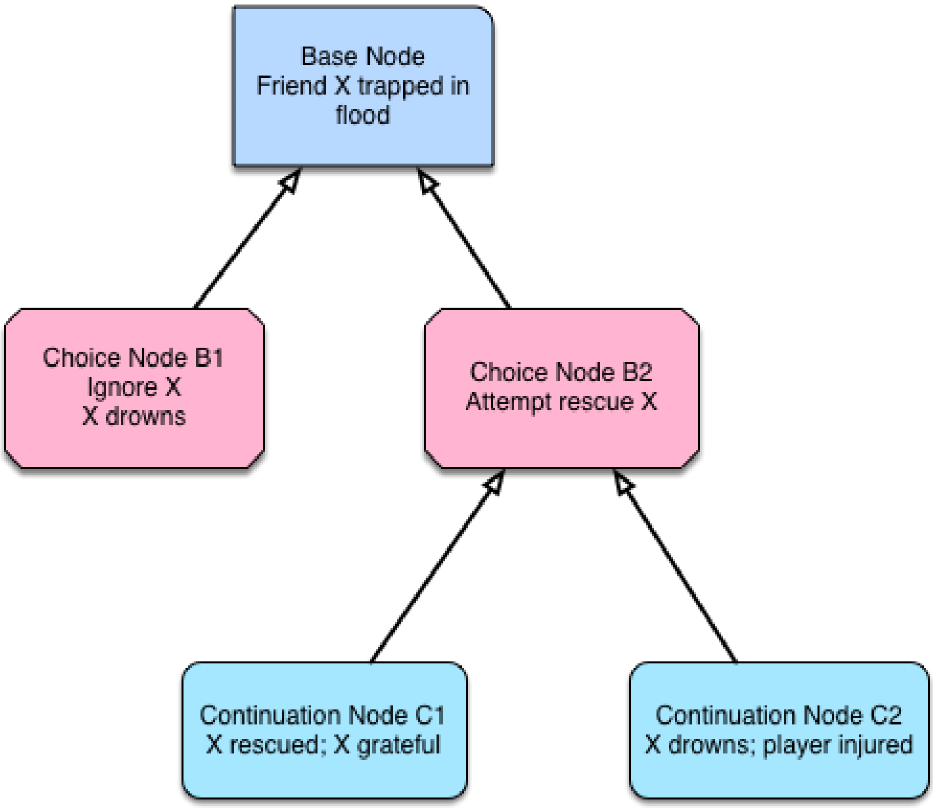

Figure 8.1: Example of a scene node-tree

1. *Preconditions* are the rules that must be satisfied in order for a scene to become eligible for display. Preconditions are specified by narrative designers. Examples of preconditions include:

    - Player is in The Woods
    - It is morning
    - A cautious NPC is in your party
    - The player just changed locations
    - The player has killed an NPC

- A Traumatic Event has occurred

2. *Bindings* are the system's way of remembering which character appeared in which scene. Since scenes are authored in a parameterized manner that leaves slots to be filled by narrative features such as characters, places, or flags. These bindings allow each slot to attach to specific eligible objects in the game world. When a scene is selected, the objects that fit the preconditions are bound to the scene and will be remembered in the event log. For example, a scene's preconditions may specify that a scene takes place between the player and an NPC in the party. The system would then try to use Sara (and then Mehmet, and Alice, etc.) as a candidate character for the NPC binding slot for this scene. It may succeed with some, all, or none of the candidates, and will return scene trees for all successful binding sets and then select one. If a scene has binding requirements, and the bindings cannot be satisfied by available characters/objects, the scene will be ineligible. Bindings particularly leverage Prolog's constraint-solving capabilities: given the current state of the world and the constraints specified, Lume will automatically try all candidate objects in each slot, trying to build viable nodes and trees that work.

3. *Instructions* are the output of the scene, either in the form of display text, functional tags, or markup to be interpreted and used by other engines. *Instructions* may be simple strings of static text, but in practice, the system is most powerful when instructions take the form of DCGs (Prolog's definite clause grammars)

which allow for text to be generated dynamically. Bindings may be passed into an instruction's DCGs to modify the text dynamically. An example instruction:

```
Instructions = [NPC1, "never loved that ",
[insultnoun, foolish]," ",NPC2]
```

might be presented as "*Sara never loved that oaf Ralph*" if Sara is bound to `NPC1` and Ralph is bound to `NPC2`, and "*oaf*" is one of the possible expansions of `[insultnoun, foolish]`.

4. *Postconditions* change the state of the world in response to the scene presented. In our system, "state" is stored as a single long list of events that have occurred in the world, and a scene's postconditions operate by appending events to this Event List. No state exists outside of this event list. Using this approach, we track not only the current state of the world, but also all of the previous states and events. Thus we can search the list for the most likely event to have caused major relationship changes or for any similar turning point. We can then refer back to these moments using Recall Phrases.

5. *Recall phrases (optional)* generated by a scene are sections of markup text that refer back to the events that took place in that scene. We might want to refer to a moment in which a major relationship change happened, a character died, a building burned down, or any other significant story event. Recall phrases allow us to query the system not just for a particular major event (e.g., the time you killed John), but for any event that fits a designer's criteria (e.g., the last time X's opinion of you changed for the worse).

In practice, different characters will refer to past events in different contexts, some positive and some negative. To handle these different tones, the system supports long (default), short, and negative recall phrases. For example, if the house burns down, we would attach recall phrases to that scene:

- Long recall: the house burned down
- Short recall: the fire
- Negative recall: you let the house burn down

An author might write

> "Sara mumbles, 'Things haven't been the same
> between us since `[ShortRecall]`.'"

The author can then link `[ShortRecall]` to a time in which a major event lowered your relationship with Sara. If the player has seen the fire scene, the line will replace with

> *"Things haven't been the same between us since* ***[the fire]****."*

If, however, a different major event lowered your relationship, Sarah might instead say "Things haven't been the same since the affair." If the line calls for a more aggressive accusation, we might choose to use the negative recall instead:

> *"Forgive!? Of course I can't forgive;* ***[you let the house burn down]****!"*

Recall phrases are entered by a human author, but may be further parameterized by DCGs (see section on DCGs below). They may be attached to a base node (e.g, “that time the church flooded”) or may be attached to choice nodes (e.g., “that time you chose to leave George in the flood”) or other child nodes.

Choice Nodes In addition to base nodes, scenes may optionally contain choices for the player to make. Each choice and its subsequent feedback to the player is contained within a choice node. Choice nodes have all of the same components as base nodes—preconditions, bindings, instructions, postconditions, and optional recall phrases—with some slight distinctions, detailed below.

1. *Preconditions* Choice nodes may have their own preconditions that must be satisfied in order for this particular choice to be presented to a player within the Scene. Unlike base nodes, if a choice node’s preconditions are not met, it does not necessarily invalidate the Scene’s eligibility (see section on Scene Eligibility).

2. *Bindings* Bindings are inherited from the base node, but choice nodes may specify additional bindings if needed. (These are inherited in turn by any children of this choice node.)

3. *Choice Text* This is the text displayed to the player before a choice is selected.

4. *Instructions* A choice node’s instructions are presented/executed if the player selects this choice. They usually take the form of feedback text offered to the player.

5. *Postconditions* A choice’s postconditions are applied after the player makes this particular choice (e.g., Penny’s relationship with the player increases by a large

amount as you chose to rescue her from the burning garage), whereas a base node's postconditions apply no matter what choice is selected (the garage has been destroyed).

6. *Recall phrases (optional)* Recall phrases attached to choice nodes refer back to the exact choice a player made in a given situation. Though recall events may be applied to any node, they are most powerful when attached to choice nodes.

Continuation Nodes Continuation nodes are used when the scene continues without a player choice. Continuation nodes may have any of the other node components. One useful way to use them is to specify multiple continuation node children, and to attach preconditions to each based on things like characters present in a party, flags raised by other scenes, etc., to branch the scene based on context. Another useful application is to place die-roll preconditions on continuation node children to add variability and unpredictability to outcomes after a choice.

8.2.2.2 Scene Eligibility

Scenes are eligible if the preconditions of the base node are met and any specified bindings can be filled. Whenever a scene node has choice nodes as children, at least two of those children must be valid for the node itself to be valid (i.e., the choice nodes' preconditions are satisfied and bindings can be filled). Likewise, if a node has continuation node children, at least one child must be valid for the parent to be valid. A node-tree is viable if the following are true:

1. All child nodes of any given node are of the same type
2. If a node has choice node children, it is only viable if at least two children are viable
3. If a node has continuation node children, it is only viable if at least one child is viable

An example node tree is depicted in Figure 8.2

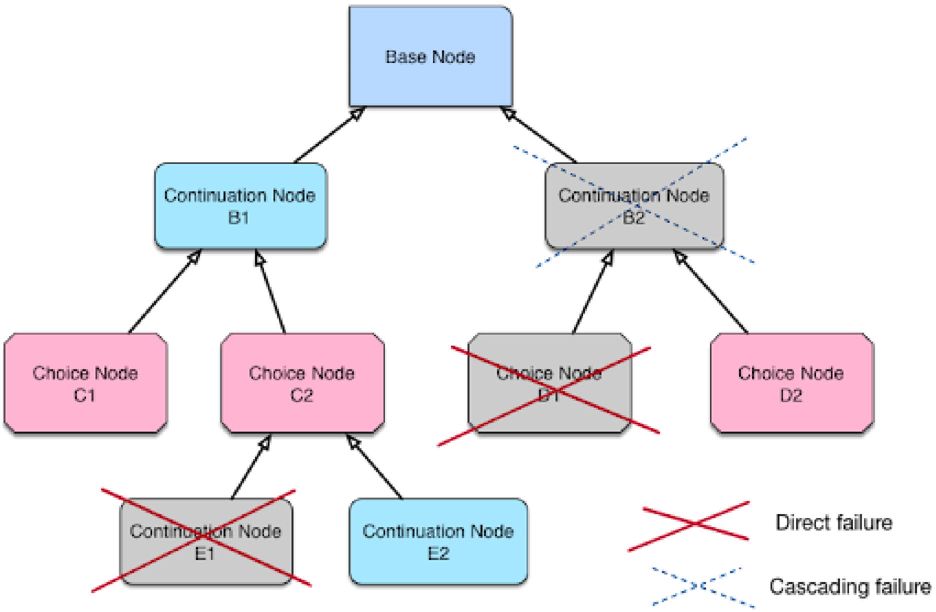

Figure 8.2: Eligibility in Lume scene node trees. This is a valid scene, because at least one continuation or at least two choice nodes are available at each level, thus a viable scene can be presented.

8.2.2.3 Scene Selection

Presuming equal weighting, scenes are drawn at random from the list of currently viable scenes, and authors may add weighting values to scenes to make them more or less likely to occur and preconditions to make certain scenes viable only in a particular context. Additionally, Lume offers some built-in mechanisms to assist with scene selection and repetition suppression:

- Scenes may be given a priority by their author. Viable higher priority scenes will be chosen ahead of any lower priority scenes.
- If two scenes have an equal priority band, the system will prioritize the scene that has been shown fewer times to the player, and will deprioritize scenes that have been shown recently.
- In practice, scenes with tighter preconditions—that is, scenes which are more specific to a particular event needing to have occurred before—are more likely to fire than common ones.
- If priority bands are equal and the player has encountered two scenes an equal number of times, the system will select one at random
- Authors can force scenes on a given turn, e.g., "Always fire Scene X as the 4th Scene in the game," or "Fire scene X on the turn after this one if it's available".
- Authors can bucket or group scenes, and can specify that scenes from a certain bucket should play on a specific turn or next or in n turns.
- Authors can combine tight preconditions with high selection priority to make a scene extremely likely to fire if its (rare) conditions are met.

Additionally, Lume features a pre-fetch system that caches all possible solutions a few nodes ahead for faster transitions from one scene to the next. Because of the possibility for player actions from a gameplay client outside of the Lume system, occasionally the system does need to recalculate the prefetch cache at major actions, however so far, in practice, these moments are infrequent and predictable.

8.3 Procedurality Within Scenes

While the Lume system offers highly variable plot structures through the ability to pull appropriate scenes from the pool of eligible scenes in any given moment, the true power of the Lume system comes from the fact that many elements of scenes are parameterized. Parameterized character bindings are extremely powerful when combined with conditional generative text. This section discusses additional features that increase combinatorial potential within scenes.

8.3.1 Definite Clause Grammars

The Lume system takes advantage of Prolog's built-in Definite Clause Grammars (DCGs), which offer a means to implement context-sensitive grammars. Context-sensitive grammars are able to carry out all of the functions of context-free grammars—structures widely used in other successful text-generation tools [33] [149]—while also allowing for more specific expansions of text based on context-specific information. In practice, this means that Lume authors can provide preconditions for text-expansions that allow for text to be expanded in certain ways if certain preconditions are met.

Additionally Lume allows for certain DCGs to override others, giving authors the opportunity to impose hierarchy and specify which DCGs they would prefer to be expanded if certain conditions are met. For example, an author could specify that a `[Greeting]` by default expands to either the phrase "Hi." or the phrase "Hello." But if the `[Greeting]` is performed by an `AngryCharacter`, override those expansions with "*How could you do this?*".

By default, eligible DCG terminals are selected at random with an equal distribution. Thus in our example above "Hi." and "Hello." would have an equal chance of selection. Authors can "weight" the expansions in certain directions by adding die rolls as preconditions.

As with other text-generation tools, DCGs can also combine to make up other DCGs. Thus we could also author `[SayName]` to expand to "My Name is `[CharName]`" or "I go by `[CharName]`" and then author the DCG `[Introduction]` to expand to the combination of {`[Greeting],[SayName]`}.

8.3.2 Pronoun Replacement and Point of View

To further parameterize the content of scenes, Lume implements a pronoun replacement system that allows authors to write text (with markup) that is more broadly applicable to multiple character bindings. Thus rather than authoring

"John killed Jane! She had it coming."

the author would instead write

```
[NPC1, " killed ", NPC2, ".  ", NPC2, " had it coming."]
```

The system interprets that in this case, `NPC2` is a woman and substitutes the appropriate pronouns. Note that the first time `NPC2` is called, it expands to `NPC2`'s name, and the second time, it expands to a pronoun. In the case in which the bindings are reversed, and `NPC1` is Jane and `NPC2` is John, the pronouns substitute appropriately. Additionally the Lume system implements a dialogue function that takes the speaker into account. Thus, if we pass the same marked up sentence into some dialogue in which John is saying the line, he will appropriately declare:

"I killed Jane. She had it coming."

This functionality is particularly useful for recall functions, in which authors can specify "`Player` killed `NPC2`" and different characters might approach the player with an accusation "*We all know [you killed Jane]*" or the player might come upon two characters gossiping "*Did you hear [John killed Jane]*?" The pronoun and point of view functionalities ensure that one piece of content can be used in as many places as possible, increasing the potential for content reuse to limit authorial burden.

We can also add recalls as anecdotal descriptions of characters. So if the player does not know John but witnessed the murder, the next time the player sees John, he can optionally be described as "the guy who killed Jane" if the author chooses that over his standard description or name/pronoun substitution.

The pronoun replacement system helps reduce the authorial burden inherent to generative systems, since authors are able to write a piece of dialogue once and reuse it multiple times across contexts. Unfortunately, it does not currently support automated fixup of pronouns that require alternate verb conjugations (e.g. they), however this is future work in which our team is deeply invested.

8.4 Recalls and the event list

In practice, the above example would likely actually be written as a Recall situation. In a node in which `NCP1` kills `NPC2`, an attached recall description might be written as:

```
[NPC1, " killed ", NPC2]
```

Lume would then store the scene's bindings, indicating that when this scenario fired, John was `NPC1` and Jane was `NPC2`.

Because our state in Lume is stored as list of entire history, we are able to search the list for moments in the past that signalled major narrative changes—the death of a character, the introduction of a character, the beginning or end of a romance, huge swings in relationship between two parties, and so on. By combining this feature with the ability to attach recall phrases to the events that cause these changes, we are able to get a huge range of the kinds of reactions above, but we can do this going back any number of turns. While intuitively, we thought that players' memories might be short-lived and certain actions would not actually resonate after many turns, in practice when we focus recalls on large swings, the result is often surprising and impactful.

8.5 Characters and Relationships

NPC Behavior is largely governed by character traits and enacted through scenes or individual nodes within a scene that specify how a character should behave. Character traits are properties that exist on a character either as a tag (without an incremental value)or as a stat (with an incremental value). Character traits govern NPC *behavior* and *reactions* to various scenarios.

Characters also have a property called their *demeanor* which governs how a character speaks and any particular mannerisms. The distinction between character traits and character demeanors allows for possibilities like the gruff rogue with a heart of gold, the upbeat but cowardly sidekick, or the friendly but calculating advisor.

Relationships among characters are represented by unidirectional values, meaning Jose might like Sally, whereas Sally detests Jose. Additionally, tags might be attached to relationships such as

- `[afriad_of, NPC]`
- `[attracted_to, NPC]`
- `[wants_revenge_on, NPC]`
- `[in_love_with, NPC]`
- `[loyal_to, NPC]`
- `[jealous_of, NPC]`

- `[protective_of, NPC]`

Certain key scenes might be authored around character moments involving traits and relationships. For example an author might create a scene in which a `shy` NPC who feels `protective_of` another stands up to an aggressor and experiences character growth (and perhaps the removal of her `shy` tag) as a result.

In practice, many scenarios key off of relationship bands (i.e., whether a relationship is good, medium, or bad) between NPCs or between an NPC and the player. Action is most interesting when most choices change relationships (though they needn't always change them by large amounts). Huge swings in relationships have the highest chance of being selected for a Recall event, and they might also open up the possibility for knock-on effects like a character becoming afraid of another during a large negative swing or attracted to or admiring another during a positive swing; such knock-ons are usually governed by personality traits.

Additionally, characters might have or acquire specific roles in the narrative. For example, a story might specify that a character who `wants_revenge_on` the player becomes her textttnemesis and the story's main antagonist. Alternatively, an author might specify that a story has an overarching antagonist, and reserve the nemesis for more of a mini-boss figure. This kind of authorship happens at the scene level and is largely left up to authors or individual stories. Built in story roles include:

- nemesis
- antagonist

- guide
- sidekick

Authors can add additional story roles to game-specific logic.

8.6 Balancing Narrative Dynamism and Coherence

One of the primary goals of the Lume system is to strike a balance between *narrative dynamism*, the feeling that what's happening is one of many paths that could have been taken, and *narrative coherence*, the feeling that narrative events are causally following from player actions or logical NPC reactions. Several of Lume's features—and especially the combination of these features—offer fruitful steps toward these goals. The coherence-enforcing systems—bindings, recalls, localized coherences through small trees—balance the dynamism of the combinatorial elements.

8.6.1 Toward Narrative Dynamism

8.6.1.1 Narrative Dynamism Through Scene Selection

Lume narratives are inherently dynamic, in that they are comprised of scenes that will always appear in different orders with different characters within them, given a large enough content pool. This architecture means that authors must take care to impose their own narrative scaffolding.

8.6.1.2 Narrative Dynamism Through DCGs

DCGs ensure that players are seeing varied content on each playthrough. Well-formed DCGs can support huge combinations of possible text output, and the context-specific nature of Prologs DCGs mean that expansions can be tailored to current game circumstances.

8.6.1.3 Narrative Dynamism Through Conditional Choices

Preconditions on choice nodes mean that players may be presented with different choices in different contexts. Yet because of the nature of the content selection mechanisms, we do not have to worry that more choices will necessarily lead to greater authorial burden.

8.6.2 Toward Narrative Coherence

8.6.2.1 Narrative Coherence Through Bindings

A careful combination of bindings and preconditions allows us to construct coherent narratives. In particular, bindings allow us to have narrative throughlines that provoke causality in the player's mind. Bindings allow us to fill the appropriate character, place, or event into a narrative moment based on the current world state.

8.6.2.2 Narrative Coherence Through Content Control Flows

In a well-formed narrative, one scene is the logical outcome of another, and indeed there may be places where authors want to specify that if X happens, Y should happen in response. Lume gives authors control of the flow of narrative on a variety of explicit specificities by allowing authors to:

- Directly control which scene should fire next
- Specify that the next scene should be from certain content pools
- Specify that a scene from a certain pools should fire in the next X turns
- Indirectly steer toward certain scene relationships by authoring post-conditions to feed into other scenes' preconditions
- Indirectly steer flow through general scene priority

8.6.2.3 Narrative Coherence Through Recall Phrases

Recall phrases are useful devices to remind players that their decisions have changed the world, and changed characters' perceptions of them. This feature directly highlights the causal relationships in the narrative, and gives players a sense that their actions have caused the current narrative events. The fact that recall phrases are dynamic furthers the feeling of a deeply-tailored narrative experience. See Figure 8.3.

Figure 8.3: Example narrative emerging from individually-selected scenes. The combination of preconditions, bindings, and recall phrases help ensure coherence.

8.6.2.4 Narrative Coherence Through Knowledge Representation

We have developed a set of rules and DCGs that allow characters to refer to other characters with generic tags if they do not know them. Thus characters could be "the man" or even "the man who [attacked us in the woods]" until he introduces himself as Robert. The knowledge representation capabilities are outside of the scope of this paper, but present interesting and promising future work.

8.6.2.5 Narrative Coherence Through Conditional Text Generation

In addition to DCGs offering a large possibility space for dynamic content, their context-specific nature means that we can ensure the most relevant expansion occurs. Thus our characters might speak in text that is specific to that character's voice, or might offer different reactions altogether based on their mood, the players' relationship, or a status. Locations might contain generic descriptions to be overridden by current story details. Narrations of different events might change based on the mood of the narrator. And so on.

I should also restate that I do not wish to claim that Lume is the first system to offer the features detailed in this section. Other systems have conditional choices, such as Inkle games or *King of Dragon Pass* [154]; *Façade* [105], and particularly *Prom Week* [109] make heavy use of recall phrases; dynamism through binding is a feature in StoryAssembler [62], *Prom Week*, and many HTN-based story generators. Instead, the particular way we have combined these features has potential for creating interesting narratives that demonstrate a high degree of narrative responsiveness.

8.7 Technical Interventions Toward Responsiveness

Though we did not have the language for it, our goal from the outset of Lume has been to create a system that could offer responsive narratives. In terms of the model of responsiveness presented in this dissertation, our interventions have largely occurred in the think and speak step of the model. Through the design of the system, we hope to leverage tight coupling between player actions and narrative feedback, offer signalled counterfactuality, and change affordances to a very high degree at narrative levels (character, plot, story, narrative roles) while keeping the possibility available for changing affordances at higher levels of abstraction (changes to narrative roles, roleplay opportunities, story genre or type of story, story world ontologies, accepted solutions to conflicts, etc). Even though the latter are next expressly encoded in the system, we have tried to architect Lume's approach to narrative and rule-authoring from the bottom-up as much as possible, leaving room for authors of particular projects to define those approaches; in other words, we have not explicitly added top-down support for these particular features, but we have also tried to ensure we leave those possibilities open for authors. Additional details are listed below, and more details about authoring top-down narrative approaches in a bottom-up system are provided in Chapter 9.

8.7.1 Think: New Approach to Story Modeling

We previously saw several approaches to story modeling in procedural narrative systems. The storylet model assumes that moments of narrative will combine together in interesting ways to create varied, salient narrative. Lume offers narrative responsiveness in its model of narrative at multiple levels as detailed in the following sections.

8.7.1.1 Modeling Story Structure

At the narrative level, dynamic node-tree structure of scenes allows us to alternate versions of scenes that respond to players' choices and the game's state. Scenes might have different choices available to the player—that is changes to affordances—based on players' past choices.

As we saw in Chapter 6, offering combinatorial affordances allows us to offer responsiveness at multiple levels of abstraction. A combinatorial narrative structure allows us to have the entire story respond to players at a fairly high level of abstraction and with a high degree of counterfactuality. This is achieved through responsiveness in plot through the narrative combination of scenes and responsiveness in narrative roles.

Lume supports responsiveness at a high level of abstraction in multiple ways:

1. It offers combinatorial affordances via:
    (a) Narrative events
    (b) Different story roles
    (c) Conditional choices
2. It supports counterfactuality to a very high degree, which is then integrated into the world model at the think step and potentially communicated by a client application in the speak step

8.7.1.2 Modeling character

The way Lume models characters supports radical changes to various characters as a result of player actions. NPCs can change their character traits—that is, they experience a character arc—as the result of player actions. Additionally these arcs might interface with the narrative roles listed above. NPCs respond differently in various narrative contexts based on these traits, thus their actions (and the affordances available to the player) can change dramatically depending on past actions.

8.7.2 Speak: New Approach to Feedback

8.7.2.1 Surfacing past choices

As we saw in Chapter 6, feedback alone can be a powerful tool toward responsiveness. Recall phrases allow for a high (and varied) degree of narrative feedback. They also provide a powerful tool toward improving diegetic signalling of affordances and assure the player that affordances are changing in response to her actions.

In addition to improving the perception of narrative causality, recall phrases as implemented in Lume are designed to improve the authorial burden of this surfacing work.

8.7.2.2 Surfacing counterfactuality

Because all past events are stored, we are able to signal state changes to players not only by the diegetic recall phrases mentioned above, but also, if desired, as UI elements. A game might decide to display certain kinds of events as UI elements, or offer an entire event history if desired.

As previously mentioned, many of the individual components of the Lume system have appeared in other systems. I argue that it is no single component that makes the system powerful, but rather how the combination of these elements allows for highly responsive narratives while reducing authorial burden to a manageable level.

8.8 Authoring with Lume

Many of the features listed in this chapter are specifically designed toward the abstract. While Lume supports templated scenario trees, recalls and event lists, basic character relationship functionality, DCGs and built-in pronoun replacement, things like narrative scaffolding, specific control flow between scenes, specific character traits and functionality, etc are left up to individual authors. While Lume is a highly adaptable system that supports a range of potential experiences, much of the burden of getting well-formed narrative to emerge from the system rests on individual authors.

Lume is a procedural system designed toward responsive narrative rather than one designed to leverage procedural narrative generation primarily to decrease authorial burden. While features like pronoun substitution and clever DCG authoring ameliorate burden, in practice, authoring a work in Lume still takes a large library of written content and a lot of design overhead in the form of managing scene authoring and design and designing and leveraging scenes for a balance of variety and reuse.

The next chapter will detail the authorial affordances and design learnings from authoring *Rumina Woods*, a prototype Unity game using the Lume system.

Chapter 9

Case Study of Rumina Woods

In the previous chapter I introduced the features of the Lume system, a procedural narrative system created to foster responsive narratives.

All systems inherently impose some constraints onto authors: some things will naturally be easier to create, and some will be harder. These represent authorial affordances [101], which ultimately influence the nature of works created with a particular system [61].

In this chapter I present a case study of authoring with the Lume system. I present a prototype of our game, *Rumina Woods*, still in development, and talk about the authorial affordances of the Lume system. Additionally, I present some of the design insights I have learned in working with it on Rumina and the prototyping work that has gone into it.

9.1 Overview of Rumina Woods

Rumina Woods is a procedural narrative adventure game in the style of a visual novel. Players take on the role of Clara in a surrealist choice-based story in an Alice/Oz-style fantasy world.

Clara is an adolescent girl who escapes a troubled family life by traveling to Rumina, her imaginary world in the forest near her house. Rumina is a world of talking plants and animals, magic items, and an unspeakable past trauma that corrupts the forest with a mysterious disease. Just as Clara's family life reaches a critical point, she meets another child in the forest, who instantly becomes fascinated with Clara and her imaginary world. Eventually Clara must choose between new friends and connections with loved ones in the real world and her friends in Rumina.

Our goal with *Rumina Woods* has been from the outset to create a visual novel with a radically, delightfully responsive narrative. We want the feel of embedded narrative—in contrast to the particular pleasures of the kinds of emergent narratives that arise from games with simulation-manipulation at their core—but one that clearly signals its own responsiveness and showcases the range of the Lume system.

Rumina Woods is an adolescent fantasy story in the genre of *Alice's Adventures in Wonderland* [29] or *The Wonderful Wizard of Oz* [14]. Independent tabletop roleplaying game *Girl Underground* [113] is a strong genre influence as well as films like *I Kill Giants* [188] and *Pan's Labyrinth* [38]—films in which a girl escapes a tough real-world situation by dealing with her problems in a strange, escapist fantasy setting. The setting provides us with a lot of forgiveness toward generator inconsistencies: most of the generated content takes place in a setting that is clearly signalled to be Clara's imagination. Unexpected, topsy-turvy behavior is expected in this setting, so odd combinations from the generator are more likely to go unnoticed by players, and as developers we can experiment with pushing the bounds of generative space without needing too much corrective rule authoring to ensure narrative coherence.

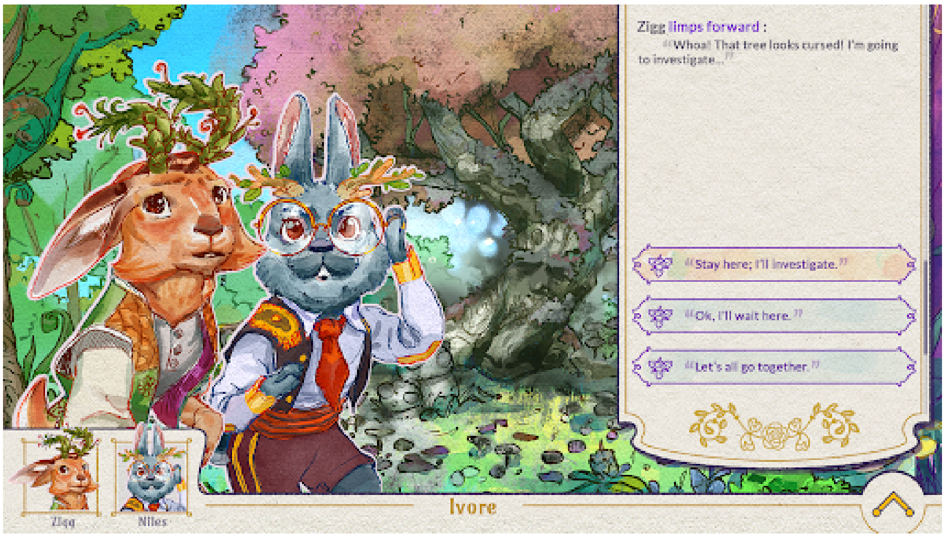

Figure 9.1: Screenshot from *Rumina Woods* prototype

9.2 Development: Integrating Lume & Unity

My collaborators and I are developing *Rumina Woods* in Unity, a popular engine for cross-platform game development. Visuals, UI logic, and game-state data is all written and stored in Unity while the Lume system stores our database of scenes and narrative logic. Unity queries Lume for the appropriate scene and game data through an API.

Unity offers a variety of libraries and modular development tools to augment the base engine's capabilities. My collaborator Carl Muckenhoupt has leveraged these libraries to create development tools and pipelines to speed up the process of authoring Lume content (see Figure 9.2. Because one of the key challenges to procedural narrative in the wild is a high degree of authorial burden [61] [140], the impact of the creation of good authoring tools that speed up procedural content production and are robust enough to stand up to real-world authoring conditions—along with all of the indecision, editing, moving, and ability to facilitate internal team communication required in realistic collaborative authoring scenarios—cannot be overstated. To this end, Carl has built a scene-authoring tool in Unity that allows authors to quickly construct scene node-trees in a visual editor, test them in the current build, and save them to the Lume database.

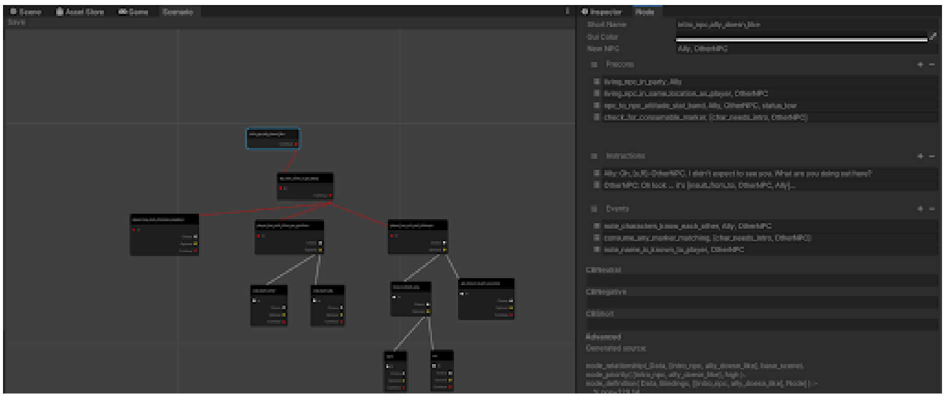

Figure 9.2: Screenshot of Lume's Scene Builder tool in Unity environment.

The scene authoring tool allows an author to quickly construct scenes, rearrange node structures, experiment with different node types, copy and paste nodes, quickly reorganize scene files, and so on. The tool also allows authors to manually set colors of nodes, which we use internally as a team to communicate things about the state of the authoring—for example, red nodes indicate a node is merely stubbed or needs an edit pass, yellow means it could be stronger but is not pressing, and so forth.

As we integrate Lume with Unity more deeply, opportunities for new uses abound. For example, recall from Chapter 8 that a scene's output primarily comes in the form of *instructions*, the text output that occurs when a node is selected, and *events* which are added to an event list. While *Rumina Woods* is a text game in which the primary content units of *instruction* are passages of narration and dialogue text, more abstract versions of instructions are possible. For example, Lume allows for markup in the instructions fields of scenes, which might pair an animation with a particular line of dialogue or a particular bucket of reactions. These findings suggest that future uses might involve Lume also serving as a system for selecting dynamic animations or voice lines, for instance. Currently such uses have received limited but promising initial tests.

9.3 Narrative Structure of Rumina Woods

Lume's model of scenes does not, by default, impose any particular hierarchy or suggestion of when individual scenes should occur. Instead, it treats every scene as a card in a deck that is equally-weighted unless a designer specifies otherwise until a scene is drawn (once drawn, scenes receive a negative weight toward repetition). Consequently, Lume does not force—or even really suggest—any particular narrative structure, and top-down narrative cohesion must be added by authors. For *Rumina Woods*, narrative structure comes in two flavors: first, we have parallel but intercepting narrative lines between the Real World (which we have shorthanded as IRL) and Clara's fantasy world, Rumina. Second, the overall story arc follows a 3-act narrative structure.

9.3.1 Mirror World

The narrative of Rumina Woods takes place between the Real World (IRL) and Clara's fantasy world, Rumina. Elements of the two bleed into each other and Clara will often need to face inner demons in Rumina before facing them IRL. Rumina is where she grows and changes as a person to gain the skills she needs in her real life.

As such, the structure of the game follows lots of parallels and back-and-forth between the two locations. What happens in Rumina affects IRL and vice-versa. For this interplay, we are drawing a lot of inspiration here from works like *Pan's Labyrinth* and *I Kill Giants* in which the trials the character faces each time she enters the fantasy world parallel and influence her understanding of the real world.

9.3.2 Fires in the Desert

Fires in the Desert is a term borrowed from Failbetter Games [54] in which you have high-impact beats of strongly-authored content at key points with generated content (that could arbitrarily be anything) in between. The metaphor comes from the idea that you are in a dark desert and can see fires in the distance, but can't see anything between you and the fire. Thus the "fires" are key narrative moments, and the dark is the generated content in between these moments. Ideally, the fires give enough of a strong backbone that our story feels well-formed, but still gives the freedom for us to leverage generated content in between such that the story as a whole feels both dynamic and tailored.

In *Rumina Woods*, we call these fire scenes *impact beats.* Impact beats typically happen in the Real World, and they seed content for the generator of the more procedural narrative moments in the fantasy world. Major decisions between the real world and the fantasy world pass back and forth and influence each other. For example, Clara's parents are at odds due to her father's drinking and both parents are vying for her affection. Clara's feelings toward her parents manifest as situations and obstacles she has to overcome in Rumina. Similarly, various abilities or growth opportunities in Rumina translate into coping strategies and abilities for Clara to handle her problems in the real world.

In general, each "run" through Rumina begins with a narrative impact beat in real life, then the player experiences several generated scenarios that detail Clara's adventures in Rumina. Upon completion of the current question or objective—which usually features both a gain and a loss or complication—the player is returned to real life and the Act/Scene structure advances.

In developing the narrative structure, I have taken inspiration from Dan Harmon's story circle, a narrative template that Harmon uses for writing popular television series *Community* and *Rick and Morty.* Each run through Rumina roughly corresponds to one loop around the circle.

Scriptwriting advice blog *Studiobinder* describes the circle both in terms of what each step represents, but also in terms of how it feels [172]:

> The Dan Harmon Story Circle is a story structure divided into eight distinct parts following a protagonist's journey. Also called "The Story Embryo" or "Plot Embryo," these 8 steps follow a character's pursuit of a goal outside of their normal world. Their inevitable return finds them changed, whether or not they achieved their goal.
>
> 8 STEPS IN THE DAN HARMON STORY CIRCLE:
>
> 1. You — A character is in a zone of comfort,
> 2. Need — But they want something.
> 3. Go — They enter an unfamiliar situation,
> 4. Search — Adapt to it,
> 5. Find — Get what they wanted,
> 6. Take — Pay a heavy price for it,
> 7. Return — Then return to their familiar situation,
> 8. Change — Having changed.
>
> [...] Why is the Story Circle a circle? Why not a straight line from Step 1 to Step 8? The circle provides an intangible momentum to the story, almost like a rollercoaster.
>
> A protagonist begins at the top must descend figuratively in the story and literally in the circle. At the bottom of the circle, they are at their literal and figurative "low point," and their rise to success in the end is likewise represented visually on the Story Circle. [172]

The circular nature of the template, its generalizability, and its alignment with momentum makes it a good fit for each run through Rumina. Thus the overall narrative structure is depicted in Figure 9.3.

9.4 Example Scene Authoring

To demonstrate the process of authoring in Lume, let us walk through how an author might create a scene. In this example, an ally in the player's party has grown

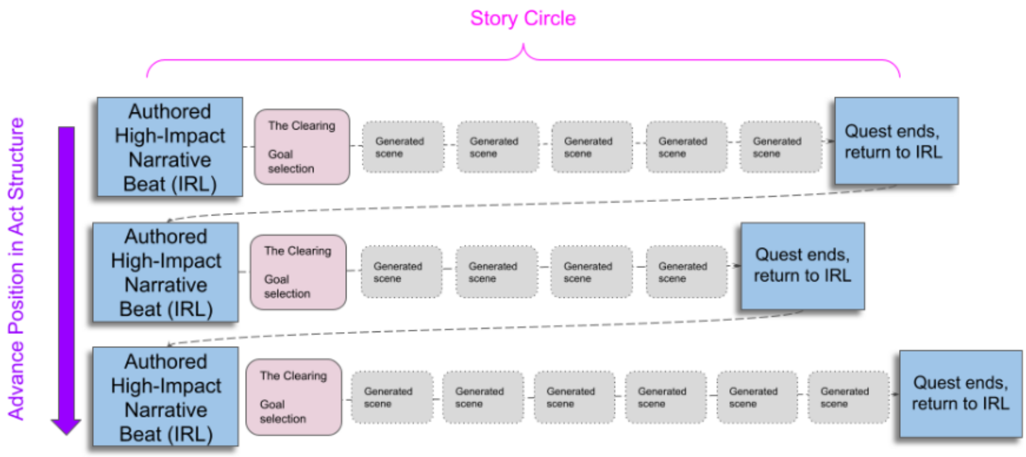

Figure 9.3: Diagram of the narrative structure of *Rumina Woods*.

upset with her actions and confronts her about leaving the group. As we author the scene, several authorial goals are likely in play:

1. We want this to be a logical, narratively consistent moment, so the scene must fire at the appropriate time.

2. We want the player to know that this is happening as a result of her choices.

3. Different characters might not leave the party in the same way. We want the scene to feel different if different characters bind to it.

4. We want this to potentially be a turning point moment that might be referenced later.

5. We want this to be a moment of real consequence. Losing a party member is a big deal in a game about friendship and relationships. The narrative situation should feel impactful.

9.4.1 Ensuring Proper Selection

The conditions of a base node determine when a scene will be selected, so they are the first things we author as we construct our scene. As an author we want to answer the following questions:

1. Where can this scene occur?

2. Who is in the scene?

3. What is happening narratively to make this occur?

For each of these questions, we want to author with the broadest possible answer in mind that will achieve coherence. Answers to these questions for our example are as follows:

1. Where can this scene occur? Anywhere (no preconditions needed)

2. Who is in the scene? An unhappy ally

3. What is happening narratively to make this occur? something has made an ally very unhappy with Clara

So to achieve selection under these conditions, we give our base node the following conditions, which we have selected from the database of existing conditions:

- `[living_npc_in_party, Ally]`

- `[npc_attitude_stat_band, Ally, status_low]`

- [find_descr_of_attitude_change_to_player, neg, Ally, Recall]

By using the same NPC variable (Ally) for multiple preconditions, we have ensured that the system will bind an NPC to this scene under the variable Ally as long as one fulfills these criteria. So if we have three party members, and all of them are happy with us, this scene will not be selected. If, however, we have a party member who is unhappy because of something we did, the scene will be eligible for selection. The last condition searches the event list for a negative change in attitude from Ally to Player (it will prefer the largest change), finds the recall phrase associated with that event, and binds it to a variable called `Recall`.

9.4.2 Signalling Consequence and Counterfactuality

Our second authorial goal is to ensure that the player knows this scene is playing out as a result of her actions. By finding an applicable recall phrase, we know we can narratively signal to the player why this is happening. In the process, we also imply that things might be different if the player had made other choices. Base nodes set the stage for the scene. To author the content of the base node, and set up the scene's premise, we type what we want players to experience in the instructions fields. We provide the following instructions as a first-pass:

```
Ally comes to Player looking unhappy.
Ally:  I've been thinking a lot about how Recall.
Ally:  I think it's best if I leave the group.
```

See Figure 9.4 screenshot of the single node and the Unity authoring environment.

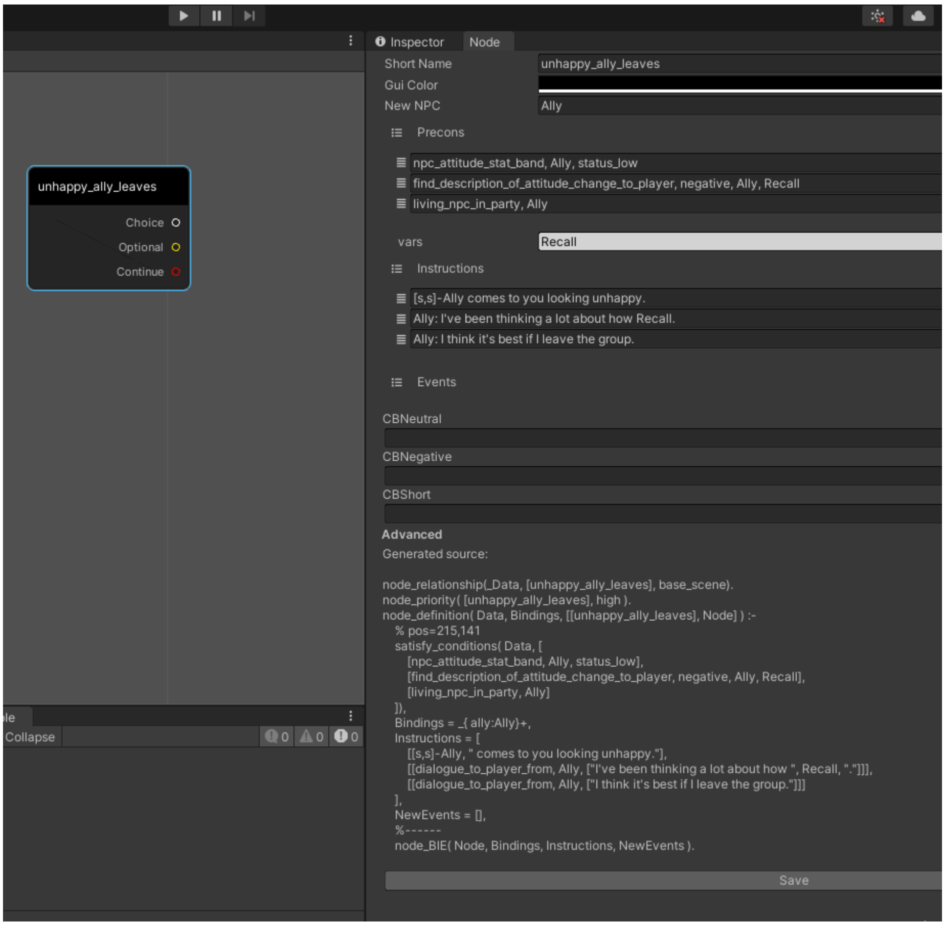

Figure 9.4: Screenshot of the Lume Scene Editor depicting a single node and the Node editor.

At this point, most narrative designers will be itching to give their player a choice—what kinds of actions should be available to the player? For simplicity, let's give the player the option to ask the ally to stay or to accept the ally's departure without a fight.

We create two choice nodes to provide these options:

- One asking to stay

- One saying goodbye

Neither of these options is a conditional choice, so they do not need preconditions on the choices. We will also fill both nodes with an initial pass of content.

For the goodbye option, we write some dialogue for the character to say once the player selects this option:

```
Ally: You're not even going to try
to stop me? You're not the person I thought you were...
Ally casts a sad look at Player over
Ally's shoulder as Ally leaves.
```

We also give this option events that should fire if this is chosen by the player:

- Ally leaves the party
- Ally likes the player even less than before.

For please stay, maybe we decide this choice should win the player back a little bit of goodwill from Ally. We author some text that provides feedback that the player's choice is acknowledged.

```
Ally looks at Player sadly.
Ally: You want me to stay?
```

For the events, we first have a small positive increase in the Ally's attitude towards the player. Additionally, this might be a small moment of positivity a Ally will refer back to in the future. We also author a positive recall phrase for this node.

Recall: `Player wanted Ally to stay`

Now at this point, we want to think about how different characters might handle this situation. The scenario graph at this point is depicted in Figure 9.5.

Figure 9.5: Screenshot of the Lume Scene Editor. The example scene is in progress (zoomed in).

9.4.3 How Different Characters React

Different characters might react to being begged to stay in different ways. At this point, different outcomes happen as a result of this choice node, only one of which will be selected at runtime. We want to make sure that at least one of the outcomes of this choice will always be eligible, so we make sure one of the outcomes has no preconditions. In some cases, we may want different outcomes:

- (Default case): the character leaves anyway.
- A compassionate character might immediately accept and stay.
- An Ally with the loyal trait might stay if the player has done something nice for them.

For the first, we display some text but the ally leaves anyway. The data for this node is as follows.

> Preconditions `(None)`
> Instructions: `Ally: I'm sorry, Player. Too little too late.`
> `Ally casts a sad look at Player over`
> `Ally's shoulder as Ally leaves.`
> Events: `[left_party, Ally]`

For the second choice, our node is as such:

> Preconditions: `[npc_has_trait, Ally, compassionate]`
> Instructions: `Ally sighs.`
> `Ally: It's been a difficult journey.`
> `I'm sure you're doing your best. .`
> `Ally: I'll stay but please`
> `try to be more considerate.` Events: `[left_party, Ally]`

The player has already regained some faction with Ally, so after some output text, no further events are necessary.

In the third case, Ally values loyalty so they want to make sure the player will be loyal to them. We decide we want to raise the stakes with this loyal character by having the player either reject them or make them a promise (promises are ripe for dramatic tension). Our node is as such:

Preconditions: `[npc_has_trait, Ally, loyal]`

`[find_description_of_attitude_change_to_player, positive, Ally, PosRecall]`

Instructions: `Ally sighs.`

`Ally: We've been through a lot together. And I really do appreciate that PosRecall.`

`Ally studies Player a moment, considering what to do.`

`Ally: Loyalty is really important to me. I'll stay if you promise to have my back.`

The scene graph at this point is depicted in Figure 9.6.

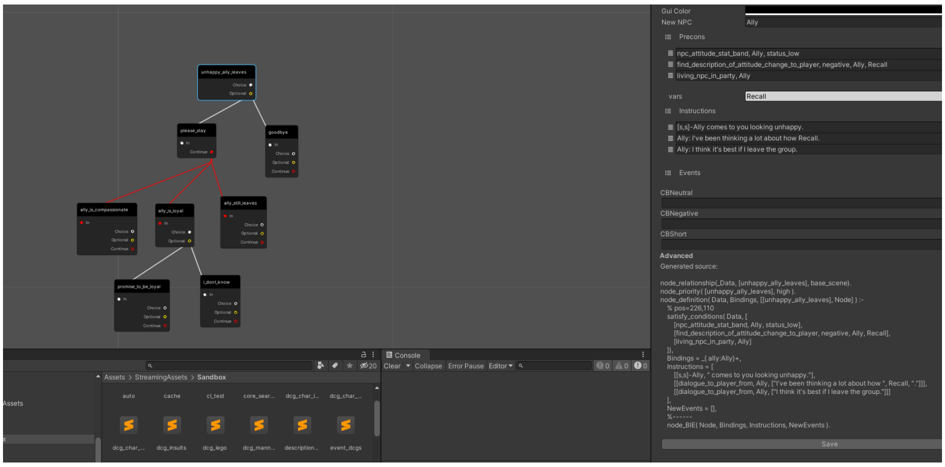

Figure 9.6: Screenshot from the Lume Scene Editor showing the final structure of our example scene (zoomed out).

9.4.4 Foreshadowing and Seeding Future Drama

At this point, if the player has asked a loyal character to stay, the loyal character then asks the player for a promise. We need to create choice nodes to reflect the player's choice options in this moment. The player will have two options:

- "I can't promise anything"

- "You have my word"

Neither of these are conditional choices, so we do not need to author preconditions for either of them. In the case the player does not promise, the loyal Ally loses further attitude toward the player.

> Preconditions: `(None)`
>
> Instructions: `Ally: If you can't promise to be loyal to your friends, you're not the kind of person I want to be friends with.I'm leaving.`
>
> `And with that, Ally disappears into the forest.`
>
> Events: `[major_decrease_in_attitude_to_player, Ally]`
>
> `[left_party, Ally]`

This is a moment that might be referred back to later. Perhaps we will encounter our former ally later, and they will be angry with us for leaving. Or perhaps someone else in the party will be glad Ally is gone. We want to make sure that we author recall phrases that the system will grab.

In practice we found that positive phrases rarely differed from neutral ones with a slightly positive tone, so we tend to use those for positive recalls. When authoring recalls I tend to phrase the positive ones as if the player is being praised for the action, or if it is being reported by a friendly third party, and the negative ones as what a character might say in an argument. Short recalls are how this incident might be referred to as a casual shorthand. Recalls usually require lots of tweaking to ensure they fit properly in a variety of situations.

> CBNeutral: `Player couldn't promise to be loyal to Ally`
>
> CBNegative: `Player doesn't value loyalty`
>
> CBShort: `Ally left`

For the other option, "you have my word", the node data looks like this:

Preconditions: `(None)`

Instructions: `Ally: Okay... I'm trusting you...`

Events: `[create_consumable_marker, [player_made_promise_to, be_loyal], Ally]`

CBNeutral: `[Player promised to have Ally's back]`

The system has now placed a flag that the player made a promise to be loyal to Ally. Down the road, we might use the existence of this promise as the precondition on a scene in which the player's loyalty will be tested.

Consumable markers allow the system to note various foreshadowing events or notable occurrences have happened so that the system can pay them off in the future.

Now let's see some sample output from this scene. Let's suppose the player has upset Zigg, a character with the `loyal` trait. Zigg is upset because Clara did not take his side in an argument against their other ally Niles. But he had a positive experience when the player let his friend Cay into the group. Example output of this exchange might look as follows:

Zigg comes to you looking unhappy.

Zigg: I've been thinking a lot about how you took Niles's side in that argument.

Zigg: I think it's best if I leave the group.

- **Please stay.**
- Very well. Goodbye.

Zigg looks at you sadly.

Zigg: You want me to stay?

Zigg sighs.

Zigg: We've been through a lot together. And I really do appreciate that you let Cay into the group because I asked you.

Zigg studies you for a moment, considering what to do.

Zigg: Loyalty is really important to me. I'll stay if you promise you'll have my back.

- **You have my word.**
- I can't promise anything.

Zigg: Okay... I'm trusting you...

Consider instead that Niles comes to you upset and wanting to leave. Niles has the trait proud and lost a lot of attitude toward Clara when she embarrassed him in front of Ponch.

Niles comes to you looking unhappy.

Niles: I've been thinking a lot about how you made that joke in front of Ponch.

Niles: I think it's best if I leave the group.

- **Please stay.**
- Very well. Goodbye.

Niles looks at you sadly.

Niles: You want me to stay?

Niles is not compassionate or loyal; the default response is selected.

Niles: I'm sorry, Clara. Too little too late.

Niles casts a sad look at you over his shoulder as he leaves.

9.4.5 Adding Variety and Voice

In the example outputs, Zigg and Niles are very different characters: Zigg has traits `loyal`, `reckless`, and `virtuous`. He has an `impatient` demeanor. Niles has traits `proud`, `vengeful`, and `insecure`. His demeanor is `haughty`. In a hand-authored scene, we could easily imbue the characters' lines with their personalities, but in our working example, much of the dialogue and exposition are shared. The characters might behave differently by different nodes of the tree activating—Zigg certainly gave us another chance, after all—but the characters do not really feel different at this point.

This is where the combinatorial magic of DCGs come in. We can write conditional text that might functionally mean the same thing, but expresses lines in different voices to reflect different characters. The trick is to balance DCG authorship: sure, every scene would be more expressive if we wrote a version that replaced every line with all of the alternate lines we might expect from different characters, but the authorial burden of that would increase exponentially with every character trait we added. The trick to to balance authoring reusable DCGs with high bang-for-buck alternates for variety.

While we use character traits to change how characters behave, and we offer many variations on those, we use a character's demeanor to determine how that character speaks, in other words, its voice. Zigg is impatient and Niles is haughty. Let's review the lines of our base node with this in mind:

```
Ally comes to Player looking unhappy.
Ally:  I've been thinking a lot about how Recall.
Ally:  I think it's best if I leave the group.
```

The first line probably already has a DCG authored for it that we can use. Descriptions of how characters look that convey their state are good candidates for DCG authorship because they can be used in lots of places. Let's presume we already have a generated description of how a character looks. We might rewrite the first line as follows:

```
[descr_approach_target_conveying_rel, Ally, Player]
```

The second line is very generic. "I've been thinking a lot" is fine, but does not convey voice or current emotional state. If Niles, a haughty, vengeful character who is upset with us delivered that news in such a measured manner, it would be very strange. So let us author this as a DCG instead.

For DCGs we always want to have a fallback case in case the system does not meet the conditions for selecting a more interesting line. So we will leave the original line as the fallback. Then, let's override the line with something that switches based on demeanor. If the ally is has a `haughty` demeanor, they will say

```
[Ally:  You know, I cannot---for the life of me---figure out why
Recall.
```

whereas an `impatient` character might deliver the line more directly even if the situation is awkward for them.

```
Ally:  Look.  I don't know how to say this.  Recall...
```

The last line is effectively the character proposing an exit. This also feels like a situation that arises in more than one place, and we can author demeanor-specific variations that could be applicable in leaving the party, a conversation, or a location. For the `haughty` character:

```
Ally:  I'll be going.  I doubt I'll miss this.
```

And for the `impatient` character:

```
Ally:  I should go.
```

We might throw in an optional physical mannerism for flavor and variety too. Now the scene feels very different between the two characters.

When Zigg (loyal, reckless, virtuous, impatient) binds to it:

> *Zigg shuffles up to you looking deflated.*
>
> *Zigg: Look. I don't know how to say this. You took Niles's side in that argument...*
>
> *He fidgets uncomfortably.*
>
> *Zigg: I should go.*
>
> - Please stay.
> - Very well. Goodbye.

When Niles (proud, vengeful, insecure, haughty) binds to it:

> *Niles waltzes in, anger flashing in his eyes.*
>
> *Niles: You know, I cannot—for the life of me—figure out why you made that joke in front of Ponch.*
>
> *He looks down his nose.*
>
> *Niles: I'll be going. I doubt I'll miss this.*
>
> - Please stay.
> - Very well. Goodbye.

Even though each of these lines functionally communicates the same thing, the addition of character voice through DCGs adds narrative context and variety to the scene. Additionally, it changes the affordances; Zigg is more sympathetic than Niles in this example, so the player is much more likely to ask him to stay than Niles. The combinatorial nature of scene construction means that affordances change in this way as a result of player choices. Additionally, constant feedback about characters' state means that players can strategize about their relationships. They can experience content with different allies or in different positions depending on their choices leading up to a particular scene.

9.5 Authorial Affordances of the Lume System

The Lume System presumes that narrative will largely emerge through combinations of storylets. As such, many of its authorial affordances arise from the assumption that in general, authors will be crafting small bits of content and assembling those into a larger narrative, with rules (via preconditions/postconditions) to provide most of the coherence and causality. In practice, additional corollary assumptions emerge from this architecture:

1. **Lume presumes reasonable pruning at the bucketing and base-node level.** In practice, Rumina's pool of scenes is on the order of hundreds, and the prefetch mechanism buys plenty of time at runtime. However, Lume traverses the entire depth of a node-tree to ensure that a scene is valid. There are many benefits to this approach from an authoring standpoint and to ensure that nodes selected are in fact legal. However, it does influence the way scenes can be authored in practice. As a first-pass, Lume presumes that the author will follow reasonable pruning practices at the base node level so that ineligible scenes are quickly removed from consideration.

2. **Lume presumes the depth of scenario trees is fairly shallow.** For the same reason, Lume presumes the depth of scene trees is fairly shallow. In practice, if a scene is becoming too deep or complex, it is often an indication that perhaps two mutually-exclusive scenarios might be a better option.

3. **Lume does not allow for complex scenario graphs** (e.g. scenario graphs with

cycles). Although the system technically supports branch-and-rejoin structures, it functionally recreates the re-used node to create branches of a tree. Therefore, while such a pattern is easier for authoring purposes (since authors can create/edit the node from a single position), the actual system behavior makes graphs with complex rejoins potentially unwieldy.

For all of these reasons, when a scene tree becomes unwieldy, it is usually easier to author two mutually-exclusive scenes than it is to try to coax overly-complex logic within a single scene. The implications of these constraints ripple out into other authoring best-practices within the system. Because of the tendency toward small bits of content combined together to create emergent narrative, the system does make some things very easy to author and some things more difficult.

9.5.1 What Lume Makes Easy

9.5.1.1 Creating Scenes

Especially since the addition of a visual authoring tool, creating and editing scenes in Lume is very simple. Experimenting with different node types, constraints, and so forth is very quick.

9.5.1.2 Authoring New Rules

Due to the nature of Prolog's logic programming paradigm, creating new narrative rules, categorizations, behaviors, etc. is very quick. Often authors will construct new rules on the fly while authoring a scene. For example, an author might realize that she needs a rule for deciding which characters would save others through various combinations of personality traits, relationship statuses, loyalty, character needs, and so on. It is very easy to create a new predicate for `npc_would_save_other` that can check these various conditions.

One potential downside to this ease of adding functionality, however, is ensuring that such on-the-fly additions are well-organized and documented among larger team members. This kind of workflow is fine for a single author or small team, so we have been fortunate to avoid logic conflicts, but more robust project management processes would be necessary with even a few authors.

9.5.1.3 Context-Dependent Procedural Text

Prolog's built-in DCG functionality makes it very easy to quickly author procedural text and to add conditions to make it interesting. Additionally, the Lume system currently offers three levels of prioritization for text expansions. In practice this means that DCGs are usually authored to ensure they cover:

1. A fallback base case

2. Wide-coverage average cases

3. Interesting specific cases

So for example, a line of dialogue might be written for a very plain fallback with little voice, some alternates of how that line might be spoken by characters with different demeanors, and a version for if something narratively relevant has just happened.

Carl Muckenhoupt and Ceri Stagg created functionality for the Unity authoring tool that also recognizes some shortcuts in dialogue and text authoring to automatically convert lines written into their markup equivalent. For example and author might write:

NPC1: How could Player do this to NPC1!?

And the authoring tool converts this to the Prolog code:

```
[dialogue_to_from, NPC, Recipient, [ "How could ", player_id, "do
this to ", [object,standard]-NPC, "?!" ] ]
```

The ultimate output displayed would expand to (supposing `NPC1` bound to Zigg in this situation):

Zigg (to Clara): How could you do this to me!?
or
Zigg (to Niles): How could Clara do this to me!?

The authoring tool uses a series of rules to try to predict that the author's use of `NPC1` in this line should be parsed as the object of the sentence. The authoring tool is accurate most of the time, but the author can also override with markup of her choosing. This can be especially useful if she wants a character to say their full name instead of relying on the pronoun system to guess.

NPC: This is my friend `[f,f]-NPCFriend`*!*

This example will force the full name of `NPCFriend` no matter where it appears in a conversation (where it would otherwise potentially substitute with he or she if `NPCFriend` is the subject of conversation. Thus it expands to:

Zigg: This is my friend Niles!

in the event that Zigg and Niles bind to the respective variables.

9.5.1.4 Surfacing the decisions the player has made

In addition to being able to quickly add procedural text of all kinds, Lume specifically makes it easy to add Recall text. To do this, an author just specifies the type of recall they want: positive, negative, or special event (e.g. death), which by default finds a noteworthy narrative event (i.e., relationship increased, relationship decreased, loved one died, etc). The author can also optionally specify the spin they want if it is different from the recall type. Spins might be positive, negative, or short. So by default the author can just call

```
[find_positive_recall, Recall]
```

and then author dialogue like

NPC: Player wouldn't do that. [Recall]. Player's my friend.

Which might expand to (for instance):

Zigg: Clara wouldn't do that. She saved me from the flood. She's my friend.

9.5.1.5 Bindings as Vehicles of Continuity and Coherence

Bindings are the real magic of Lume. They are the stitching agents that help the system remember who saved whom in the flood. Thus searching the event list for the character that should appear in this scene—the player's most trusted ally, or perhaps a gullible loyalist—easily yields dramatically-interesting situations. By allowing the author to specify the most dramatically-interesting person to appear in the scene at an abstract level, bindings really do a lot of the heavy lifting on continuity and coherence.

9.5.2 What Lume Makes Difficult

9.5.2.1 Debugging in general

By virtue of the fact that Lume is selecting the best content unit (whether scene or DCG) from a pool, it makes it easy for things to fail silently. In situations where a scene is not firing, it is easy to presume the selection mechanism has just chosen something else. Sometimes, we have not noticed a scene had been broken until one of us casually remarked we hadn't seen the mutiny scene in a while. In cases in which an author is certain a scene is broken, it is still difficult to tell which particular condition is causing the failure.

Though we have debugging tools that allow us to quickly get characters and locations into particular states to test scenarios, we do not currently have a good way to highlight when a scene *should* be eligible but the scene selector is not choosing it for some reason. Once we realize a scene is not firing, we also have debugging tools that let us manually force scene selection. Still, it is hard to look for something you do not see breaking.

9.5.2.2 Ensuring coverage

Visualizing the possibility space of a generator is a difficult problem that is known to the procedural narrative research community [63]. While branching narrative systems have built up a set of conventions over the last few decades, procedural narrative still does not have any commonly-accepted visualization techniques.

We do not current have good tools for the Lume system to highlight where scene coverage is needed; in other words, we do not have a good way to visualize scenes that should exist but do not.

9.5.2.3 Voice-Over

Voice-Over (VO) dialogue is commonly used in narrative games, especially ones of greater fidelity. It is also an expensive part of the development process, and the cost scales based on how many lines voice actors must record. Much of the power of our system depends upon lines assembled from generated text, voice-over is one of the largest hurdles to still to overcome; for *Rumina Woods* it is too difficult a problem to be feasible.

A naive solution might be the brute force approach: asking actors to record every possible line in every possible configuration. However, the power of things like recall phrases comes from our ability to stitch together sentences from atomic content units of phrases that are smaller than sentences but often larger than single words. Consider the fact that each character and point of view adds an order of complexity onto the space of lines we need, and the brute force approach quickly becomes intractable for a small team. The problem is lessened if we restrict generated text atoms to the sentence-level, however this approach would cost us many of the generation features we currently leverage toward cohesion and expression of character.

Other considerations might be the increasinging quality of generated text-to-speech options such as machine learning approaches. While these are rapidly improving in quality, and could suffice for non-human characters where the fidelity bar is lower, they are not currently at a level of quality where we might reliably put them in the game without fear of unintended side effects.

9.5.2.4 Localization

Because aspects of the text generation systems, such as the pronoun replacement logic, depend on English-specific syntax, localization is a much more complex problem for our system than for most games, which can simply hand a text file off to translators. For each language we decide to localize, we will need to consider language-specific text generation rules; in effect we will need to tweak the text-generation system for each language we support.

Like the VO problem, the impacts of the localization problem can be mitigated somewhat by authoring text content units to be larger than a sentence, but again, the larger the content units, the more dynamism we lose.

9.5.2.5 House of Cards Narrative Structure

Lume was structured for each scene to function like a card being pulled from a deck. By design, it does not offer inherent narrative scaffolding to enforce top-down narrative structure. Consequently, while I would not categorize this as a difficulty per se, it should be noted that significant design overhead must go into structuring narrative rules and ensuring coherence. The approach I detailed Section 8.3 is one approach, but others should work as well. This is both a strength and a difficulty of authoring with

Lume.

Designing narrative for Lume is effectively like designing a card game: top-down rules can simplify the narrative system so the author can wrap her head around it, but a more bottom-up approach also leaves more room for emergence and the opportunity for the creator herself to be surprised by the dynamics that arise in the narrative. Lume does very well with a bottom-up approach to narrative design, and in many ways really shines the more bottom-up the author goes. The more piecemeal the narrative is designed to be, the more room for surprising responsive potential, but ensuring logical transitions, follow ups, and payoffs requires thought. Also like a card game, changes to a single card can have rippling effects on other cards that were designed to work synergistically or antagonistically with this one.

Lume does not offer inherent bridging—either between scenarios or between scenes and mechanics. All of that still needs to be designed, and its design will naturally influence the kinds of scenarios etc that need to be authored. For designers that favor a bottom-up approach, this will feel like a great degree of flexibility. For those that prefer top-down narrative design, this might feel very finicky and cumbersome.

In *Rumina Woods* specifically, there is somewhat of a tension between top-down design and bottom-up design. While major impact beats and overall narrative drive feel very top-down, the generated content in the forest is extremely bottom-up. Marrying the two together in a way that suggests causal interplay is challenging, but extremely rewarding when it works well.

9.6 Design Insights

Over the course of developing *Rumina Woods*, we have discovered many design insights. Many of these are specific to *Rumina Woods*, but many can be generalized to other procedural narratives as well.

9.6.1 Stationary vs Location-Based: Trade-offs in Scene Bucketing

I mentioned above that one of our earliest conceptions of what the Lume system could offer was a game that was something very close to *Reigns* but with a richer narrative throughline. One advantage *Reigns* has that we did not consider is that *Reigns* offers a narrative in which you never have to be embedded in a place. You are the monarch; everyone comes to you with their troubles.

As we introduced a sense of place very early, the question of how to bucket content arose with it: on the one hand, bucketing content by location makes each location feel distinct. For a game about exploration, this feels very natural. On the other, more disparate buckets offer less opportunity for reuse of content; they decrease the amount of content in each bucket, therefore players are more likely to see scenes repeated more often.

Ultimately we landed on a variety of buckets for a mixed approach. Some scenes are very general, have few preconditions, and can more or less occur any time. These are lower-priority scenes but they do offer a sense of familiarity, repetition, or ritual when experienced. Some scenes are location-specific. These offer some distinct flavor to locations. Some are relationship-specific, so these can occur anywhere, but will usually be some of the higher-impact moments in character relationships. And then we have our impact beats. Though there are fewer impact beats, these serve as the backbone of the story, provide motivation and context for the generated action, and offer a sense of canon for players to latch onto.

9.6.2 Size of content unit

The size of our content unit was a decision we kept returning to over numerous prototypes. On the one hand, larger content units skirted or minimized some of the problems above. It allowed us to ensure better-formed narratives and simplified the stitching process between scenes. It also undermined a lot of the strengths of the system.

Smaller content units that were more numerous actually allowed us to move faster in early stages when it was critical to get prototypes together as quickly as possible. It also, better-leveraged Lume's abilities and gave us some truly hilarious and delightful emergent behaviors that we could not have predicted. It also increases the burdens of the house of cards problem. The smaller the content units, the more authorial overhead is needed to structure good selection rules, preconditions and coverage checks.

I noticed a shift for us: we started with a very combinatorial bottom-up approach (with smaller content units). As we moved more toward forms of gameplay that centered the narrative experience entirely and minimized the resource management aspects of previous prototypes, we found ourselves reaching for larger and larger content units. Then as *Rumina Woods* moved back into highlighting simulation elements of the system, we again started moving back to smaller scenes. For us, the more "game-like" the game felt, the it made sense to create smaller content units that more directly interface with the state of resources and other mechanical considerations.

9.6.3 The Uncanny Valley of Representative Fidelity

In previous prototypes with the system, we experimented with higher fidelity graphics, 3D models with procedural animations and so forth. We quickly reached an uncanny valley of expected fidelity: certain constraints of the system (e.g., the lack of VO) forced some elements of the game to represent symbolic actions; we could not get all the way to fully acted-out scenes, and the in-between hit a very uncomfortable uncanny valley. The gap was much wider than we thought. Ultimately we ended up scaling back our visuals and settling on a visual style that is familiar to other narrative simulation games (e.g., *Long Live the Queen*).

9.6.4 Work With The Generator Rather Than Against It

For all of my writing about using rules to force coherence or wrangling the system into a more coherent narrative, I was to stress that most of the best procedural work I have seen works with the generator rather than fighting it to produce something that would have been easier to create with branches.

There are many ways to work with the generator: Dietrich Squinkifer has done excellent work using generator text for non sequitur humor in *Coffee: A Misunderstanding* [167]. Aaron A. Reed's "permutational novel" *Subcutanean* [141] uses the fact that his novel is generated to mirror the books themes of uncertainty and fear of the abyss.

For *Rumina Woods*, the fantastical, surrealist setting smooths over many seeming incongruities and sometimes leads to delightful accidental emergent behaviors. For example, I had forgotten to note that plant characters cannot move, so when a friendly tree walked up and said, "Hello!" while I was testing a scene one day, I was confused and then delighted. It was an excellent example of using the generator to augment the world rather than to push against it.

9.6.5 Narrative and the Oatmeal Problem

The oatmeal problem in procedural narrative [32], mentioned in Chapter 6, describes a situation in which a generator produces many variations of what are effectively the same thing. For alternate lines of dialogue, sometimes meaningful difference can be difficult to pin down. Consider variations on a line of dialogue.

> *Zigg: Clara wouldn't do that! She saved me from the flood!*
> or
> *Niles: Clara would never stoop to such depths! And besides, what should we make of the fact that she saved me from the flood then?*

Functionally, these two lines carry the same information, but they convey different things about the characters in question. They vary slightly in form. They are really two lines that fulfill the same function, and their meaning is effectively the same. We might also increase the level of abstraction from this particular line to `[expression_of_incredulity]` and find other ways to make such an expression.

Are changes in voice really different then? Are changes in tone? Yes and no. As direct-substitutions for one another, no they will probably not avoid players' pattern-matching faculties. On the other hand, they are the kinds of meaningful differences we would expect if the line were spoken by one character vs another.

In practice, we found that the oatmeal problem can be mitigated through a variety of strategies:

1. Vary form as well as presentation. Change sentence structure, length of the expression, etc.

2. Better still, vary content. Ask what point the line is serving in the scene and then author some alternate ways to serve that point. These could be functionally different lines, but they could be a character performing an action as well.

These insights also held at higher levels of narrative. The same scene is instantly recognizable if the dialogue is functionally the same each time. But what is that dialogue trying to communicate? Are there alternate ways to communicate that same thing?

Though players will eventually pattern-match, these strategies have helped to increase the number of repetitions it takes players before they do. And authoring variation must be balanced by those alternate expressions being reusable in multiple places.

9.6.6 Put Oatmeal to Work

The oatmeal problem is not always a problem. Generation in which the details change but the thing is fundamentally the same is making an argument about sameness. Dietrich Squinkifer's *Interruption Junction* [168] does a great job of using generated text to express the unending sameness to the gossipy conversations the players' friends make. The content is not the point; the endless sameness is.

Similarly, Twitter bots that tweet episode plots of a cancelled show are making an argument about the templated sameness of the show's storylines. Sometimes embracing this argument is actually what an experience needs; especially if the author is trying to evoke boredom, ritual, habit, apathy, or pointlessness.

9.6.7 Apophenia is Your Friend

Apophenia is the human tendency to see order and meaning in randomness. In an article about designing procedurally generated personalities, procedural designer Tanya X. Short wrote about the phenomenon:

> Mammals, and humans in particular, enjoy recognising patterns. Pattern recognition is so fun, that we'll look for them on our own where none seem to exist ("apophenia"). Colloquially, I've taken to calling it the Constellation Effect(3), which is a bit of a Promised Land of game design – if you can get players immersed and engaged enough to start seeing meaning in chaos, you've got them hooked. [163]

Sometimes writers want to answer all questions and narrative designers only want paths through the narrative that make perfect, coherent sense. However, I am continually surprised by how much apophenia players will actually inject into causal relationships between narrative events. Even in playtests when we had glaring bugs, players performed mental gymnastics and read all sorts of causality into things. Writing things a little bit vaguely or evocatively to leave room for the player can do wonders toward not only their perception of the story's coherence, but also toward their sense of ownership and satisfaction at *solving* it.

9.7 Looking Ahead

Theoretically Lume has the potential for other narrative functionality of which we are only scratching the surface. We have authored the output of each scenario to be text so that we can prove out the system's capabilityes toward narrative at the levels of plot, dialogue, location, and character. But theoretically, Lume might also be used to search for animations to match particular character reactions, lighting or envinronment assets to match a particular scene's mood, and so on.

One key area of research going forward will be how to build upon the libraries of character behaviors, social rules, and narrative structures we have already constructed to extend and improve Lume's ability to also integrate other aspects of storytelling.

Chapter 10

Conclusions

To return to the quote introduced at the beginning of this dissertation in which Brenda Laurel asks why creating videogames was an obvious thing to do with the unparalleled power of computers so early in their lifecycle, it is, I believe, the allure of interactivity. It is no coincidence that so many researchers and practitioners have independently identified conversation as the core metaphor between player and game, and that this engagement is considered a pleasure core to games as a medium; we crave interaction, and in order to achieve something that feels like a complete conversation, we must have response.

I hope that the introduction of a model of responsiveness might guide our thinking toward how we architect systems toward interactivity and why it matters. This research has engaged responsiveness through a combined approach to theoretical, design, and technical explorations of its properties, each of which have iteratively influenced the others. This research offers contributions along all three fronts.

10.1 Contributions

Presented here are the contributions to theory, design, and technical understandings of system responsiveness. As theoretical contributions, I have offered the following:

- A model of interactivity that positions agency as a component property of interactivity and disambiguates the two terms (see sections 3.1 and 3.2).
- An analysis of the current literature on agency. The synthesized understanding of agency as positioned in this dissertation accounts for seeming contradictions in how the term has been used across discourse. In particular, the understanding of agency as existing at multiple levels of abstraction simultaneously through nested affordances reconciles some of the confusion between agency and its part in roleplay that has plagued game studies discourse (see 3.3).
- A definition of a system's *responsiveness* as the degree to which a system changes its affordances and feedback as a result of player actions and the positioning of responsiveness as the system's counterpart to agency within an interactivity loop (see chapter 4).

For design contributions, I have offered the following:

- An aesthetic analysis of responsive narrative experiences. I explored how different affordance and feedback changes have been utilized to create different aesthetic experiences.

- A case-study of the design patterns of a game designed toward responsiveness and our design learnings from that experience.

And for technical contributions, I have offered:

- A deep dive into the technical design of the Lume system for procedural narrative.
- An exploration of the authorial affordances and trade-offs of authoring with the Lume system in *Rumina Woods*.

As a final point, I would like to use insights about technical interventions in the LTS loop we have discussed previously to look ahead toward how these might inform research toward responsive narrative systems in the future.

10.2 Open Problems in Responsive Narrative Systems

As we have explored technical interventions at each step of the loop, our situating of research efforts as improvements to listen, think, or speak allows us to frame open problems more precisely, particularly in how they might be directed toward narrative responsiveness. Some immediate open problems are listed below.

- **Listen: improvements to input modeling:** while efforts in NLP research and computer vision have improved our capacity to recognize inputs, how game developers can utilize those improvements remains unclear. We have increased the capacity for embodied recognition of verbs—systems can correctly detect more bodily actions and speech utterances than ever before—but making those verbs actionable requires more advanced modeling of the verbs as meaningful within a narrative system. As with many problems in AI, this is as much a design problem as a technical one.

- **Think: improvements to player (character) modeling:** Improved modeling of players has huge potential for delightful feedback that confirms a narrative system has heard players in new and unexpected ways. In addition to the ongoing research around modeling players for improved matchmaking, monetization, harassment prevention, and so on, the space for improvements to modeling players toward narrative goals is an area ripe for research efforts. Areas of particularly low-hanging fruit include:
  - *Modeling of player character mannerisms*: mannerisms are often used as a shorthand for character traits in other media. The ability of a system to recognize when players take similar actions as mannerisms that are significant to their character is low-hanging fruit that could make players feel heard in interesting ways.
  - *Deeper modeling of player choice histories in the world* - player choices often offer counterfactuality by either branching a bespoke narrative plot moment or by incrementing a value on the player's character sheet, but alternate approaches offer low-hanging fruit. Developers might choose to categorize choices along different axes than we currently see in games (e.g. rather than "good" or "evil", "collectivist vs autonomous" and have characters who res-

onate with certain values or react to players acting contrary to the character's perception of their values), or they might have certain choices or actions change how all characters of a certain type react to them.

  - *Player character perception/knowledge/psyche*: many games model the player's "knowledge" of events by keeping track of content the player has seen, but room for deeper modeling of knowledge could offer interesting avenues. Additionally, systems that change the presentation of the game (lighting, art assets or visual effects, color palettes, etc) to align with the player character's mood or psychological state are still underexplored.

- **Think: improvements to story structure modeling:**

  - *Adaptations from one narrative structure to another based on player actions*: We have only begun to explore the space of narrative adaptations, and most successful efforts have followed a western dramatic structure that is roughly descended from Aristotle. But it seems that varying form as well as content can stave off the Oatmeal problem to some degree. Thus a system that might adapt to entirely different narrative structures based on player inputs while still maintaining a narrative that feels well-formed is a very interesting problem, but we know very little about the design challenges that cross-structural adaptations might present. Work that seems ripe to explore these questions (such as work on story sifters by Ryan [147] and Kreminski [83]) have either not integrated player choices as a component of their generation or curation processes [146], thus the generation is intended to provide situations for the player to react to rather than providing systemic adaptations to the player's choices, or in cases where player actions are integrated into the generation process, the researchers have not specifically reported on changes to story structure in their sifting processes [83]. The interchange between story sifting and responsiveness feels incredibly rich, and I am very excited to potentially see this work leveraged specifically toward responsiveness.

– *Narrative focus adaptations in response to player pursuit*: Generative narrative systems that select content via recommendation systems seem like an interesting area of study that has not reached successful adoption. Projects like *AI Dungeon* [86] show that there is an appetite among players for games that offer a wide selection of story types within a single experience, detect the kinds of stories players seem to opt into, and present more of those. Though I am wary of narratives that adapt solely based on indirect metrics like "taste" algorithms, and overfitting is a common problem in recommendation systems, nevertheless I would love to see research that explores narrative adaptation with more of an eye toward drastically varied player experiences, and the impacts of various algorithms and methods as they relate to narrative design. One could imagine a game in which the player refuses the call of the hero's journey, and the game instead adapts from pushing her toward an epic quest to instead offer lower-stakes problems within her community as she tends her garden, with the war she refused to join as a backdrop setting instead of the plot. Combining more robust player modeling with generative mechanics approaches could yield interesting experiments toward offering narrative affordances at different levels of abstraction and narrative stakes at different magnitudes based on player choices.

- **Speak: improvements to surfacing responsiveness:** The game industry has spent much design effort on how to communicate systems to players. Often this comes in the form of communicating rules, systems, and the immediate impacts of player actions, but comparatively little effort is spent on things like: how the system recognition of a player's action couples with narrative feedback; how to best communicate affordance changes, especially within generative systems; or how (much) to highlight counterfactuality within generative systems. Each of these is as much a design problem as a technical one, but they all intersect with the problem of authorial burden that technical solutions might help ameliorate.

This list is not meant to be exhaustive, nor is it meant to obfuscate the very good work toward these goals that is already underway. However, research that explicitly situates its technical improvements along these various dimensions of listening, thinking, and speaking in terms of how these improvements facilitate (1) the ability of the system to change affordances, (2) how the system impacts the player's understanding of such changes and (3) how her actions caused them offers a path for designers to understand how specific systems might be used to enable improved responsiveness and foster player agency.

More broadly, I would love to see more research efforts connect very local improvements to broader use-cases. If work claims that it will help "tailor experiences to players" or "create more reactive/dynamic/responsive narratives" this framework could be used to discuss how the research will facilitate that or offer a step toward that ultimate goal. After all, designers are generally looking for ways to get players to make interesting choices and perform interesting actions. It is important for them to understand not only how a given technical solution works, but also how they can use it to create new experiences.

10.3 Future Work

Though I offer Lume as a potential technical step toward increasing responsiveness in choice-based narrative games, it is far from the only way to integrate the principles and design understandings presented in this work. Lume is just one system, built for a very specific kind of game. Different mechanics and aesthetic goals will necessitate different approaches. Different experiences will want to optimize for different affordances and different types of affordance changes. Different developers with different backgrounds will require different authorial affordances and development optimizations in their tools. Lume is only one approach to responsive narrative, and I am excited to see others; great work to this end is already underway.

Personally, this work is far from over for me. Lume in its current form feels like only the first step toward its full potential. Further use will necessarily change it, potentially in very radical ways. It is very important to our team that we adapt our tools to the experiences we want to create rather than constrain ourselves into ill-fitting tools. As such, we have barely seen what kinds of mechanics pair well with Lume and have likely only just begun to "use it in anger", as my collaborator Ceri Stagg would say to describe the unpredictable difficulties of real-world development conditions. We have seen one potential use for it and are excited to explore others.

We have several exciting future applications and improvements planned:

1. **Broaden instructions to include other forms besides text.** As I mentioned before, the output of Lume's selection efforts is currently text and we have experimented with markup, but tests toward coupling these with generative animations have been incredibly promising. The impacts on authorial burden for such uses remains untested.

2. **Testing, automation, and debugging are areas of immediate improvement.** It is currently difficult to understand why particular content has not been selected and improvements to reporting on condition satisfaction as well as content coverage are difficult but necessary next-step improvements.

3. **My collaborator Carl Muckenhoupt is currently working on a DCG authoring tool** that will speed up the authoring of procedural text. Currently DCG authoring still requires an author to retype strings multiple times and tediously requires ensuring correct syntax. The process is not friendly to writers who are not comfortable working in code or markup. The new tool will output prolog code

automatically (as our Unity-based scene-authoring tool does), ensuring proper syntax and eliminating redundancies.

4. **For a second game, we will need to determine which rules are reusable, and which are game-specific.** It remains to be seen how much of our narrative logic is adaptable across games. Surely things like character-specific traits are specific to this story and setting, but what about our rules for item use? Relationships? Social interactions? Nemeses and antagonists? Our ultimate goal is to build up a database of rule modules that can be added piecemeal by game as developers see fit. This approach means that the engine will continue to become more robust as we develop with it, and over time we hope to use Lume across different game genres and experiences.

The commercial games industry is currently undergoing radical economic changes that will surely affect the kinds of works that will be produced in the next few years. Increasing shifts toward small-team development, movement toward livestreaming and influencer playthroughs as the dominant marketing techniques, and continuing emphasis on games as a service economic models mean that the ways games have previously told stories is more and more difficult to justify under the current economic realities. Games are looking for new ways to tell stories; narrative designers are looking for new tools and technologies for inspiration. Games are finally starting to recognize their appeal to audiences across demographics and are looking for new voices, new verbs, and new core pleasure-experiences. All of these will require new design paradigms and new ways to integrate stories and play.

I remain excited by the potential of these new paradigms. Stories and games were never at odds, as any small child proves every time they change the rules in their game of makebelieve to suit their preferred narrative affordances, but now more than ever our industry and academic research communities seem ready to acknowledge it. We have work ahead of us, but I am excited by its potential. In addition to new technical innovations, we will need new designs and new theories. But I cannot wait to meet with them and to see how we might create the magic of responsiveness in new forms yet undiscovered.

Bibliography

[1] Espen J. Aarseth. 1997. *Cybertext: Perspectives on Ergodic Literature.* Johns Hopkins University Press, Baltimore, MD, USA and London, UK.

[2] Jem Alexander. 2018. *Spirit AI's technology could bring Westworld-style interactions to video games sooner than you think.* https://gamedaily.biz/article/129/spirit-ais-technology-could-bring-westworld-style-interactions-to-video-games-sooner-than-you-think

[3] Leigh Alexander. 2012. GDC 2012: Sid Meier on how to see games as sets of interesting decisions. (2012). https://www.gamasutra.com/view/news/164869/GDC_2012_Sid_Meier_on_how_to_see_games_as_sets_of_interesting_decisions.php

[4] Leonardo Electronic Almanac. [n. d.]. *About LEA: History.* https://web.archive.org/web/20160303190346/http://leoalmanac.org/about/history.asp Page accessed through the Internet Archive, archive date Aug 16, 2006.

[5] Sam Kabo Ashwell. 2014. *A Bestiary of Player Agency.* https:

//heterogenoustasks.wordpress.com/2014/09/22/a-bestiary-of-player-agency/

[6] Sam Kabo Ashwell. 2015. *Standard Patterns in Choice-Based Games.* https://heterogenoustasks.wordpress.com/2015/01/26/standard-patterns-in-choice-based-games/

[7] Atari. 1980. *Missile Command.* Arcade, Atari 2600, others. Playable Media Application.

[8] The Game Band. 2020. *Blaseball.* Web Broswer. https://www.blaseball.com/ Playable Media Application.

[9] Nate Barker, Anala Prabhu, and Steve Webb. 2019. Monetization Strategies That Drive Value. *Game Developer's Conference 2019.*

[10] Sam Barlow. 2015. *Her Story.* MS Windows, Mac iOS, Mac OS X, others. Playable Media Application.

[11] Belinda Barnet. 2013. *Memory Machines The Evolution of Hypertext.* Anthem Press, London; New York; Delhi. http://www.jstor.org/stable/j.ctt1gxpbkb

[12] Roland Barthes. 1978. The Death of the Author. In *Image-Music-Text.* Hill and Wang.

[13] Richard Bartle. 1999. *Hearts, Clubs, Diamonds, Spades: Players Who Suit MUDs.* https://mud.co.uk/richard/hcds.htm

[14] Frank L. Baum. 1987 (orig 1900). *The Wonderful Wizard of Oz*. Harper Collins, Singapore.

[15] Mark Bernstein. 2001. Card Shark and Thespis: Exotic Tools for Hypertext Narrative. In *Proceedings of the 12th ACM Conference on Hypertext and Hypermedia (HYPERTEXT '01)*. Association for Computing Machinery, New York, NY, USA, 41–50. https://doi.org/10.1145/504216.504233

[16] Mark Bernstein. 2002. Storyspace 1. In *Proceedings of the Thirteenth ACM Conference on Hypertext and Hypermedia (HYPERTEXT '02)*. Association for Computing Machinery, New York, NY, USA, 172–181. https://doi.org/10.1145/513338.513383

[17] Mark Bernstein and Diane Greco. 2008. Designing a New Media Economy. *Genre: Forms of Discourse and Cultre.* 41 (2008), 59–82. Issue 3-4.

[18] Bill Bly. 1997. *We Descend*. Mac OS, MS Windows. Playable Media Application.

[19] Ian Bogost. 2017. Video Games Are Better Without Stories. (2017). https://www.theatlantic.com/technology/archive/2017/04/video-games-stories/524148/

[20] Ian Buchanan. 2010. fabula and sjužet. (2010). https://www.oxfordreference.com/view/10.1093/oi/authority.20110803095807354 Online excerpt: Oxford Reference.

[21] Matthew Byrd. 2018. The Walking Dead: The False Promises of Telltale's Masterpiece. (2018). https://www.denofgeek.com/games/the-walking-dead-false-promises-telltale/

[22] Galactic Cafe. 2011. *The Stanley Parable*. MS Windows, Mac OS X, Linux. Playable Media Application.

[23] Roger Callois. 2001 (orig 1958). *Man, Play and Games*. University of Illinois Press, Urbana and Chicago, IL, USA.

[24] Capcom. 2016. *Street Fighter V*. Sony Playstation 4, MS Windows. Playable Media Application.

[25] Rogelio Enrique Cardona-Rivera. 2019. *A Model of Interactive Narrative Affordances*. Doctoral Dissertation.

[26] R. E. Cardona-Rivera and R. M. Young. 2013. A Cognitivist Theory of Affordances for Games. In *DiGRA Conference*.

[27] Drew Carey. 2011. *Improv-A-Ganza*. Television show. Scene with performers Colin Mochrie and Ryan Stiles. Season 1, Episode 1. Original Air Date April 11, 2011 on Game Show Network.

[28] Drew Carey. 2011. *Improv-A-Ganza*. Television show. Aired from April 11, 2011 to June 3, 2011 on Game Show Network.

[29] Lewis Carroll. 2015 (orig 1865). *Alice's Adventures in Wonderland*. Penguin Random House, New York, NY, USA.

[30] Seymour Chatman. 1980. *Story and Discourse: Narrative Structure in Fiction and Film.* Cornell University Press, Ithica, NY, USA.

[31] Coilworks. 2014. *Cloudbuilt.* MS Windows. Playable Media Application.

[32] Kate Compton. 2016. *So you want to build a generator...* http://www.galaxykate.com/buildagenerator-kcompton.pdf

[33] Kate Compton, Quinn Kybartas, and Michael Mateas. 2015. Tracery: An Author-Focused Generative Text Tool. 154–161. https://doi.org/10.1007/978-3-319-27036-4_14

[34] Robert Coover. 1992. "The End of Books". *The New York Times* (June 21, 1992 1992).

[35] Chris Crawford. 2005. *Chris Crawford on Interactive Storytelling.* New Riders, Pearson Education, San Francisco, CA, USA.

[36] Extra Credits. 2014. *Design Clug - Super Mario Bros: Level 1-1 - How Super Mario Mastered Level Design.* https://www.youtube.com/watch?v=ZH2wGpEZVgE

[37] Mihaly Csikszentmihalyi. 2008 (orig 1990). *Flow: The Psychology of Optimal Experience (Harper Perennial Modern Classics).* Harper Perrennial, New York, NY, USA.

[38] Guillermo del Toro. 2006. *Pan's Labyrinth.* Warner Bros. Pictures.

[39] Nintendo Entertainment Analysis & Development Division. 1971. *The Legend of Zelda: Majora's Mask*. HP 2100. Playable Media Application.

[40] Quantic Dream. 2018. *Detroit: Become Human*. Sony Playstation 4, MS Windows. Playable Media Application.

[41] Roger Ebert. 2010. *Video games can never be art*. `https://www.rogerebert.com/roger-ebert/video-games-can-never-be-art`

[42] Blizzard Entertainment. 2010. *Starcraft II: Wings of Liberty*. Windows, Mac OS. Playable Media Application.

[43] Blizzard Entertainment. 2016. *Overwatch*. MS Windows, Microsoft Xbox One, Sony Playstation 4. Playable Media Application.

[44] Dontnod Entertainment. 2015. *Life is Strange*. MS Windows, Mac OSX, Sony Playstation 3, Microsoft Xbox 360, others. Playable Media Application.

[45] Liam Esler. 2016. *Finding the Right Branching Narrative Structure For Your RPG*. `https://www.gdcvault.com/play/1023386/Finding-the-Right-Branching-Narrative`

[46] Richard Evans, Emily Short, and Graham Nelson. 2014. *Versu*. `https://versu.com/about/` Software Application.

[47] Clara Fernandez-Vara. 2019. *Taxonomy of Narrative Choices*. `https://clarafv.com/2019/05/01/taxonomy-of-narrative-choices/`

[48] Cyril Focht. 2019. *Sonder: an examination of how choice creates meaning and the narrative effects of agency.* Masters Thesis.

[49] Quantic Foundry. 2019. *Gamer Motivation Model.* https://quanticfoundry.com/wp-content/uploads/2019/04/Gamer-Motivation-Model-Reference.pdf

[50] FromSoftware. 2014. *Dark Souls II.* Playstation 3, Xbox 360, MS Windows, Others. Playable Media Application.

[51] Fullbright. 2013. *Gone Home.* MS Windows, Mac OSX, Linux, Others. Playable Media Application.

[52] Doublespeak Games. 2013. *A Dark Room.* Web browsers, Android, Mac iOS, Nintendo Switch. https://adarkroom.doublespeakgames.com/ Playable Media Application.

[53] Failbetter Games. 2009. *Fallen London.* Web browsers. https://www.fallenlondon.com/login Playable Media Application.

[54] Failbetter Games. 2010. Echo Bazaar Narrative Structures, part three. (2010). https://www.failbettergames.com/echo-bazaar-narrative-structures-part-three/

[55] Failbetter Games. 2014. *Sunless Seas.* Microsoft Windows, Mac OSX, Others. Playable Media Application.

[56] Hanako Games. 2012. *Long Live the Queen.* MS Windows, Mac OS X, Linux. Playable Media Application.

[57] Irrational Games. 2007. *BioShock.* MS Windows, Microsoft Xbox 360, Sony Playstation 3, others. Playable Media Application.

[58] Supergiant Games. 2011. *Bastion.* Nintendo Switch, MS Windows, Others. Playable Media Application.

[59] Telltale Games. 2012. *The Walking Dead.* Windows, Mac OS X, Linux, Others. Playable Media Application.

[60] Telltale Games. 2017. *Batman: The Enemy Within.* Windows, Mac OS X, Linux, Others. Playable Media Application.

[61] Jacob Garbe. 202. *Increasing Authorial Leverage in Generative Systems.* Doctoral Dissertation.

[62] Jacob Garbe, Max Kreminski, Ben Samuel, Noah Wardrip-Fruin, and Michael Mateas. 2019. StoryAssembler: An Engine for Generating Dynamic Choice-Driven Narratives. In *Proceedings of the 14th International Conference on the Foundations of Digital Games (FDG '19).* Association for Computing Machinery, New York, NY, USA, Article 24, 10 pages. `https://doi.org/10.1145/3337722.3337732`

[63] Jacob Garbe, Aaron A. Reed, Melanie Dickinson, Noah Wardrip-Fruin, and M. Mateas. 2014. Assistance Visualizations for IceBound , A Combinatorial Narrative.

[64] John Gardner. 1991 (orig 1983). *The Art of Fiction: Notes on Craft for Young Writers*. Vintage Books, New York, NY, USA.

[65] Alex Garland. 2014. *Ex Machina*. Film 4, DNA Films.

[66] William Gaver. 1991. Technology Affordances. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 79–84. https://doi.org/10.1145/108844.108856

[67] James J. Gibson. 2015 (orig 1986). *The Ecological Approach to Visual Perception: Classic Edition*. Taylor & Francis, New York, NY, USA.

[68] Gary Gygax and Dave Arneson. 1978. *Dungeons & Dragons*. Series of Tabletop Game Rule Books.

[69] Juho Hamari and Janne Tuunanen. 2014. Player Types: A Meta-synthesis. *Transactions of the Digital Games Research Association* 1 (03 2014), 29–53. https://doi.org/10.26503/todigra.v1i2.13

[70] Dan Harmon. 2009 - 2015. *Community*. Television show. Aired from September 17, 2009 to June 2, 2015 on NBC (Sony Pictures Television).

[71] Porpentine Charity Heartscapre and Brenda Neotenomie. 2007. *With Those We Love Alive*. Web Browser. http://slimedaughter.com/games/twine/wtwla/ Playable Media Application.

[72] Johan Huizinga. 2016 (orig 1949). *Homo Ludens: A Study of the Play-Element in Culture*. Angelico Press.

[73] Robin Hunicke, Marc LeBlanc, and Robert Zubek. 2004. MDA: A formal approach to game design and game research. In *Proceedings of the AAAI Workshop on Challenges in Game AI*, Vol. 4. San Jose, CA, 17–22.

[74] Inkle. 2014. *80 Days*. iOS, Android, Microsoft Windows, Mac OSX. Playable Media Application.

[75] Shelley Jackson. 1995. *Patchwork Girl*. Mac OS, USB Stick. Playable Media Application.

[76] Spike Jonze. 2013. *Her*. Annapurna Pictures.

[77] Michael Joyce. 1995. *Of Two Minds: Hypertext Pedagogy and Poetics*. University of Michigan Press, Ann Arbor, MI, USA.

[78] Michael Joyce. 1997. Nonce Upon Some Times: Rereading Hypertext Fiction. *MFS Modern Fiction Studies* 43, Fall (1997), 579–597. https://doi.org/10.1353/mfs.1997.0061

[79] Jesper Juul. 2013. Without a goal. In *Videogame, Player, Text*, Barry Atkins and Tanya Krzywinska (Eds.). Manchester University Press. https://doi.org/10.1007/978-3-319-02756-2_3

[80] Isaac Karth. 2014. *Ergodic Agency: How Play Manifests Understanding*. Brill, 205–216. https://doi.org/10.1163/9781848882959_020

[81] Hartmut Koenitz. 2014. Five Theses for Interactive Digital Narrative, Vol. 8832. 134–139.

[82] Hartmut Koenitz, Christian Roth, Noam Knoller, and Teun Dubbelman. 2018. " Clementine will remember that"-Methods to Establish Design Conventions for Video Game Narrative.. In *DiGRA Conference*.

[83] Max Kreminski, Melanie Dickinson, and Noah Wardrip-Fruin. 2019. *Felt: A Simple Story Sifter*. 267–281. https://doi.org/10.1007/978-3-030-33894-7_27

[84] Max Kreminski and Noah Wardrip-Fruin. 2018. *Sketching a Map of the Storylets Design Space*. Springer International Publishing, 160–164. https://doi.org/10.1007/978-3-030-04028-4_14

[85] Ben Kybartas and Rafael Bidarra. 2017. A Survey on Story Generation Techniques for Authoring Computational Narratives. *IEEE Transactions on Computational Intelligence and AI in Games* 9, 3 (2017), 239–253. https://doi.org/10.1109/TCIAIG.2016.2546063

[86] Latitude. 2019. *AI Dungeon*. Web Browser. https://play.aidungeon.io/main/home Playable Media Application.

[87] Brenda Laurel. 2014. *Computers as Theatre (Second Edition)*. Addison-Wesley, Upper Saddle River, NJ.

[88] Jonas Linderoth. 2013. Beyond the digital divide: An ecological approach to gameplay. *Transactions of the Digital Games Research Association* 1 (04 2013). https://doi.org/10.26503/todigra.v1i1.9

[89] Judy Malloy. 1956-2010. Judy Malloy papers, 1956 - 2010. (1956-2010). https://archives.lib.duke.edu/catalog/malloyjudy

[90] Judy Malloy. 1989. *its name was Penelope.* DOSBox Emulator and GWBASIC. Playable Media Application.

[91] Judy Malloy. 2012 (orig 1988). *Uncle Roger.* DOSBox Emulator and GWBASIC. https://collection.eliterature.org/3/works/uncle-roger/uncle.html#basic Playable Media Application.

[92] Judy Malloy. 2021. *Judy Malloy - Electronic Literature, Hyperfiction, Generative Poetry, Polychoral Literature.* http://www.judymalloy.net/

[93] Molly Maloney. 2017. *Out on a Limb: Practical Approaches to Branching Story.* https://www.youtube.com/watch?v=4-MGT6SZWWo Presentation given at Konsoll 2017: Bergen, Norway. October 5-6, 2017.

[94] Cat Manning. 2018. *Successful Reflective Choices in Interactive Narrative.* https://catacalypto.wordpress.com/2018/06/19/successful-reflective-choices-in-interactive-narrative/

[95] Stacey Mason. 2013. On Games and Links: Extending the Vocabulary of Agency and Immersion in Interactive Narratives. In *Interactive Storytelling.* Springer International Publishing, 25–34. https://doi.org/10.1007/978-3-319-02756-2_3

[96] Stacey Mason. 2018. *Clip: "Oh my God, that should have been*

a choice!". https://www.twitch.tv/cerebralarcade/clip/LivelySparklingCormorantWutFace?filter=clips&range=all&sort=time Recording: gameplay performance for Cerebral Arcade stream series.

[97] Stacey Mason and Mark Bernstein. 2019. On Links: Exercises in Style. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media (HT '19)*. Association for Computing Machinery, New York, NY, USA, 103–110. https://doi.org/10.1145/3342220.3343665

[98] Stacey Mason, Ceri Stagg, Noah Wardrip-Fruin, and Michael Mateas. 2019. Lume: A System for Procedural Story Generation. In *Proceedings of Foundations of Digital Games Conference* (August 26–30, 2019). ACM, San Luis Obispo, CA, USA. https://doi.org/10.1145/3337722.3337759

[99] Michael Mateas. 2001. A preliminary poetics for interactive drama and games. *Digital Creativity* 12, 3 (2001), 140–152. https://doi.org/10.1076/digc.12.3.140.3224

[100] Michael Mateas. 2002. *Interactive Drama, Art and Artificial Intelligence*. Doctoral Dissertation.

[101] Michael Mateas. 2003. Expressive AI: A semiotic analysis of machinic affordances. In *3rd Conference on Computational Semiotics for Games and New Media*, Vol. 58. Citeseer.

[102] Michael Mateas. 2021. *UC Santa Cruz, Computational Media. Lecture slides CMPM148: Interactive Narrative.* "Free will and Mateas's Demon.

[103] Michael Mateas and Andrew Stern. 2002. A Behavior Language for Story-based Believable Agents. In *AAAI Technical Report SS-02-01.* https://www.aaai.org/Papers/Symposia/Spring/2002/SS-02-01/SS02-01-014.pdf

[104] Michael Mateas and Andrew Stern. 2003. Façade: An Experiment in Building a Fully-Realized Interactive Drama. (04 2003).

[105] Michael Mateas and Andrew Stern. 2003. Façade: An Experiment in Building a Fully-Realized Interactive Drama. (04 2003). https://www.cc.gatech.edu/~isbell/reading/papers/MateasSternGDC03.pdf

[106] Michael Mateas and Andrew Stern. 2005. *Façade.* MS Windows, Mac OS. Playable Media Application.

[107] Peter A. Mawhorter. 2016. *Artificial Intelligence as a Tool for Understanding Narrative Choices.* Doctoral Dissertation.

[108] Maxis. 2000. *The Sims.* MS Windows, Mac OS, Others. Playable Media Application.

[109] Josh McCoy, Mike Treanor, Ben Samuel, Michael Mateas, and Noah Wardrip-Fruin. 2011. ¡i¿Prom Week¡/i¿: Social Physics as Gameplay. In *Proceedings of the 6th International Conference on Foundations of Digital Games (FDG '11).*

Association for Computing Machinery, New York, NY, USA, 319–321. https://doi.org/10.1145/2159365.2159425

[110] Joshua McCoy, Mike Treanor, Ben Samuel, Noah Wardrip-Fruin, and Michael Mateas. 2011. Comme il Faut: A System for Authoring Playable Social Models. *Proceedings of the 7th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2011.*

[111] Joanna Mcgrenere and Wayne Ho. 2000. Affordances: Clarifying and Evolving a Concep. *Proceedings of the Graphics Interface 2000 Conference*, 179–186.

[112] Robert McKee. 2009. *Story: Style, Structure, Substance, and the Principles of Screenwriting.* MIT Press, Cambridge, MA, USA.

[113] Lauren McManamom and Jesse Ross. 2019. *Girl Underground.* Tabletop Game Rule Book.

[114] Nick Montfort. 2011. Curveship's Automatic Narrative Style. (06 2011). https://doi.org/10.1145/2159365.2159394

[115] Judd Morrissey and Lori Tally. 2000. *The Jew's Daughter.* Macromedia Flash. https://collection.eliterature.org/1/works/morrissey__the_jews_daughter.html Playable Media Application.

[116] Steve Mullis. 2013. In The Stanley Parable, Finding The Story Is The Game. (2013). https://www.npr.org/sections/alltechconsidered/2013/11/19/246213397/in-the-stanley-parable-finding-the-story-is-the-game

[117] Daniel Muriel and Garry Crawford. 2020. Video Games and Agency in Contemporary Society. *Games and Culture* 15, 2 (2020), 138–157. https://doi.org/10.1177/1555412017750448 arXiv:https://doi.org/10.1177/1555412017750448

[118] Janet H. Murray. 1997. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. Simon & Schuster, New York, NY, USA.

[119] Janet H. Murray. 2016 (orig 1997). *Hamlet on the Holodeck, Revised Edition: The Future of Narrative in Cyberspace*. Simon & Schuster, New York, NY, USA.

[120] Graham Nelson. 2006 - 2015. *Inform 7*. Windows, Mac OS X, Linux, Others. http://inform7.com/ Software Application.

[121] Nerial. 2016. *Reigns*. Android, iOS, others. Playable Media Application.

[122] Neryssa. 2021. Overview of Zone Changes Between Classic and Cataclysm. (2021). https://classic.wowhead.com/guides/zone-changes-from-classic-to-cataclysm

[123] Jonathan Nolan and Lisa Joy. 1987 - 1994. *Westworld*. Television show. Aired from September 28, 1987 to May 23, 1994 on CBS (Paramount Television).

[124] Donald Norman. 1999. Affordance, conventions, and design. *Interactions* 6 (05 1999), 38–42. https://doi.org/10.1145/301153.301168

[125] Donald A. Norman. 1988. *The Psychology of Everyday Things*. Basic Books, New York, NY, USA.

[126] Donald A. Norman. 2013 (orig 1988). *The Design of Everyday Things:Revised and Expanded Edition.* Basic Books, New York, NY, USA.

[127] Nicholas O'Brien. 2021. Opinion: When choice becomes a metric, narrative design suffers. (2021). https://www.gamasutra.com/view/news/378193/Opinion_When_choice_becomes_a_metric_narrative_design_suffers.php

[128] Joseph Osborn, James Ryan, and Michael Mateas. 2017. Analyzing Expressionist Grammars by Reduction to Symbolic Visibly Pushdown Automata.

[129] Critical Path. 2017. *Ernest Adams - Agency vs Story.* https://www.youtube.com/watch?v=G4G88LzOP_8

[130] Bernard Perron and Dominic Arsenault. 2008. In the Frame of the Magic Cycle: the Circle (s) of Gameplay. *The video game theory reader* 2 (2008), 109–131.

[131] Kojima Productions. 2008. *Konami.* Sony Playstation 3. Playable Media Application.

[132] Alex Proyas. 2004. *iRobot.* 20th Century Fox.

[133] Punchdrunk. 2009. *Sleep No More.* Interactive Theater Performance. Attended performance October 20, 2009. Old Lincoln School: Brookline, MA, USA.

[134] Matt Purslow. 2021. Disco Elysium: Bringing a Million Words to Life for the Final Cut. (2021). https://www.ign.com/articles/disco-elysium-the-final-cut-voice-acting-console-release

[135] The Quinnspiracy. 2013. *Depression Quest*. Windows, Mac OS X, Linux, Others. Playable Media Application.

[136] Don Rawitsch, Bill Heinemann, and Paul Dillenberger. 1971. *The Oregon Trail*. HP 2100. Playable Media Application.

[137] Nintendo R&D4. 1985. *Super Mario Bros.* Nintendo Entertainment System. Playable Media Application.

[138] Nintendo R&D4. 1986. *The Legend of Zelda*. Nintendo Entertainment System. Playable Media Application.

[139] Aaron Reed and Jacob Garbe. 2016. *The Ice-Bound Concordance*. Mac iOS, Physical Book. Playable Media Application.

[140] Aaron A. Reed. 2017. *Changeful Tales: Design-Driven Approaches Toward More Expressive Storygames*. Doctoral Dissertation.

[141] Aaron A. Reed. 2020. *Subcutanean*. Procedurally generated novel, print-on-demand. Documentation at https://aaronareed.net/subcutanean-book/.

[142] Mark Riedl and Vadim Bulitko. 2013. Interactive Narrative: An Intelligent Systems Approach. *AI Magazine* 34 (03 2013), 67–77. https://doi.org/10.1609/aimag.v34i1.2449

[143] Gene Roddenberry. 2016 - current as of 2021. *Star Trek: The Next Generation*. Television show. Aired from beginning to October 2, 2016, currently running as of 2021. on HBO (Warner Bros. Television).

[144] Jonathan Rowe and James Lester. 2010. Modeling User Knowledge with Dynamic Bayesian Networks in Interactive Narrative Environments. *Proceedings of the 6th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE 2010*.

[145] Steve Russell. 1962. *Spacewar!* PDP-1. Playable Media Application.

[146] James Ryan. 2018. *Curating Simulated Storyworlds*. Doctoral Dissertation.

[147] James Ryan and Michael Mateas. 2017. *Simulating Character Knowledge Phenomena in Talk of the Town*. 433–448. https://doi.org/10.4324/9781315151700-37

[148] James Ryan, Ben Samuel, Adam Summerville, and Jonathan Lessard. 2015. Bad News: A Computationally Assisted Live-Action Prototype to Guide Content Creation.

[149] James Ryan, Ethan Seither, Michael Mateas, and Noah Wardrip-Fruin. 2016. Expressionist: An Authoring Tool for In-Game Text Generation. https://doi.org/10.1007/978-3-319-48279-8_20

[150] Marie-Laure Ryan. 2015. *Narrative as Virtual Reality 2: Revisiting Immersion and Interactivity in Literature and Electronic Media*. Johns Hopkins University Press, Baltimore, Maryland, USA.

[151] Anastasia Salter and Stuart Moulthrop. 2021. *Twining: Critical and Creative Approaches to Hypertext Narratives*. Amherst College Press, Amhesrt, MA, USA.

[152] Ben Samuel. 2016. *Crafting Stories Through Play.* Doctoral Dissertation.

[153] Jason Schreier. 2011. Review: Narrator's Constant Voice Adds Human Touch to Bastion. (2011). https://www.wired.com/2011/07/bastion-review/

[154] A Sharp. 1999. *King of Dragon Pass.* MS Windows, Mac OS. Playable Media Application.

[155] Emily Short. 2000. *Galatea.* MS Windows. , Z-Interpreter pages. Playable Media Application.

[156] Emily Short. 2014. *Blood & Laurels.* Mac iOS. Playable Media Application.

[157] Emily Short. 2016. *Mark Bernstein on Hypertext Narrative.* https://emshort.blog/2016/04/28/mark-bernstein-on-hypertext-narrative/

[158] Emily Short. 2019. *Links and Structures from Michael Joyce to Twine.* https://emshort.blog/2019/07/27/michael-joyce-on-hypertext-links/

[159] Emily Short. 2020. *Casual Games and Storylets: Or, How to Make Game Mechanics Express Choice.* https://emshort.blog/2020/01/09/casual-games-and-storylets-or-how-to-make-game-mechanics-express-choice/

[160] Emily Short. 2021. *Emily Short's Interactive Storytelling.* https://emshort.blog/

[161] Emily Short. 2021. *Personal communication via email.*

[162] Emily Short, John Cater, Rob Dubbin, Eric Eve, Elizabeth Heller, Jayzee, Kazuki Mishima, Sarah Morayati, Mark Musante, Adam Thornton, and Ziv Wities. 2009. *Alabaster*. http://iplayif.com/?story=http%3A%2F%2Fwww.ifarchive.org%2Fif-archive%2Fgames%2Fglulx%2FAlabaster.gblorb Playable Media Application.

[163] Tanya X. Short. 2016. Procedurally Generating Personalities. (2016). https://www.gamasutra.com/blogs/TanyaXShort/20161216/310387/Procedurally_Generating_Personalities.php

[164] Tanya X. Short and Tarn Adams (Eds.). 2019. *Procedural Storytelling in Game Design*. Taylor & Francis, Boca Raton, FL, USA.

[165] Square. 2002. *Kingdom Hearts*. Sony Playstation 2. Playable Media Application.

[166] Dietrich Squinkifer. 2013. *Imposter Syndrome*. Web browsers. https://www.deirdrakiai.com/impostor/ Playable Media Application.

[167] Dietrich Squinkifer. 2014. *Coffee: A Misunderstanding*. Live performance.Documentation at https://squinky.me/2014/06/23/coffee-a-misunderstanding/. Playable Media Application.

[168] Dietrich Squinkifer. 2014. *Interruption Junction*. Web Browser. https://www.2015.differentgames.org/interruption-junction/ Playable Media Application.

[169] Game Studies. 2001 - 2021. *Game Studies: The International Journal of Computer Game Research*. http://gamestudies.org/2102/archive

[170] Bethesda Game Studio. 2011. *The Elder Scrolls V: Skyrim*. MS Windows, Sony Playstation 3, Microsoft XBox 360, others. Playable Media Application.

[171] Paradox Development Studio. 2020. *Crusader Kings 3*. MS Windows, MacOS, Linux. Playable Media Application.

[172] Studiobinder. 2020. How the Dan Harmon Story Circle Can Make Your Story Better. (2020). https://www.studiobinder.com/blog/dan-harmon-story-circle/

[173] Twinbeard Studios. 2012. *Frog Fractions*. Adobe Flash, MS Windows. Playable Media Application.

[174] David Sudnow. 1983. *Pilgrim in the Microworld*. Warner Books, New York, NY, USA.

[175] David Sudnow. 2001 (orig 1978). *Ways of the Hand*. MIT Press, Cambridge, MA, USA.

[176] Adam Summerville, Chris Martens, Sarah Harmon, Michael Mateas, Joseph Osborn, Noah Wardrip-Fruin, and Arnav Jhala. 2019. From Mechanics to Meaning. *IEEE Transactions on Games* 11, 1 (2019), 69–78. https://doi.org/10.1109/TCIAIG.2017.2765599

[177] Beat Suter, Rene Bauer, and Mela Kocher (Eds.). 2021. *Narrative Mechanics - Strategies and Meanings in Games and Real Life.* Transcript Verlag: Media Studies.

[178] Steve Swink. 2009. *Game Feel: A Game Designer's Guide to Virtual Sensation.* Elsevier, Burlington, MA, USA.

[179] Karen Tanenbaum and Theresa Jean Tanenbaum. 2010. Agency as commitment to meaning: communicative competence in games. *Digital Creativity* 21, 1 (2010), 11–17. https://doi.org/10.1080/14626261003654509

[180] Pitapot The Blaseball Community: Contributors Nesblitt, ItsSteve. 2021. *Beginner's Guide to Blaseball.* https://www.blaseball.wiki/w/Beginner%27s_Guide_to_Blaseball

[181] Thekla. 2016. *The Witness.* MS Windows, Sony Playstation 4, Microsoft XBox one, others. Playable Media Application.

[182] David Thue, Vadim Bulitko, and Marcia Spetch. 2008. PaSSAGE: A Demonstration of Player Modeling in Interactive Storytelling.

[183] David Thue, Vadim Bulitko, Marcia Spetch, and Eric Wasylishen. 2007. Interactive Storytelling: A Player Modelling Approach. 43–48.

[184] Walter F. Tichy, Paul Lukowicz, Lutz Prechelt, and Ernst A. Heinz. 1995. Experimental evaluation in computer science: A quantitative study. *Journal of*

Systems and Software 28, 1 (1995), 9–18. https://doi.org/10.1016/0164-1212(94)00111-y

[185] Rowan Tulloch. 2014. The construction of play: Rules, restrictions, and the repressive hypothesis. *Games and Culture* 9, 5 (2014), 335–350.

[186] Valve. 2004. *Half-Life 2*. Windows, Mac OS X, Linux, Others. Playable Media Application.

[187] Valve. 2010. *Portal*. MS Windows, Sony Playstation 3, Microsoft Xbox 360, others. Playable Media Application.

[188] Anders Walter. 2018. *I Kill Giants*. RLJE Films.

[189] Noah Wardrip-Fruin. 1997. *Expressive Processing: Digital Fictions, Computer Games, and Software Studies*. Harper Collins, New York, NY, USA.

[190] Noah Wardrip-Fruin. 2005. *Playable Media and Textual Instruments*. http://www.dichtung-digital.de/2005/1/Wardrip-Fruin/index.htm

[191] Noah Wardrip-Fruin. 2020. *How Pac-Man Eats*. MIT Press, Cambridge, MA, USA.

[192] Noah Wardrip-Fruin and Michael Mateas. 2014. *Envisioning the Future of Computational Media: The Final Report of the Media Systems Project*. Report. Center for Games and Playable Media at the University of California, Santa Cruz.

[193] Noah Wardrip-Fruin, Michael Mateas, Steven Dow, and Serdar Sali. 2009. Agency

Reconsidered. In *Proceedings of DiGRA 2009 (Breaking New Ground: Innovation in Games, Play, Practice and Theory.)*. Citeseer.

[194] Ben Weber. 2016. *Game Analytics & Machine Learning.* Guest Lecture: UC Santa Cruz, Computational Media Dept. CMPM146: March 2016.

[195] Joseph Weizenbaum. 1966. ELIZA—a computer program for the study of natural language communication between man and machine. *Commun. ACM 9, 1* (1966), 36–45.

[196] ZA/UM. 2019. *Disco Elysium.* MS Windows, Microsoft Xbox One, Sony Playstation 4, Google Stadia, others. Playable Media Application.