UNIVERSITY OF CALIFORNIA

SANTA CRUZ

SEARCHING WITHIN AND ACROSS GAMES

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

DOCTOR OF PHILOSOPHY

in

COMPUTATIONAL MEDIA

by

Barrett Anderson

December 2020

The Dissertation of Barrett Anderson
is approved:

_____

Professor Noah Wardrip-Fruin, chair

_____

Professor Adam M. Smith

_____

Professor Sri Kurniawan

_____

Henry Lowood, Ph.D.

_____

Quentin Williams

Acting Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

List of Tables

Abstract

Searching Within and Across Games

Media like videogames (and novels, film, music, etc.) play an important role in most people's lives, and creation and comprehension of these works benefit from critical study. Scholarly analysis of media has been facilitated by information retrieval technology, which not only saves time, but also makes it possible to ask new kinds of questions. These benefits are relevant for any students or educators whose work includes the creation and interpretation of interactive media. However, current approaches treat games as monolithic artifacts because they are not capable of ingesting and crawling the content of games, and even modern search engines are better at capturing secondary text-based sources about games such as criticism, guides, and reviews. Moments that can take place within a game could instead be treated like pages in a book or like sites on the web. In this dissertation I describe research that identifies the needs of scholars and other experts for videogame moment retrieval, the development of two interactive systems that can meet a subset of those needs, and a publicly available corpus of student games that can enable artificial intelligence, information retrieval, and digital humanities research. My work contributes a new technical method to game studies, and a new domain of complex and culturally relevant data to information retrieval. Finally, my work contributes to game education with the design of search and visualization tools that could be used to analyze student game projects, to inspire future student work, and potentially to be incorporated directly into future student work processes.

Acknowledgements

This work would not have been possible without the assistance and encouragement of many individuals. My advisor, Adam M. Smith, helped me find a new academic home in his Design Reasoning Lab, and his personal support has been unwavering even when my confidence in myself has not. Adam's aplomb, adaptiveness, and positivity encouraged me to stay engaged, and to pick myself up after every setback. I do not believe I could have had a better mentor, and I am excited to the be the first of what I am sure will be many successful students.

Before I met Adam, Noah Wardrip-Fruin made me feel welcome in the Computational Media department. He introduced me to research communities, that I had not even imagined to exist, that would value the kind of work that I wanted to do. Elizabeth Swensen was also incredibly supportive in this transitional time. I believe that I would not have made it this far if not for the encouragement I received from my first academic mentor, Philippe Goldin. Adam, Noah, Elizabeth, and Philippe's examples of patience, supportiveness, kindness, and playfulness go beyond supporting me as a researcher, and have been a touchstone for me in deciding the kind of person I want to be.

This era of my life has been a joyful one because I have been able to share the journey. Christy Starr, thank you for being my friend, for all our deep conversations, and for being there for me in so many ways. Thank you to Lee Taber for our Zoom

writing sessions. Consistently applying myself to writing while working from home would have been much more difficult without them. Thank you also to Adam Rausch, Melanie Dickinson, Katie Green, Max Kreminski, Robert Kutner, Johnathan Pagnutti, and Stella Mazeika, for all of the games, conversations, and camaraderie. Thank you to Jasmine Otto, for your patience, excitement, understanding, and companionship.

I am grateful to my mother Linda Lees Dwyer, my father Scott Anderson, and my grandmother Anne Anderson, from whom I have received ongoing material and emotional support. I am confident that I have done more than I otherwise would have due to this support, because of the worries I did not have.

My grandfather, David Anderson, passed away before I came to UC Santa Cruz, but without him no part of this dissertation would have come into existence. Before I went to college, he and I spent time together studying world religions and philosophy, which started me out on the life of the mind. He provided some of my first examples of curiosity, compassion, and intellectual playfulness, and I would like to think that if he saw where I am today that he would have been proud. This has been a particularly challenging year, marked by violence against students from the UC administration, the loss of homes to the wildfires in California, and the ongoing disruption to so many aspects of life in a global pandemic. It is good to be able to bring it to a close with feelings of accomplishment, and to share the immense gratitude I have for those who made these accomplishments possible.

Chapter 1. Introduction

Media can inspire empathy, improve comprehension of a novel situation, contribute to the development of critical skills, and have a variety of other positive (and negative) impacts on those who engage with it (Sherry, 2013; Subrahmanyam et al., 2001; Young et al., 2012). Games are a popular form of media, with 2.5 billion players worldwide (2018 Video Game Industry Statistics, Trends & Data - The Ultimate List, 2018) and including players in more than 70% of all US households ("2019 Essential Facts About the Computer and Video Game Industry," 2019). Scholarly work contributes to our understanding of the personal and social effects of all forms of media, which is important both for consuming media critically and for thoughtful media creation. Relative to traditional media such as text and film, games and other kinds of playable media are understudied, both because of their novelty and because the ways that they differ from traditional media may require new approaches. For example, only recently have standards been developed for cataloguing and citing games in an academic context (Kaltman et al., 2016, 2017). Even at the level of citing an entire game, the proliferation of hardware and software variations mean that these standards face substantial challenges (Altice, 2015). Further, there are also currently no broadly accepted ways of citing specific moments within games, such as one might cite specific pages within a book or a specific time-stamped moment within a film.

Information retrieval capabilities of modern search engines include a comprehension of the semantic content of traditional media (e.g., the text of a book or

an article, the identity of an object or person in a photo, etc.). The latter is a more recent innovation, and current search systems cannot yet do the same for videogames. Therefore, they are limited to providing users with secondary sources of information (e.g., linking to an article about the game, a player's guide, or to a recorded play session, aka a "Let's Play" video, etc.). Recent work has demonstrated that it is possible to ingest and index a subset of the possible moments within a game (Zhang et al., 2018), which can then be referenced in the way one might reference pages in a book or a timestamped section of a film. In addition to *searching across* many games, which is imperfectly supported by keyword-based search systems in app and game stores, this dissertation is concerned with enabling *searching within* individual games.

Specific projects described in this dissertation include the initial efforts to identify user needs for videogame moment search, development of two novel search tools to meet those needs, evaluation of the effectiveness of those tools, and a growing annotated collection of student games which can support further research in this area. For the first of these projects, I interviewed scholars and other experts on games (developers, streamers, speedrunners, educators, etc.). My subsequent projects creating videogame search tools to enable and accelerate game search were guided in part by these expressed needs. The first of these tools is flexibly designed to work on videos of gameplay from a variety of sources. In the course of developing the second of these tools, I also began to interact with and contribute to a growing collection of student work, which creates additional research and educational opportunities. One of those

opportunities was the development of the second search tool, which builds on this collection of student games, using structural features of a common text-based game format (Klimas, 2009). Overall this work uses computational techniques grounded in a user-centered design process, and my hope is that it will enable both experts and students to do their work more efficiently and empower them to explore new approaches. In terms of intellectual contribution by field, I intend for these efforts to contribute to games studies by facilitating close and distant readings, to information retrieval by introducing and advancing a new domain for the field, and to game education with the development of visualizations and tools for student work that add insight and inspiration.

In addition to growing as an industry, games (and other forms of playable media) are growing as a field of study. In his introductory textbook, Frans Mäyrä defines game studies as "a multidisciplinary field of study and learning with games and related phenomena as its subject matter" (2008, p. 6). A recent review indicates that game studies includes at least seven distinct research communities spread across dozens of research venues and a large variety of distinct themes (Melcer et al., 2015). The limitations of modern commercial search systems (e.g., Google) are a challenge to the continuation of this growth because, while they are capable of covering many of the byproducts of interactive media (e.g., textual and video game reviews, walkthroughs, criticism), they are not well equipped for ingesting the primary sources (i.e., the games and other playable artifacts themselves). The commercial aspect of many games is part

of the challenge here, where sometimes derivative works (like "Let's Play" videos) are seen as free advertising, and sometimes as copyright infringement (Lastowka, 2013), but the availability of the games themselves is restricted due to fear of piracy. Scholars and other domain experts are frequently interested in working with these primary sources, and in developing insights that require looking both within individual games and making comparisons across a large variety of games. While this work is possible, it can be very labor intensive in ways that limit the scope of current investigations. To give one example, I spoke with a scholar interested in the history of games who framed the query "Are there any platforming games from before 1980 that combine side-scrolling mechanics and a blue background, on any platform?" (B. R. Anderson & Smith, 2019). Answering this question while relying on current search systems would require the effort of a time consuming exhaustive search conducted by the individual scholar, but it is not difficult to imagine a future search system that would enable this (and similar questions) to be answered with minimal effort. For this specific case a large enough corpus of games and a robust list of filters might suffice. Technical approaches that can accomplish this and meet other scholarly needs can be derived in part from existing information retrieval methods (e.g., a vector space model (Salton et al., 1975), domain specific taxonomies (Vallet et al., 2005)), but will need to be adapted to meet the unique challenges of games.

The next chapter of this dissertation will cover relevant background, including a brief summary of the way that current search systems work and work that is

specifically relevant to game search. The following chapters will describe the design and contributions of a game-search requirements analysis, several experiments in developing videogame search systems based in part on those requirements, and the development of two videogame search systems that build on these experiments. Between describing these systems, I will recount the development and content of a collection of student games, as a resource for further such work. The ultimate intent is to provide an empirically grounded articulation of the user needs of scholars, educators, and other experts for videogame moment retrieval, to create novel videogame search systems, and to provide guidance for the future development of these systems. The final chapter of the dissertation will discuss the broader intellectual contributions of this work, specifically in the context of games scholarship, information retrieval, and game education, and offer some directions for future research.

It is worth taking a moment here to clarify what is meant by moments in this work, because of the importance of what is being left out. Specifically, there is an actual historical and social context to moments that is not captured by the technical approaches I am exploring here. For example in the game of Go between Lee and Google's AlphaGo system a specific moment of play, the 37th move of the second game, was identified as reflecting a kind of unprecedented non-human play by expert commentators (Metz, 2016; Moyer, 2016). Similarly, a speedrun that takes place in the context of a contest between live players is very different from a pre-recorded run, even if the runs are technically identical. The first time a boss is defeated in a popular MMO,

a decisive moment in an esports tournament, or a moment that is being shared and given context by a videogame streamer all include elements that are not being captured by the technical systems described in this dissertation. A moment of play that changes how a game is perceived, that reflects a high level of individual or group achievement, or that in some other way is contextually important, is significant in ways that a simple record of a playthrough does not capture. A "moment" as I will be using the term does not take into account any novel historical or social context, but only the moment as understood by the game's technical systems. That said, while a record of game states does not capture this external context, it can enrich our understanding of these significant moments. To some extent this is a matter of degree. For example, speedrunning streamers often incorporate elements like physiological data (e.g. heart rate), or current and personal best times. This information then becomes part of their broadcast archives, attached in some way to those moments of play. A system that understands gameplay moments as I am describing them could be the basis for one that incorporates this, or a selection of other elements to, provide a broader context of information.

Contributions

Videogame Moments in STEAM Education

STEAM is an acronym for the fields of Science Technology Engineering Art and Math (Feldman, 2015; Yakman, 2008), which has been presented as an alternative to the more common STEM acronym which was coined to promote national investment

in these fields (Bybee, 2010). Search tools designed explicitly for interactive media (e.g. games, apps) can meet needs identified by STEAM educators (e.g., developing classroom video demonstrations, and providing students with interactive examples) that current search systems cannot. For students, such tools could support their own design work by exposing them to a greater variety of related works for contrast and comparison, and enabling advanced analyses such as exploring analogical relationships between points in that design space. Developing systems for searching within and across works of interactive media expands the field of information retrieval into a new domain, requiring novel representations that capture not only the content of that media, but also aspects of its structure. By structure I mean the elements that define the space of possible content (e.g., the links between passages an interactive story, the variables capturing a system's state, etc.) that distinguishes interactive and traditional media. This work has transformative potential because of the growth of interactive media (in both popularity and quantity), the influence this media has on those who engage with it, the inadequacies of current search systems for these domains, and the potential of search technology to meet the expressed needs of STEAM educators and to scaffold advanced student work.

I believe that the primary potential social benefit of the projects described in this dissertation will be their contribution to STEAM pedagogy. Specifically, this will primarily provide benefit to programs that teach interactive media creation or analysis. The undergraduate games corpus (Chapter 5) and projects derived from it will aid

students in seeing and responding to work by a broad and diverse array of classmates, rather than being limited to a much narrower canon of examples, or what is available through the bottleneck of commercial release. The search tools will enable students to begin exploring advanced analysis work, facilitating their distanced readings at the scale of representative samples or even the corpus as a whole. Exposure to a variety of peer work could provide a replacement for lost peer connections, especially during the remote teaching response to the COVID-19 pandemic. More exposure to peer work might also increase personal relevance and encourage engagement, particularly providing examples that are more relevant to novice designers than comparisons with commercial releases. These tools and resources will be incorporated into a classroom intervention. Outcomes of this work includes a public student games corpus, advanced interactive media search tools, and the beginning of evidence-based pedagogical guidance for each.

This project is inherently a contribution to STEAM (and not STEM alone) specifically because it exposes students to a technical system designed to enhance the creation and criticism of media artifacts, and it asks them to be critical not just of those artifacts but also of the system itself. This can also be considered a contribution to the media literacy part of MESH education (Wise, 2019) (MESH stands for Media Literacy, Ethics, Sociology, and History). In this way it complements the more rounded approach that characterizes games education programs which combine the

development of interpretive, creative, and technical skills into a more complete and modern form of liberal arts education (Brockman, 1996).

A common failure of imagination in incorporating games into education is not seeing the student only as a potential player and not also as a potential creator or a potential critic. The mistake of observing players engaged with a game and imagining you can capture that engagement in the social context of a classroom has been repeated many times, often leading to the games characterized as "chocolate-covered broccoli" (Bruckman, 1999, p. 76). This contributes to a poor reputation of educational videogames (Van Eck, 2006), with "edutainment" often being a term of derision (Egenfeldt-Nielsen, 2011; Salter, 2016). Many mediocre learning games have been released (Shuler, 2012), disparaged by critics as based on a flawed model of education as simple content delivery (Gee, 2013). the flexibility of games means that they have the potential to implement almost any given theory of education (Plass et al., 2015), but doing so effectively requires deep knowledge of both the subject being taught and of game design.

However, there are many ways to engage with a game other than as a player. A few examples from outside of an academic context include the parasocial interaction of participating in a game stream, the labor of making and sharing guides and walkthroughs, the creation of game mods and fan games, and the critical analysis that is not uncommon on game forums and discussion groups. To focus on just one aspect that can be brought to the classroom, as the tools to do so have become more accessible,

educators can incorporate game-making as a way of knowing into more subjects. One of the best ways to really learn something is to teach it, and for this purpose the computer is an effectively obstinate pupil. A student who generates a model that reflects a personal reality has done important expressive work, both demonstrating and building on their own understanding. This creative work may ultimately be a more effective way of achieving engagement and learning than traditional modes of play. The tools and resources I will describe are best equipped to be incorporated into an approach to games education that emphasizes the creative and the critical.

Videogame Moments and Information Retrieval

This work will (a) generate knowledge in the domain of information retrieval, and (b) contribute to the development of a public archive of student work that will support future research.

Information Retrieval: Even small works of interactive media potentially contain vast spaces of interactivity within them. Capturing this potential requires going beyond how text and images are represented for search, and developing effective representations of both *structure* and *content*. Tools for exploration and analysis will depend on these representations for effective navigation across and within interactive media. It is an enormous engineering challenge for commercial search engines to keep track of the trillions of pages on the web, but it is potentially an even harder challenge to account for the space of technically distinct moments that are possible within even a single, simple, videogame. Examining even the single sense of moments in games

considered in this dissertation is advancing information retrieval concepts beyond the quickly enumerable and discrete pages of other, non-interactive, forms of media.

Student Games Corpus: Student game collection efforts will support future work including new kinds of research (e.g. game studies, machine perception, etc.) and pedagogical support for several novel approaches (e.g. corpus level analysis, design trajectory exploration, etc.). The size of this collection and the variety work within could also be used to stress-test the representations and tools described above. The student games corpus can be used to advance machine perception for interactive media, an area of research that has primarily focused on processing static inputs (e.g. fixed images and texts), in the same way that image datasets have advanced the capabilities of computer vision (e.g. ImageNet (Deng et al., 2009)). In the short term, this could help moment search systems judge when two technically distinct moments are similar enough for human perception purposes. In the longer term, this might contribute to automated game playing research, and the development of systems that can crawl games and to harvest their unique and salient moments.

Research Questions

In the projects described in this dissertation, I address the following research questions:

- **RQ1 [task definition]**: Within works of interactive media, what are people (in the categories of scholars, hobbyists, and students) looking for? Which

collections of interactive media do they want to look in? How do they want to find what they are looking for?

- **RQ2 [operationalization]**: How can we create search systems that match or build on these task definitions?
- **RQ3 [opportunity]**: Once we have created such systems, how does it enable further actions that users wish to take? (e.g., teaching game design, improving game development, technical games research, etc.) What guidance can we provide for the development of future systems to more effectively enable these actions?
- **RQ4 [impact]** : What are the opportunities that this creates for games education?

Research Paradigms

The projects described in the subsequent chapters include elements of human-centered research, the development of computational systems, and integration into educational practices (See Figure 1).

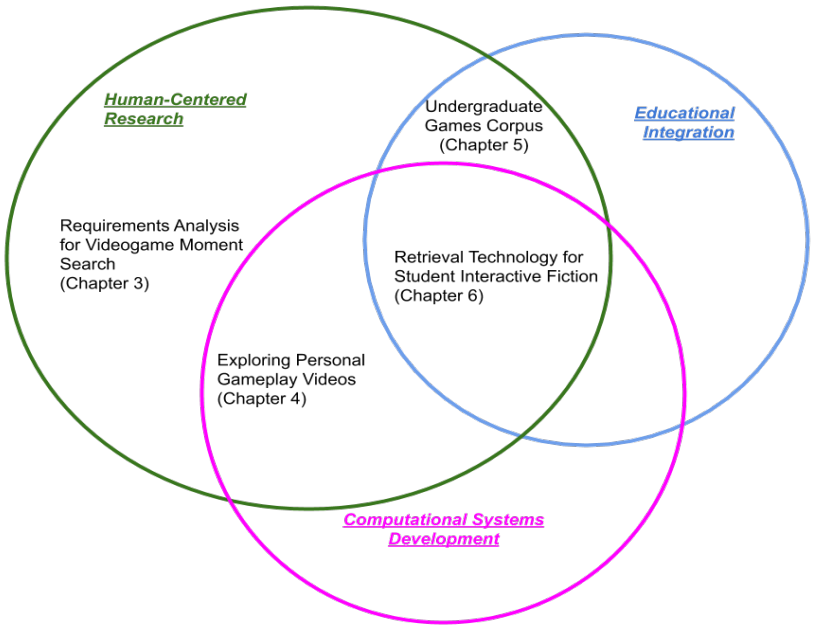

Figure 1. Dissertation projects positioned in relationship to human-centered and computational systems research methods, and educational integration.

The human-centered research described below has included interviews, focus groups, surveys, cognitive walkthroughs and other methods that explore the impact of design on human users. Development of computational systems was guided by insights gained from this research, as well as recent technical breakthroughs. Educational integration here includes choosing domains that are relevant to current students, and the starting to bring these new systems into an ongoing educational practice. The final outcomes of my search system projects ultimately belong to the central intersection of these paradigms.

Chapter 2. Background

In this chapter I review how current search engines work, what is (and what is not) different for search within and across games, and highlight foundational and inspirational prior work related to game search. The search background is intended to be a brief sketch rather than an exhaustive explanation. My intent is not to provide a complete reiteration of the technical details of information retrieval, but to provide sufficient reference for the scope of the systems proposed in later chapters, to introduce relevant information retrieval vocabulary, and to present a few common search system evaluation metrics. The discussion of previous games search work in this section illustrates scholarly needs for games related research tools and establishes some related technologies and ideas. This discussion of specialized search systems and the needs of game scholars provides context and motivation for my own work.

How Search Engines Work

Information retrieval has been defined as "finding material (usually documents) of an unstructured nature (usually text) that satisfies an information need from within large collections (usually stored on computers)" (Manning et al., 2008, p. 1). Before considering how information retrieval applies to other media, it will be useful to review concepts and vocabulary from the traditional domain of text search. The entire collection of documents that the system is going to be able to search is sometimes also called the corpus. In a specialized search system this might be the articles in an

encyclopedia,[1] a collection of scientific papers,[2] or the works of a specific community of authors.[3] For a web search engine, this might be the content of a large subset of all web pages, which was on the order of billions to trillions of words a decade ago (Manning et al., 2008, p. 3) and continues explosive growth with the proliferation of social media. For any of these systems, the user conducts a search by providing a text query. In this context it is important to distinguish the user's information need, which is what they want to find, from the query they provide, which is the text that they actually type into the system (imperfectly hinting at the underlying need).

A modern web search engine provides immediate and detailed results, usually in response to a brief text query, often in ways that suggest a strong awareness of the users' context. These increases in speed and relevance of search results are possible because, despite search results exhibiting few surface level visible changes in the past few decades (Hearst, 2009, p. 2), these systems have undergone a great deal of development (Manning et al., 2008, p. 10; Sanderson & Croft, 2012). For example, modern systems can incorporate user logs and other contextual information to distinguish when the same query indicates different informational needs. However, there are some general features that have remained constant. For search to be possible,

[1] https://en.wikipedia.org/wiki/Help:Searching#Under_the_hood

[2] https://search.proquest.com/psycinfo/advanced

[3] http://www.ohnorobot.com/

an engine must somehow have ingested and indexed the corpus of documents that are to be searched. While early search systems focused on explicitly and meticulously curated corpora for a variety of specialized expert-user systems (e.g., collections of scientific papers, legal documents, news articles, etc.), the current canonical information retrieval case for search is the World Wide Web (Hearst, 2009, p. 4). This adds the requirement of crawling a constantly growing corpus (Hubert, 2013), many elements of which are often incompletely or incorrectly formed, sometimes even in ways that are actively adversarial to this venture (c.f. search engine optimization).

The analogy of crawling and indexing a growing web of documents to the process of ingesting interactive artifacts suggests some useful technical approaches, but more significantly it implies that the project of cataloging spaces defined by interactive media for search purposes is not infeasible in scope. Not only are there systems that effectively ingest the web (at the scale of trillions of distinct documents), they also manage to generally keep up with a web that is always changing. Given this, we can imagine a system that is capable of automatically exploring and indexing a growing corpus of modern games, and several of the necessary components of such a system have already been demonstrated by others in my research group (Zhang et al., 2018).

The simplest approach that a search system could follow is checking each document in the corpus for an exact match for each of the query terms. However, this can very quickly become infeasible for anything but the smallest corpus. To make search more efficient (e.g., to find documents relevant to a query in a reasonable

amount of time), each item in the corpus can be summarized in a compact data structure, in a process known as indexing (Manning et al., 2008, p. 6). For example, a traditional inverted index contains a record for each word in the corpus, and each of these records contain a list of the documents in the corpus that the term appears in. In this case, the system attempts to return the most relevant documents by finding the entry in the index that matches the user's query, and returning the documents listed in that record. This system could handle Boolean modifiers (e.g., AND, NOT) by, for example, finding the intersection between the results returned for two queries. There is some additional complexity behind even this simplified example, because most systems would need to be able to deal with words that contain special characters, very common words, words that should be combined because they are slight variations in spelling, etc. Another level of complexity that is omitted in the previous example, by counting only presence or absence of words, is the importance of each word in the document it appears in. This can be approximated with some understanding of the structure of the document (e.g., if the term appears in the title or headings), and by counting the frequency of the term within the document and the corpus as a whole.

Vector Space Model for Information Retrieval

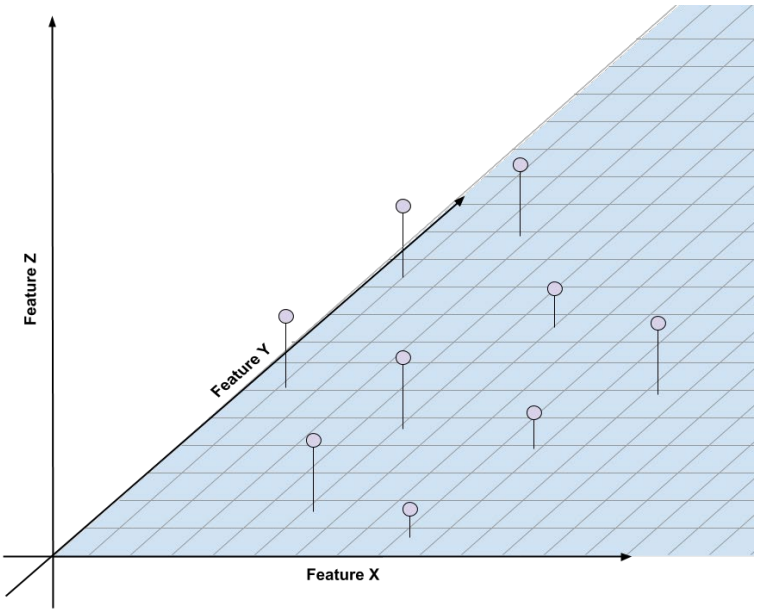

Figure 2. Documents represented as vectors, which can be thought of as points in an n-dimensional space.

If we flip the representation in the preceding example around, looking at vectors for each document (with a dimension for each word) rather than for each word (with a dimension for each document), we can also convert the user's query into an appropriate vector (See Figure 2), which allows us to find relevant matches by proximity in an abstract search space (See Figure 3). Once we have done this, we can start to consider the usefulness of more abstract dimensions and weightings, such as those produced by machine learning, or refined by user provided relevance feedback. There are also search

systems that allow for other kinds of search query input (e.g., searching for images with an image query), that function in essentially the same way. Another way to think about this is that all documents available to the search engine are represented as vectors, and the distance between these vectors in an abstract search space approximates the semantic similarity of the documents. Given a point in this search space defined by a query (whether that query is text, an image, or an even richer document), the goal is to find the points that exist in the set of all documents indexed by the search engine that are closest (most similar) to that point. This abstract description could apply to search in many other domains, even those with no obvious sense of discrete search terms.

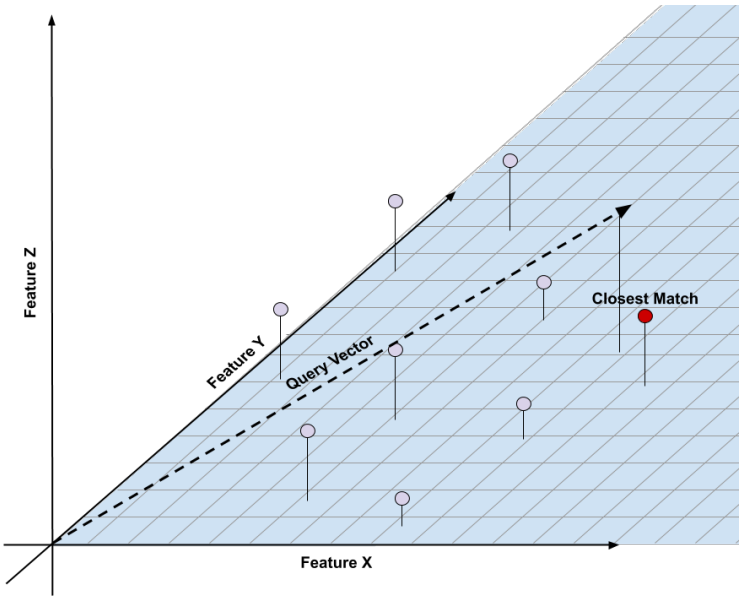

Figure 3. Query vector based retrieval.

There are several ways to define the distance between vectors in a space like this. A common approach is cosine similarity, which looks at the difference between

the angles but ignores differences in magnitude. This approach is one way of dealing with challenges created by the shape of a space defined by a large number of dimensions, which multiplies the amount of empty space between possible points of interest, by ignoring a kind of distance that is believed to be less meaningful. For example, we hope that a short query phrase identifies the topic of the documents we want to find (regardless of their word count); we don't mean to literally search for results that are themselves short phrases. A user would not want their results to be identical to the query, in this way we can return relevant information in the results that is longer than the search term itself so long as it covers topics with a similar relative frequency.

I follow a vector-space retrieval approach in the systems described in later chapters of this dissertation, because vector space models have already proved useful for search across a variety of domains (e.g., for text (Manning et al., 2008; Sanderson & Croft, 2012), for images (Rui et al., 1997; Thomee & Lew, 2012), video (Shih-Fu Chang et al., 1998; Zheng et al., 2006), for audio (Barrington et al., 2007; Typke et al., 2005), and for computer program code (Sachdev et al., 2018)) and because many games contain content that is already part of these well tested domains. The relatively advanced development of text search motivates a selection of text heavy games for initial games search experiments, which will be discussed in more detail in Chapter 6 below. All of the aspects of a videogame moment contain information that can be defined as vectors in a high-dimensional space. In fact, each of these aspects (text,

visuals, state, etc.) is individually in its own separate high-dimensional space. Fortunately, there are several techniques that can be used to meaningfully visualize even the very large and complex space implied by combining all of these elements into a single moment.

Visualizing High-Dimensional Information

In several of my projects to visualize the relationships between videogame moments, I have been using the t-SNE (t-Distributed Stochastic Neighbor Embedding; Maaten & Hinton, 2008) and UMAP (Uniform Manifold Approximation and Projection (McInnes et al., 2018)) techniques, which can create useful two or three dimensional representations of high dimensional data. These are popular data visualization techniques useful across a variety of domains (Wattenberg et al., 2016), including the visualization of embedding vectors developed for image and natural language processing (Chan et al., 2018). The t-SNE technique has also been applied to videogames, both the whole game level and at the level of individual moments (Zhang & Smith, 2019). Compared to the traditional PCA (Principal Component Analysis) method for analyzing high-dimensional data, the recent t-SNE and UMAP algorithms have several advantages. These include the ability to map nonlinear relationships, to avoid the problem of points crowding together (aka “the crowding problem” or “the curse of dimensionality”), and the ability to create a mapping that reveals local structure. PCA’s advantages over t-SNE are that it is a more intuitively simple algorithm, and (more importantly) that creating a t-SNE representation does not also

result in a function that can be used to reduce the dimensionality of new data points. This means that while you can use t-SNE to create a visual representation of existing data you cannot map new data points into the space that it has defined. It is also non-deterministic, meaning that even without changing the data applying t-SNE more than once will result in different visualizations, although the new mapping should still reflect the same underlying relationships. UMAP is a newer visualization technique that was designed to specifically remedy some of the limitations of t-SNE, and includes the ability to add new points to an existing space. Visualizations like these are useful, and often necessary, first approaches to understanding data (Yanai & Lercher, 2020). Further insights can be gleaned by supplementing them with descriptive and inferential quantitative analyses.

Low dimensional projections (2D/3D) of a complex high dimensional space like videogame moments can facilitate other kinds of searching across and within. They enable forms of spatial interaction that are more akin to using a map than to using a typical search engine, for example seeing moments that are near to others or exploring routes between moments, rather than providing a flat list of results. Different projections, or even choosing a different perspective on the same projection, can provide different insights about the relationships between moments. These projections should be thought of not as abstract illustrations of the space, but as functional search tools in their own right.

Quantitative Evaluation of Information Retrieval Systems

The effectiveness of an information retrieval system is traditionally assessed by quantifying its *precision* and *recall*, which are technical terms in this context (Manning et al., 2008). Precision is defined as the fraction of the results returned that were relevant to the user's information need, out of all of the results that the system returns (perhaps just on the first page of results). Recall is defined as the fraction of relevant documents that were returned, out of all of the relevant documents that exist in the system's collection. For either of these measures to be useful there needs to be some evaluation of relevance, which can be a time consuming and expensive process. This is a more significant limitation for measures of recall, which require this relevance evaluation be done for each and every document in the retrieval system's collection. Relevance judgments need to be considered in the context of a specific set of documents and queries, but the precision and recall measures that they enable can be discussed more generally.

Two common scoring metrics for precision rely on the existence of canonical best matches for a set of queries. These might be results explicitly curated by the search engine's human designers, or the results of some existing comparable system that the designers are attempting to replicate. In a way, this is the most obvious approach to improving the quality of a search system, by directly asking the target users (or your best approximation of them) to provide information about their preferences. One metric is mean reciprocal rank (MRR), where the score for a query is the reciprocal of the rank

of the best correct response (Manning et al., 2008, Chapter 12), (e.g., if the best response is the second search result the score is ½, if it is the third result the score is ⅓, etc.). This scoring method ignores the quality of all results beyond the first relevant match for each query. An approach without this shortcoming is normalized Discounted Cumulative Gain (nDCG), which assumes that several useful results should be near the top of the rankings provided by the search system (Manning et al., 2008, Chapter 12). This generates a score that is sensitive to the position a document appears in, resulting in lower scores when a document is lower in the list of results than expected, based on its relevance. A normalized score makes it possible to compare the quality of results between different queries, because the length of the results list for each query will vary. An nDCG metric was used internally at Microsoft during the development of their Bing search engine (Shum, 2012), and nDCG is used in this dissertation in the evaluation of both novel videogame moment search systems (Chapters 4, and 6).

Quantitative evaluations can aid in the development and improvement of search systems, but there are some important caveats when considering how to best use them. For one, it is possible to optimize too well for any single quantitative metric. This would result in a search system that exactly and only reproduces examples that, by definition, were already available, since they were the exact example set of queries and best results. This means that a sample of novel results need to be evaluated for relevance to produce meaningful scores, and that this evaluation process cannot be entirely automated (because if it could, that evaluation would better be incorporated into the

search system itself). Ultimately, the evaluation of a search system requires observation of real users. However, with some evaluative labor, high scores on metrics can give you a sense of whether a system is approaching generalizability, or the ability to provide novel relevant results, which is frequently the most desirable property of a search tool. Having established some context for how search systems work and are evaluated, I next turn to inspirational prior work related to game search.

Videogame Search

Academic interest in videogame search, and in videogame moment search, is reflected in work that has taken place in labs around the world. Many of these efforts are oriented toward providing resources and tools for archivists and other games scholars. For example, several projects in Jin Ha Lee's Game Research Group[4] at the University of Washington School of Information are directly related to information retrieval for games, such as standardizing and evaluating metadata about games (Lee, Clarke, & Perti, 2015), the development of a controlled vocabulary for videogame visual styles (Donovan et al., 2013) and for videogame plots (Welhouse, Lee, & Bancroft, 2015). On the other side of the world at the University of Aalborg, recent qualitative analyses of videogame search requests on a social media site (Jørgensen & Bogers, 2020) gave insights about what potential users might want from a game search

[4] https://gamer.ischool.uw.edu/

system and provided an example of another approach to videogame search. My direct inspiration, however, came from research projects that took place at UC Santa Cruz.

The Game Metadata and Citation Project (GAMECIP)

My research on videogame moment search builds on foundational efforts of the The Game Metadata and Citation Project (GAMECIP), a multi-year IMLS-funded joint initiative between the UCSC Library, UCSC Computer Science, and Stanford University Library. This was a project aimed at improving institutional practices around games, including efforts aimed at improving the ways that games are cited for games and making the citation of game moments possible (Kaltman et al., 2014). Artifacts that were developed as a part of this project enable a certain kind of search across games. One such project, GameSpace created an interactive 3D star map of games, allowing the user to fly through and observe constellations of related games (Ryan et al., 2017). Another, GameSage takes natural language game descriptions as queries, and returns a list of likely games that are often reasonable guesses (Ryan et al., 2016). These are both useful and interesting tools, but it should be noted that they both work by leveraging the textual content of Wikipedia pages about games. This means that these prior systems are using information about games available in secondary sources. In contrast, the projects described in this dissertation operate directly on the content of the games themselves.

Game Moments as Primary Sources

Part of GAMECIP also involved the development of a format for citing moments within games, within the Game Interactive Software Scholarship Toolkit (GISST) tools (Kaltman et al., 2017). These tools enabled the citation of game software, game performance recordings, and specific moments of emulated play, all supported with a stable citation database. Going even further into the contents of game moments themselves, Zhang et al. (2018) crawled and indexed them in the same way that a search engine crawls and indexes pages on the web. For videogame moments, as for web crawling, this sampling approach is a useful way to characterize spaces that are too vast to explore exhaustively. These projects were critical steps forward, acting on actual game content instead of relying on secondary sources, and the information retrieval work described in this dissertation was inspired by and builds on them both. One of the primary benefits of taking an information retrieval approach to increase the searchability of game moments is that it will make it easier for scholars (or other potential users) to find and cite gameplay moments that they cannot reach through their own play. This is helpful for several reasons. Reaching many game moments can be a time-consuming and uncertain process, and can depend on both luck and ability. Some moments are rare, are only reachable by highly skilled players, or even only by highly skilled teams of players. Search also enables more exploratory investigations, such as those that did not begin with a specific moment in mind, or that require the discovery and comparison of related moments.

Chapter 3. Requirements Analysis for Videogame Moment Search

Project Summary

One way to approach the creation of game search tools begins with an effort to develop an understanding of the needs of those who might use them. The first step in this process is to identify possible and probable users. I began this work with the expectation that experts and students of various kinds would be the primary stakeholders for game search projects. To understand the needs of the former, I engaged in conversations with experts who were able to identify how these tools might fit into their existing work, and provided a variety of specific and grounded suggestions. The experts that I interviewed included a variety of scholars, streamers, and developers. While it is not a substitute for talking with students directly about their concerns, many of the experts I spoke with were also educators, and from that perspective they were able to speak to an important subset of student needs. These interviews built on pre-existing work that had already demonstrated the technical foundations of a specific kind of game moment search (Zhang et al., 2018). My intent was to validate that work's assumptions, to increase the diversity of potential user perspectives considered, and to engage in a preliminary exploration of unanticipated user interests.

In this chapter I discuss a research project that involved synthesizing recommendations for the development of game moment search from expert interviews, and consider how those recommendations might apply to student work. The

intersection of student relevance, technical feasibility, and the expert user needs identified in these conversations guided exploration in subsequent game search projects, which are described in the later chapters of this dissertation. The current chapter expands on findings originally published in the proceedings of the 2019 Foundations of Digital Games conference (B. R. Anderson & Smith, 2019). Detailed notes about the conversations that this work is based on are also provided (See Appendix – Videogame Search Requirements Analysis Interviews). The co-author listed in this publication directed and supervised the research which forms the basis for this dissertation chapter.

Research Questions

Enabling rich search capabilities for videogame moments has implications for many different types of users. Speedrunners, streamers, or general consumers might bookmark and share moments in a leisure or hobby context. Individuals in these user profiles are usually interested in well-known commercial games that have been played by many others. Scholars, educators, translators, developers, and critics need to reference specific moments as part of their professional work. This might involve relating moments in less well-known games, or linking other games to the specific game they are developing or critiquing. In an industrial context, the existence and reachability of certain events is important to app store operators and quality assurance testers. Whether certain moments are possible in a soon-to-be-released game may impact the game's release schedule and its later commercial success. The study reported

here explored the needs of individuals in a subset of this broad spectrum of possible user profiles.

Zhang et al. (2018) examined the possible methods and applications of videogame moment search from a technical capability perspective. The expert interviews described here complemented this approach, seeking to identify the stakeholders of videogame moment search and identify their needs with a user-centered design methodology. The way that a scholar, teacher, or hobbyist might want to search for a moment varies depending on the context and purpose of the search. For example, an educator seeking a moment to best represent a certain game mechanic during a lecture may have different interests than the expert speedrunner using search to review records of their own past play. I interviewed a range of users who could benefit from a videogame moment information retrieval system: speedrunners, streamers, educators, developers, and scholars. For each of these user profiles, I addressed the following research questions (RQs):

RQ1 What kinds of game moments do various types of users seek?

RQ2 How do they want to identify those moments in a search?

RQ3 What will they do with retrieved moments?

Following the expert interviews, I synthesized answers to these questions into specific recommendations for the next technical steps in interactive media information retrieval research, highlighting a role for systems that work across many game

platforms, search with textual queries, ingestion of bulk video data, and applications of retrieval to personal information management.

Related Work

The interactivity of videogames, with their branching paths and dynamic content, presents a new challenge for information retrieval. Conversations about games often need to include references to moments to answer fundamental questions: What can players do in the game, how do they do it, and what effects do those actions have on the game world? The problem of identifying specific moments has been long solved for more traditional media (e.g. a page number in a book, a timestamp for a film), but games present several unique obstacles. For example, because games can be nonlinear, two playthroughs of the same game might see the same moment at wildly different times. Because games are dynamic, and players can make different choices or have different levels of skill, it is possible for two playthroughs of the same game to contain very different collections of moments with very little overlap. Kaltman et al. (2017) recently introduced a system for meeting these challenges, primarily for scholars, which defines a citation format that can identify a unique game moment unambiguously. However, this system only helps users communicate about unique moments in a game that they are able to reach by their own play. Search and retrieval systems could help scholars find many other moments they might wish to refer to, replicating the successes of rich-media search in other domains (Thomee & Lew, 2012). The study reported in this chapter was the first to frame the question of how best

to use information retrieval (IR) for videogame moments as a question for human-computer interaction (HCI) research methods. It presented the results of an interview study, partially anchored on a demonstration of current techniques for content-based retrieval of videogame moments, and suggested several new applications both within and beyond game studies.

This chapter contextualizes the investigation of human-computer interaction concerns in information retrieval, the newly emerging scholarly study of moments in games, and existing strategies for search used with rich (non-textual) media.

HCI of IR

There are many possible interfaces to search engines that vary based on their context of use and the type of media they are being used to retrieve. HCI research in IR seeks to identify the distinct patterns of use and the interface designs that best support them. For example, Andre et al. (2009) analyzed differences in use between image and text search engines and recommended new interfaces for searching non-textual media. Work like this is important for setting the direction of new technical research on the basis of user needs discovered through empirical research methods. Within the human-computer interaction and information retrieval community, recent work by Ince et al. (2018) used a semi-structured interview process to understand the search practices of doctoral students in scholarly workflows. Rather than investigating how subjects engaged with a specific retrieval system, they sought to understand patterns of use

across different tools. The work described here adopts the same research methodology and similar investigative goals. Where Ince recruits subjects all within a specific early-career academic role, I recruited subjects spread across several different contexts of use for retrieval technology in an effort to represent the breadth of stakeholders. In long-form media outside of games, recent work by Bogers et al. (2018) considered similarities and differences in how users seek books and movies. The attributes by which these media artifacts are recalled differ from how users seek webpages or research articles. In particular, users often search for the larger book or movie by referring to specific scenes, characters, or even precise lines of dialog within it. This is evidence for the value of ingesting and indexing the content of moments for a variety of media, which has positive implications for videogame moment search.

Indexing Moments in Games

Current games researchers analyze games using a variety of methods, ranging from analysis of pre-recorded video and screenshots to careful investigative play using specially prepared game platform emulators. Even at the level of referring to whole games, the lack of standardization around vocabulary and reference formats has meant scholars work with conventions developed in hobbyist communities, or need to do the work of developing their own reference rules (Altice, 2015). In some scholarly contexts it may not be enough to provide a reference of the style "*Super Mario World* (1990)," in the face of the many versions of any given game (e.g. international releases, minor revisions, complete remasters, etc.), including some variations that were created or

modified by fans (e.g. mods, total conversions, etc.). Kaltman et al. (2017) led an effort to make it possible to cite specific moments in specific games. In particular, those moments that are captured by the complete state (e.g., of system memory and processor register contents) of a game platform emulator. This same notion of moments is the focus of the work I was building on, and the source of the examples provided, during the course of these interviews. In contrast to laboriously described identification of a game moment by text (such as instructions for reproducing it in manual gameplay), these emulator snapshots are compact and precisely replicable.

Zhang et al. (2018) introduced the problem of crawling, indexing, and retrieving moments in interactive media as the key processes in a new kind of search engine. Zhan et al. (2019) identify several different strategies for extracting a diverse collection of (emulator snapshot) moments from a given videogame on platforms ranging from the 8-bit Atari 2600 to the 32-bit Nintendo 64. Some of these strategies involve replaying expert gameplay traces while others involve running algorithms to automatically explore the space of interaction.

In the interviews reported below, I spoke with some of the individuals who produce these kinds of expert interaction traces (speedrunners). I also interviewed individuals in a game developer role who might gain insight from moments uncovered via automated exploration that they (or testers on their team) would have been unlikely to find themselves.

Rich-media Search Engines

A search engine for moments in videogames is situated in the context of previous work on search for rich (i.e., non-textual) media. One component of a game moment is the visual image it presents to the player, and there are several paradigms for visual image search. Users of existing tools like Google Image Search[5] or TinEye[6] search seeking image results. For Google Image Search, this might mean using a descriptive text query, and often involves learning the particular vocabulary necessary to produce the desired results. For TinEye the query is always an image itself, and the sought after result might be the original web-page source and context of that specific image, variations on the image, or a more canonical or higher resolution example. The Shazam[7] mobile app offers a sound-based search engine where users query the system with a short snippet of recorded audio and expect to see results that list which popular music track they have identified. In part of my interview protocol, I remind subjects of these systems to uncover parallels between audiovisual search and interactive media moment search (which might use rich media artifacts as query or result items).

Another approach to rich-media search considers searching for objects that are not themselves images or audio. In the ProductNet visual search system (Bell & Bala, 2015), images of home furnishings are used to retrieve instances of products that are

[5] https://images.google.com
[6] https://www.tineye.com
[7] https://www.shazam.com

available for purchase. These results are actual products, not merely visually similar images. Many of these rich-media search systems work by training a neural network (or other function representation) that can embed images of similar objects into nearby locations in a latent space while mapping dissimilar objects to distant locations. Zhan & Smith (2018) introduced a similar technique to represent moments in games by relating screenshot images to savestate memory configurations. Compared to product search, which retrieves possible results from a finite catalog, the space of possible moments that might be returned by a game moment retrieval system can be astronomically larger. My interviews suggest it may be useful to retrieve moments that no human player has previously seen or reached via direct (either ordinary or tool-assisted) play.

Interview Methods

To understand the needs of users in videogame moment retrieval, I conducted a series of semi-structured interviews. I contacted individuals who use game moments in a variety of different contexts. Although I do not cover every possible user profile or context of use sketched below, I believe the sampling in this initial study on the topic already provides a broad perspective.

Recruitment

The way that a scholar, teacher, or hobbyist might want to search for a moment will vary depending on the context and purpose of the search. For example, an educator

seeking a moment to best represent a certain game mechanic during a lecture may have different interests than the expert speedrunner using search to review records of their own past play. In this project I conducted a series of semi-structured interviews with a range of users who could benefit from a videogame moment information retrieval system: speedrunners (who compete to complete a game as quickly as possible), streamers (who share live videos of their gameplay online), educators (who teach others to create and critique games), developers (who make games), and scholars (who interpret and critique games and their contexts). Potential subjects were contacted directly, either via previously established personal connections, or through publicly available contact information. Out of twelve individuals contacted, ten agreed to be interviewed. User categories were initially generated by a brainstorming process, and further refined by informal discussions with other games researchers, and by directly addressing these roles within the interviews themselves.

Before recruitment, I generated a set of hypothetical user profiles, which was refined and expanded throughout the study. Potential interview subjects needed to have a connection to one or more of these profiles, as judged by the authors. Potential subjects were contacted directly, either via previously established personal connections, through my own networks, or through publicly available contact information. Out of twelve individuals contacted, ten agreed to be interviewed. Of those ten, three came from outside of any academic institution. All subjects were recruited for their specific qualifications relevant to each user profile. I was not able to recruit

subjects for each profile, and these missing perspectives (e.g., game translator) are considered below.

Because each subject's individual background is integral to interpreting their needs, their comments are quoted below with attribution. I introduce each subject in the section in which they are first referenced. Of the many contexts of use, my subjects offer coverage of hobby and professional use of videogame moment search, but we were not able to connect with subjects that could speak to industrial uses (e.g., app store operators). All subjects consented to being recorded and to being quoted with identifying attribution. The complete list of subjects and their association with user profiles is shown in Table 1.

User Profile	Profile Description	Interview Subjects
Speedrunners	Competing to complete games as fast as possible	BH, SAM, AR
Streamers / YouTubers	Broadcasting gameplay video with commentary	NJ, SCM, SAM
Educators	Teaching about the design, history, or interpretation of games	NA, HL, ES
Developers / Designers	Creating new videogames or tools to support videogame creation	BA, TL, ES
Scholars	Doing critical or interpretive work with videogames	NA, HL

Table 1. User profiles and their representative interview subjects.

[Q] Subject identification: Identify yourself in one or more of the user profiles in the list provided. Describe how your work or other activity relates to videogames.

[Q] Current search methods: How do you currently search for and within games now, or instruct others to?

[P] Current search methods: Review examples of Google text search, Google Image search, Steam text search and browsing, YouTube text search and browsing, Reverse image search (e.g. TinEye), and Maps search.

[Q] Speculation: How do you think a hypothetical videogame moment search engine would work? What are the inputs and outputs? What should it do for you?

[P] Demonstration: View a demonstration of Zhang's system review an explanation of existing technology.

[Q] Capabilities: What could / couldn't you do with this system or one like it?

[Q] Queries and Results: What game platforms needs to be indexed for system to be useful to you? What needs to be extractable from search results? How would you share results?

[D] Open-ended discussion

Figure 4. Semi-structured interview protocol from videogame moments requirements analysis interviews. Q marks questions for subjects, P priming demonstrations intended to guide discussion, and D subject-specific discussion. Items are paraphrased here for brevity. Reproduced from Anderson & Smith, 2019.

Interview Protocol

Interviews were conducted following a specific protocol (See Figure 4) intended to avoid biasing participants while also eliciting complex responses. Interviews were approximately 40 minutes in length, and were conducted face-to-face

or via video-conferencing software. In one case a subject requested an interview to be conducted via text chat. All other interviews were video or audio recorded. Subjects were encouraged to self-identify among several possible user profiles[8] or to describe their work with games in their own terms. Subjects were asked how they currently search for or within games, and how they might wish to do so. After their initial responses they were reminded about several forms of search beyond free text (e.g., reverse image search, map based search, video recommendation algorithms, user generated tags, etc.) and an existing videogame moment search system was demonstrated (Zhang et al., 2018). Further discussion after these anchoring demonstrations elicited additional suggestions. Combining and categorizing responses from these conversations, I generated a set of specific recommendations for future game search systems (see Table 1).

Anchoring Demonstration

An existing prototype videogame moment search engine was demonstrated about halfway through each interview, followed by an explanation of the underlying technology. This demonstration and explanation served to prompt additional suggestions and ideas, now grounded with a working example. The system is described in detail elsewhere (Zhang et al., 2018), but in brief it allows searching for in-game

[8] User Profiles: Speedrunner, Streamer, Educator, Scholar, Developer, or Designer.

moments using screenshots as input, currently for a select few games for the Super Nintendo Entertainment System (released in the early 1990s).

About halfway through the interview, and after I had already elicited several suggestions, I asked my interview subjects to watch a demonstration video[9] of a prototype game search system. In this video a screenshot from the game *Super Mario World* (1990) being dragged and dropped onto a web-based search engine interface. In response, the system supplies an initial list of potentially relevant search results. Each result item identified what game the moment was from and showed a screenshot thumbnail. Next, the video demonstrates tagging some of the results as examples of relevant and irrelevant results. The system's relevance feedback feature is applied to refine the next batch of results under these suggestions. All moments seen in search results are sampled from a recording of non-speedrun gameplay from a student volunteer (rather than from any specific dataset with which the subject might already be familiar).

Analysis

During the course of the interview, the interviewer made occasional reflective notes. This process of memoing (Glaser, 1992, Chapter 4) continued with a subsequent review of the complete recording. In this process the interviewer looked specifically for patterns that could be attached to our broader user profiles as well as for unique

[9] https://drive.google.com/file/d/1eGkx1mh_Nry1hHP2S4j0p4HVL2pMCzTy

insights from individual subjects. These notes and portions of the interview were presented to the research team, who provided additional observations or suggestions for interpretation. The analysis was conducted with an emphasis on identifying specific future directions for technical research (particularly within the artificial intelligence techniques used when constructing new search engines). This marks contrast with analysis purely focused on finding generalizable patterns within a population. Finally, the audio or video recording was reviewed again to find specific quotes supporting the observed patterns or insights.

Findings

User Profile	Moments Sought (RQ1)	Query Formation (RQ2)	Ultimate Goal (RQ3)
Speedrunners	Individual speedruns, useful glitches and tricks	Screenshots, Short Video, Text (with named tricks and descriptions)	Learn from examples and practice for their own runs
Streamers / Youtubers	Moments that support specific arguments or commentary, moments from personal play	Text (with game specific tags and parameters), Screenshots, Video	Resume play from a moment to explore variations or to skip between game sections
Educators	Conceptual and aesthetic examples, Examples of mechanics	Text (with a controlled vocabulary containing examples of game mechanics), Citations, Screenshots	Resume play from a moment to share experiences of game mechanics with students, or produce classroom video demonstrations
Developers / Designers	Moments linked to player reactions, moments where design tools would be useful, bugs and unexpected behavior	Player Reactions (e.g., excitement, boredom, etc.), Text (with an awareness of moment types, e.g., "tutorial," "boss fight," etc.), Screenshots	Finding "fun" or "pain" points for monetization opportunities, communicating game design concepts within development team, finding additional examples of a known bug
Scholars	Rare or unique moments, examples of superior/aesthetic play, historically important moments	Text (with descriptions of known events, or named mechanics), Complex filtering (e.g., "before 1985"), Citations, Code	Validate the possibility of a moment happening in a given game, resume play from a moment for exploration, discovering what game elements specific code enables

Table 2. Summary of videogame moment search requirements analysis findings organized by user profile. **RQ1**: What kinds of game moments do various types of users seek? **RQ2:** How do they want to identify those moments in a search? **RQ3**: What will they do with retrieved moments? Reproduced from Anderson & Smith, 2019.

To report findings from interviews with subjects in many overlapping roles, I have organized the results by user profiles rather than by subject or research question. The background of specific subjects is included in the first section for which they are introduced. A summary of findings can be seen in Table 2.

Speedrunners

This first user profile section examines one category of stakeholder who would use videogame moment search in a hobby context. Speedrunners are an online community of players who compete to complete games as fast as possible. Tracing its roots to the release of id's *Doom* in 1993, this subculture has seen substantial growth, with an increase in dedicated websites (e.g. NESVideos in 2003, which became TASVideos[10] as its scope extended beyond Nintendo games, and the launch of Speedrun.com in 2014). The charity speedrunning convention Games Done Quick (GDQ[11]) has also been growing in attendance, viewership, and fund-raising since its inception in 2010. Speedrunning combines aspects of human performance and technical achievement. Some speedrunners prefer to keep things as close to unaided as possible, while others are interested in using any mode of play (including the use of automation) to find the fastest times possible. To understand the needs of this community we spoke with individuals who participate in and contribute to different aspects of the hobby.

Speeedrunner Interview Subjects

Alexis "Protomagicalgirl"' Ross is a speedrunner, a member of several online speedrunning communities, and one of the organizers of the GDQ convention. Her

[10] http://tasvideos.org/
[11] https://gamesdonequick.com/

Speedrun.com profile contains over a hundred runs across a large variety of games,[12] and she has dozens more on a site dedicated specifically to the popular *Mega Man* game series.[13] For GDQ she has been a member of the pre-show and interview teams since 2016, and she's served as staff for a small number of other speedrunning events. She prefers, when possible, to do her speedrunning on original hardware.

Bernd “Aran Jaeger / ED” Huber is a speedrunner and a member of the TASVideos community. His contributions there have primarily been related to the game *Super Metroid*, a popular game released for the Super Nintendo platform in 1994. Instead of running the game himself, he looked for exploitable glitches in games that could be used to produce a faster run, which sometimes involves deducing how a game is using the system's memory, and shared his insights on a wiki dedicated to speedrunning the game.[14]

Stella Mazeika is a hobbyist speedrunner. She is a moderator of a speedrunning forum, and streams her own runs on Twitch.[15]

What kinds of videogame moments do speedrunners seek?

The speedrunners I spoke with are interested in sharing individual runs, as well as tips, tricks, and useful glitches. Ross and Mazeika explained that a moment in a

[12] https://www.speedrun.com/protomagicalgirl

[13] http://megamanleaderboards.net/index.php?page=runner&name=protomagicalgirl

[14] https://wiki.supermetroid.run/ED\%27s_stuff

[15] https://www.twitch.tv/leggystarscream/

game, for a speedrunner, might refer to a specific technique or exploit, and that the speedrunning community has given these tricks names. One example given by Ross was of trying to share a “frame perfect” trick for killing the first boss in *Castlevania*, or finding another specific named trick that would be useful in completing a particular run (e.g. “bat crit”" in *Castlevania* or “teleporter room early” in *Paper Mario: The Thousand Year Door*. Ross expressed interest in a system that can handle many types of classic console games (e.g. *Castlevania*), a small number of indie PC games (e.g. *Cave Story*), and arcade games (e.g. *Tetris the Grandmaster*). Huber described a speedrunner sharing on online forums a moment where a runner was stuck, where the game was crashing, or some other rare event was occurring, in order to solicit advice. He also suggested that in the course of searching for new paths through the game, a speedrunner might be preferentially interested in moments that are further away from ordinary play.

How do speedrunners want to identify those moments in a search engine?

For all of the speedrunners I spoke to, the first idea suggested was using a screenshot to find an example of the trick they were interested in reproducing. This suggestion was offered prior to the anchoring demonstration video. One subject also considered using a short gameplay video as search input. The second common suggestion was using the community's text name for the trick or technique, perhaps with a few additional qualifiers (e.g., with the addition of a location or time in the game) if searching for a particular variation. Speedrunners also expressed interest in making

sure they can filter results to those that are relevant to their community specifically. This might mean restricting results to speedrunners, but also the ability to apply additional restrictions (e.g., human-only, tool-assisted, played by a specific person, or played with a specific set of additional restrictions or rules). This last set of restrictions is important because even within speedrunning the meaning of completing a game as fast as possible can vary depending on what is considered necessary for a game to be "complete." For example, the popular speedrunning title *Super Mario World* has 11 categories on Speedrun.com that each define different requirements (e.g., the "0 exit" category that requires reaching the end of the game while passing through none of the typical level exit points).

What will speedrunners do with retrieved moments?

A speedrunner working on improving their own runs by mastering a new technique would want to find an example to follow. This example might be a set of sequential images or a short video clip, ideally paired with a walkthrough of the necessary steps. For some tricks, the exact frame-by-frame sequence of button presses is also important. A save state is even more useful, if a player wishes to practice by skipping directly to the point in the game where the technique is relevant. Some games already have systems in place to extract and share specific moves. For example, Ross described the community-created tool Fumen,[16] which can be used to create abstract

[16] http://harddrop.com/fumentool

representations of certain moves, as a useful training and discussion tool for *Tetris* players. For players that are already used to using such tools for *Tetris* or other games, compatibility would be useful. Mazeika also suggested she would be interested in uploading several of her own runs, to allow side-by-side comparison of the moments they contain.

Streamers and YouTubers

Videogame players who broadcast their play online, often along with their own commentary, are called streamers. These streamers might be trying to showcase a high level of play, or provide entertainment in other ways such as humorous commentary or insightful analysis. Streaming is also a hobby that has seen a great deal of growth since the mid-2010s, both on YouTube and the dedicated game streaming website Twitch.com. There are a variety of different kinds of streams, such as esports, "Let's Plays," and academic commentary. Esports are examples of competitive game performance which differ from speedruns because players are competing with each other directly. Let's Plays are generally more casual playthroughs of games which are often unfamiliar to the player. Academic streams are similar, with the primary difference being how they are framed around analysis rather than entertainment.

Separate from but related to this kind of streaming, there are also web-based videos (pre-recorded, rather than broadcast live) which consist of reviews, textual analysis or cultural commentary about games (e.g. the YouTube channels "Extra

Credits"[17] or "Tropes vs Women."[18] These videos often involve demonstrations of a concept that exists across a large variety of games, or less frequently a deeper dive into one specific game. This type of video was popularized on YouTube, which is still its primary venue, so there is not yet an accepted generic term other than YouTuber for these kinds of creators.

Streamer and YouTuber Interview Subjects

Stacey Mason is a PhD candidate in computer science at UC Santa Cruz and an academic streamer who provides commentary and analysis on her blog and Twitch channel, both called Cerebral Arcade.[19] She is one of the founders of ScholarsPlay,[20] an academic streaming channel.

Stella Mazeika, mentioned as a hobbyist speedrunner above, is also a contributor to ScholarsPlay.

Nick Junius, another contributor to ScholarsPlay, is the author of various videogame close reading and textual analysis videos available on YouTube.[21]

[17] https://www.youtube.com/channel/UCCODtTcd5M1JavPCOr_Uydg

[18] https://www.youtube.com/playlist?list=PLn4ob_5_ttEaA_vc8F3fjzE62esf9yP61

[19] http://cerebralarcade.com/ and https://www.twitch.tv/cerebralarcade

[20] https://www.twitch.tv/scholarsplay/

[21] https://www.youtube.com/channel/UCSunwmLFUsUTfovdQ4JyxXQ

What kinds of videogame moments do streamers and YouTubers seek?

While the primary focus of streamers is usually broadcasting live play, there are circumstances where seeking out particular moments can be beneficial. An esports player might be interested in reviewing their own play. A streamer or YouTuber who is providing an academic analysis or cultural commentary will be interested in searching for moments that serve to support their argument. For example, in preparing his analysis videos Junius looks for moments in games that he is already very familiar with.

Mason seeks moments she has not played through herself: “What I'd most want is a parameterized save state, which is easier to get than by playing myself.” She said that for certain kinds of streams she would like to construct demonstrations of play by “Julia Childs-ing”'---jumping between prepared locations rather than showing the tedious path between them, analogous to how the television chef would skip past time-consuming phases of a recipe demonstration.

How do streamers and YouTubers want to identify those moments in a search engine?

Streamers can vary in how familiar they are with the games they are playing, depending on the type of stream they are doing. They could search for play similar to their own by uploading their own screenshots or video. Searching for evidence to support a particular argument is more likely to be served well by a text search, if the

content of the video has been tagged or transcribed. Anything that adds more semantic information about the moment's content can help support performing this kind of analysis.

For Mason's quick navigation, she wants to use game-specific tags and parameters to identify target moments: “my character at a point in the game with particular set of items or with a particular set of bosses killed, etc.” For narrative-driven games “like *Detroit: Become Human*,” she desires a browsable visualization: “a plot map (a tree map of the plot) would be useful.” Such a visualization would be useful for both the streamer user of the retrieval system as well as their live audience, to help understand how the space is being navigated.

What will streamers and YouTubers do with retrieved moments?

Streamers want to continue playing from the point of a retrieved moment and demonstrate the gameplay possibilities at that point for their audiences. Mason said “it could be useful as a way to play through variations, particularly for story games to look at variations on a particular moment.”

For those streamers and YouTubers who make arguments, Mason and Junius suggested that finding moments ahead of time can help produce a variety of supporting examples. Junius commented that “When I wanted to talk about an event two-thirds or three-quarters through *Mass Effect 3*, I didn't have a save file and couldn't find one that I wanted. It took me about 15-20 hours or play to get 3-4 minutes of footage.” It took

this amount of time despite Junius using an existing community resource (specific to the game in question) to find an appropriate save file from which to begin again.[22]

Educators

Educator, here, refers to anyone who is teaching about the design, history, or interpretation of games. To better understand the needs of this community I spoke to several educators who have taught these topics at the university level.

Educator Interview Subjects

Elizabeth Swensen is a game designer and educator. She has taught courses on game design, both digital and physical, and on interactive media at the graduate and undergraduate level at the University of Southern California and at UC Santa Cruz.

Henry Lowood is a digital curator, scholar, and educator. He has taught undergraduate courses on the history of digital games and graduate level courses on the curation of new media at Stanford University, at San José State University, and at UC Santa Cruz.

Nathan Altice is a scholar and educator. He teaches undergraduate courses on game design and games systems as well as graduate courses on media archaeology at UC Santa Cruz.

[22] http://www.masseffectsaves.com

What kinds of videogame moments do educators seek?

The educators that I spoke with primarily expressed interest in the moments in games that can serve as examples for the concepts that they are teaching. Swensen suggested that for her lessons she might search for a particular kind of jump in a platforming game, systems for tracking the morality of a player's choices in various popular games, or how different games have handled text-input using analog sticks. Altice discussed sharing moments that convey a particular aesthetic experience. When teaching game design, it can be useful to trace an example in several ways. Swensen suggested that an educator might be interested in the most famous or prototypical example of a mechanic as a prototype to present to their class (e.g. jumping and running in *Super Mario Bros.*), and Altice suggested that they might wish to connect this example to more current incarnations, which might be more familiar to their students (e.g. jumping and running in *Super Meat Boy*. Altice, Swensen, and Lowood all expressed interest in the ability to trace mechanics backwards to find more prior examples (e.g. jumping and running in *Donkey Kong* or *Frogger*).

How do educators want to identify those moments in a search engine?

The educators I spoke with currently find the moments they are looking for using text-based search engines (e.g. Google or YouTube). Swensen and Altice both said that for more obscure examples this means hoping someone has already done the work of capturing and cataloguing it, or searching for a moment that they know to be close to the one they are looking for and scrubbing through videos to find it. Swensen

indicated that she might reach out to her friends or peers for help. Altice said that he has infrequently done the work to capture a moment himself, but often this is impractical because it is a time-consuming process or because it requires rare specialized hardware. Swensen suggested that educators who are interested in searching for mechanics across games might find a controlled vocabulary helpful (e.g. with keywords for jumping, text input, morality systems, action drafting, etc.). Altice and Swensen both suggested (prior to priming in the interview protocol) that a screenshot might be sufficient to connect to a moment, similar to conducting a reverse image search. Although Lowood is one of the contributors to the Game and Interactive Software Scholarship Toolkit (GISST), a component of the GAme MEtadata and CItation Project (GAMECIP), other interview subjects also expressed a desire for the ability to identify moments using citations (Kaltman et al., 2014, 2017).

What will educators do with retrieved moments?

Swensen and Altice wanted their students to be able to replay the moments they have found, to have the opportunity to experience those moments for themselves. A video alone can be useful, but, according to Swensen, video of an interactive moment is more like a description of the experience than it is like the experience itself. She notes: “Unfortunately, what's amazing there is not what you see but what you feel, and how seamless it is.” Altice also mentioned referring during one of his lectures to a sequence in the game *Silent Hill 2* (2001) where a player rowed a boat to a lighthouse, using analog sticks as a proxy for physical oars, where the difficulty and pace of this

sequence contribute to the game's atmosphere. He lamented that classroom facilities typically do not have the capabilities to facilitate this kind of experience, which would require at minimum some kind of emulation system available to all students. He also expressed a concern about finding a clean example of a moment for classroom demonstration, saying: “One of the ironies of teaching game design is that we have to rely on secondary sources, e.g. a Let's Play with some jackass narrating over.” Swensen suggested that prototypical or classic “best” examples (from platforms students may not even have access to) can be useful for conveying the core of a concept to students, while exploring older examples as a way of tracing the origin of an idea or mechanic that students are developing can provide useful context and suggest interesting alternatives to investigate for student projects. Altice offered that modern examples, and examples from platforms that host games actually played by the students, can make the material more immediately relatable. In either case, both Swensen and Altice emphasized that it is important that students be able to see and experience concrete examples of illustrative moments rather than just read about or hear descriptions of those moments.

Developers and Designers

Developers and game designers, here, are those whose work includes the creation of videogames and the tools to support videogame creation.

Developer and Designer Interview Subjects

Ted Lim is a game developer for NetEase NA, where his work has included producing feature design documentation, conducting user tests, collecting player feedback, and designing interactions. NetEase NA is focused on developing free to play mobile games, including games in the massively multiplayer online; action role-playing game; card game; endless runner; augmented reality; and puzzle game genres.

Batu Aymetiz a PhD student in computational media at UC Santa Cruz and a game designer. His research involves investigating the way that games are designed, including work on a variety of mixed-initiative tools that help designers do things like create better tutorials, or balance competitive games. He has designed a variety of physical and digital games for entertainment, education, and artistic purposes.

Elizabeth Swensen, first mentioned in the section on educators above, is also a game designer. Her work has included educational strategy games such as Application Crunch (2011), an Indiecade finalist that teaches students about accessing higher education.[23] Her other educational games have focused on metacognitive development, using strategy to teach historical thinking, self-advocacy, and encouraging career aspirations in STEM fields.

[23] http://elizabethmakesgames.com/application-crunch/

What kinds of videogame moments do developers and game designers seek?

Lim, as a developer of mobile and free to play games, was interested in understanding specific moments in the player's life-cycle that cause “fun or pain.” He also expressed interest in looking for similar moments in competitor's games. Unlike most of the examples from the interviews I have described so far, these moments are defined not just within the game itself, but by the context of player reactions. Swensen suggested that a developer would be interested in finding moments in their own game in development that were examples of a bug. Aytemiz, whose work includes developing tools to assist in game development, expressed interest in finding moments where those tools could be useful. One example he gave was a tool to help create better tutorials, which would be well served by finding the first moments in a game where the player needs to use a new skill to progress.

How do developers and game designers want to identify those moments in a search engine?

Lim offered that developers are often interested in player reactions that are tied to specific moments, and would like to be able to search for these reactions. One way to obtain these reactions would be linking moments to other data recorded during their own playtests. For competitor's games, Lim also suggested collecting reactions from online play (e.g. in Let's Plays on YouTube, or streaming on Twitch), which would then need to be linked to specific moments from the attached gameplay video. These reactions will vary with player experience, so it would also be useful to have

information about the player, particularly including their degree of experience with the game (e.g. a new vs. experienced player). Lim also indicated that, for their own games, developers are often well equipped to define moments at a high level of specificity (e.g. at this point in the level progression, with these items in your inventory, etc.). For tracking down a bug, Swensen suggested searching from a screenshot of an example. Swensen and Aytemiz both expressed interest in using a controlled vocabulary to identify specific kinds of moments (e.g. tutorials, or boss fights) in the course of development.

What will developers and game designers do with retrieved moments?

Swensen suggested using moments to find more examples of a bug, to understand and eliminate it. Lim cautioned that this approach would need to compete with or complement existing quality assurance practices. Lim also suggested that linking player reactions to specific moments can help a developer decide what needs to be cut, added, or changed. This could mean finding a moment that players enjoy, figuring out why, and changing the game to make similar moments more likely, or finding a moment that players find unnecessarily frustrating, and eliminating it. Alternatively, he offered that for a free-to-play game both kinds of moments can be opportunities to add monetization. On a different note, Lim also suggested that being able to search for specific moments might be helpful for communication within a team during development to make sure everyone has a common understanding of the team's goals, particularly when an older game is being referenced for inspiration. He also

offered that easy familiarity with popular moments and reactions from competitor's games might suggest certain directions for development.

Scholars

For this user category, I defined scholars as those whose work with games is primarily critical or interpretive. This includes any category of work that seeks to understand the history or future of games, game platforms, or the people who play them.

Scholar Interview Subjects

Nathan Altice, in addition to his role as an educator mentioned above, is a games studies scholar. This work has included writing about the history of computing, platforms, and culture. Altice is the author of I AM ERROR (2015), a book about the history of the Nintendo Entertainment System (NES) which emphasized its cultural role and the impact of its physical affordances and limitations.

Henry Lowood, whose work as an educator was mentioned above, is also a digital curator and scholar. He is the curator for History of Science & Technology Collections and Film & Media Collections in the Stanford University Libraries. His scholarship has included writing about the history of games, machinima, and documenting gameplay as a kind of athletic performance.

What kinds of videogame moments do scholars seek?

Lowood and Altice indicated that scholars are interested in the moments that can help them develop a deeper critical look at a game. Altice said that this would include moments from game platforms that are not as well studied, and those that do not have good emulators (e.g. the Bally Astrocade). Lowood said that scholars might be interested in moments that are examples of superior or athletic play, or moments that are part of historical events, and those moments that changed how the game was played. One example of such a moment, in a non-videogame context, took place recently in the famous Go match between (human) expert Lee Sedol and (machine) AlphaGo. A certain moment (Move 37) represented a pivotal event, defined by its relation to societal context: the machine chose an action that would have been rated as next to useless by human master players.[24] The challenge for information retrieval is knowing that the relevance of this moment is not captured in the state of the game board but instead almost entirely in the context.

How do scholars want to identify those moments in a search engine?

Scholars indicated they would be likely to use (Lowood said) "99% text based search," including text descriptions of a known event (e.g. "the gnome demonstration" for a specific protest that took place inside the game *World of Warcraft*). Text search might also include description of game mechanics, general or specific (e.g "the rocket

[24] https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/

jump", or "*Super Mario* zero [minus] world glitch"). Altice also provided several examples of ways that scholars might wish to add complex filtering to their searches (e.g. "all examples before 1985," or even "all horizontally scrolling games with a blue background from before 1985, in all consoles and arcades," "Mario on Yoshi in an underwater level with 8 lives remaining,", "all games with bosses," or "vertical scrolling in all games before 1985, across consoles and arcade"), and said that he would also be interested in adding platform specific characteristics to his search (e.g. a specific way the audio processor is being manipulated). For some of this work, Altice indicated there are already communities which have existing standards, e.g. junk headers in NES ROMs, while for other less popular platforms creating a search system might mean inventing these standards. Lowood, and Altice expressed an interest in being able to search for specific moments from screenshots, and Altice also expressed an interest in being able to take video as search input as well. Altice was familiar with the GISST citation system (and Lowood assisted in developing it), where a persistent URL link can connect to a particular moment. Lowood remarked that enabling that kind of search also means having a repository built around the tool. Lowood also suggested that in rare cases a scholar might want to search with code itself, to discover what mechanics or play the code enables.

What will scholars do with retrieved moments?

Lowood and Altice expressed interest in being able to validate a moment from a screenshot. This means being able to determine if it represents a possible, achievable

state, and if possible linking it not just to the game it came from, but to a particular version (ROM) of that game. Altice gave the example of taking a screenshot from a magazine, and finding out if it came from an early prototype or variant of a game rather than the released version. Altice had also suggested using specific platform characteristics for search, and one example he gave for how this would be useful was in early Nintendo games, where ways of manipulating the audio processors which were hallmarks of certain (uncredited) composers. Lowood and Altice both expressed interest in being able to continue play directly to a moment in a format as accurate as possible to the original. This could mean emulation in the web browser, or a downloadable save-state which can be loaded into an emulator. For this second possibility Altice noted that there is no canonical or ideal save state, and each save state is specific to an individual emulator. Lowood also noted that for some kinds of games this will not be possible at all, such as a massively multiplayer online game which requires external servers to play. Lowood expressed interest in knowing what specific code enables, and if possible understanding what happens when you change it. Lowood also suggested that a search system that knows how to find a "rocket jump," if it understands it at a semantic level, could also help a researcher find games with mechanical precursors.

Lowood explained that the scholar's need to make new discoveries is in some ways at odds with search. Anything that is made available by a search has, in some sense, already been found. Following the examples given above, if a scholar conducted

a text search using the terms “rocket jump” or “gnome protest” to find a particular moment, any moment the search engine returns must have already have been tagged (in inferred to be intended as tagged) as containing that mechanic, or belonging to that event. For this reason, and because they are often looking for unexpected or unknown connections, scholars are interested in having as much information as possible attached to each moment (including information about the social context of the moment for players and for other scholars), and in the ability to conduct the kind of complex compound searches that might never have been anticipated.

Findings Summary

I started this work by asking: What kind of videogame moments do various types of users seek? Varied responses indicated that speedrunners wanted to see examples of tricks, streamers and YouTubers sought out moments they knew well, educators looked for prototypical examples, developers looked for moments with clear positive or negative engagement, and scholars looked for whole categories of moments.

Next, I asked how people wanted to identify those moments to a search engine. Interestingly, there was a good deal of commonality here, with many user categories interested in using screenshots or text search. When the screenshot search suggestion was offered by speedrunners, it was before they had been presented with the anchoring demonstration. Some users also were interested in search via online video, which has some overlap with search by image. There was some variation in how users wanted

text search to operate, but generally that was a matter of wanting a controlled vocabulary, and one that understood terms already in use by the community.

Finally, I asked: What will they do with retrieved moments? Here again I saw a great deal of variety. Speedrunners seek moments for self improvement, streamers and YouTubers seek moments to support an argument, educators seek moments to provide context and examples, developers seek moments to understand player reactions, and scholars seek moments to find interesting and unusual connections.

While I saw a great deal of variety in both the kinds of moments that users are interested in, and in what they want to do with those moments once they have found them, I saw more commonality in the ways that users want to search for moments. This is promising for the development of systems for videogame moment retrieval that seek to serve many types of users.

Technical Recommendations

Findings from the interviews I conducted suggested further technical work in game moment search, such as the inclusion of additional game platforms, more intelligent text queries, extracting gameplay from videos in the wild, and adding searchability to personal data sources. These interviews directly inspired further work in my own research group to support natural language queries for videogame moments from two modern videogames (Zhang & Smith, 2019).

New Platforms

While some speedrunners and scholars are interested in search on classic game consoles, everyone I spoke with was interested in search for other platforms. Our interview subjects were primarily interested in PC games, and in some cases also interested in mobile (e.g. Android and iOS games). For applications to be relevant to new or in-development games, retrieval systems cannot rely on the existence of pre-recorded interaction traces. In web-search parlance, future systems will need to be able to crawl the content in full-scale games for modern game platforms automatically.

There is ample work on automatically playing videogames (Yannakakis & Togelius, 2018), but to meet the needs articulated in the interviews above the challenge will be crawling the contents of a wide variety of culturally-impactful games. This would ideally be done with little to no game-specific engineering, in order to be useful across a large variety of past and future games. Another challenge for automated game exploration will be the inclusion of those games for which there is no obvious score to maximize or quantifiable winning state to achieve, e.g. *Mario Paint* (1992), *Animal Crossing* (2001), etc., because they often rely on reward-driven approaches.

Textual Queries and Domains-specific Vocabulary

All types of users were interested in text search, although they had different ideas about how they would use it. Speedrunners, developers, and educators were all interested in search based on vocabulary that was specific to their community, or

developed to support this kind of search. Scholars, streamers, and others were also interested in being able to search with free-form descriptions. Altice's suggested query "Mario on Yoshi in an underwater level with 8 lives remaining" presents the specific challenge of building a system that can understand game-specific distinctions that might not have been considered at the time moments were indexed.

Based in part on the issues raised during these interviews, Xiaoxuan Zhang (whose work introduced the problem of videogame moment search) completed a master's degree with a project specifically focusing on allowing users to search with natural-language queries (Zhang, 2019). This initial work combined ideas from natural language processing and computer vision to extract several different senses of keywords from screenshots taken from gameplay videos (as downloaded from YouTube). While this system could not understand domain-specific vocabulary, it could identify moments referencing everyday concepts (e.g., horses, sunsets, and selfies) for games that did not mention these concepts in their title and description, and it was not available in other paratext (Rockenberger, 2014) or in other writings about them (e.g. Wikipedia articles).

Videos in the Wild

Developers and speedrunners expressed interest in being able to search using video data in the wild. Most speedrunners share their runs via online video. Developers would want to pair this with player engagement and reactions to find out which of their

(and their competitor's) ideas are being enjoyed by their audience. Scholars asked for as much information as possible, and this would serve their interest in understanding moments in their cultural context. Transcripts of live game play streams would also be useful for enabling more effective text searches, by making it possible to search for both on-screen text and for player reaction vocalizations. Many videogames do not inherently possess a textual component, so there is the problem of aligning the distribution of moments in a game with the specific vocabulary familiar to specific classes of users. In the case of streaming applications, there is the additional technical challenge of processing ephemeral data as fast as it is produced. Future research should find ways of making use of bulk gameplay video content including on-screen text, player reaction vocalizations or facial expression, and audience comments on the video.

Personal Data Sources

Speedrunners, streamers, and YouTubers were also interested in a form of introspective search, being able to upload data from their own playthroughs. For a speedrunner or streamer this might be a way of evaluating or sharing their own performance, comparing their play to world records or to their own previous attempts. This is similar to how users may want to search within their own photos or other personal information, and future work in this area should build on the existing body of Personal Information Management (PIM) research (Jensen et al., 2018).

Research Recommendations

This chapter describes preliminary work in the interactive media retrieval domain, using an interview-based qualitative research method. There are several other methods that previous research in the user-needs analysis community suggest would also be fruitful. For example, researchers interested in book and movie recommendations coded the content of relevant forum posts (Bogers et al., 2018). Researchers examining image search interfaces supplemented their informal interviews with system designers with a large scale analysis of search logs (André et al., 2009). Additionally, the interdisciplinarity of game studies suggests a range of alternative methods that could potentially be used to supplement these results (Mäyrä, 2008, Chapter 8), which could include those drawn from a range of humanistic, social science, and design research traditions.

Staying within the interview methodology, when this project concluded each interview conducted was still generating novel recommendations. This indicates that future interviews would likely produce additional fruitful recommendations, up to the point of saturation where responses will become repetitive.

There are also additional user categories identified during the course of the interviews to explore beyond those included in this initial study (e.g., distributors, platform holders, translators, fans, etc.), and research methods beyond interviews (e.g., surveys, participatory observation) which might lead to additional insight. In the course

of preparing this work I had to select specific individuals to contact, and was only able to report based on those who were able to respond to my queries. This means my results are missing several important stakeholder perspectives. One of these is the game distributors who manage storefronts where potential players can discover and purchase games, along with their consumers. I can speculate that they might be interested in finding the best possible moments or images to sell games, in measures that might make discovery in the context of a massive catalog of games easier, and perhaps in moments that are a better representation of a game than the developer's own carefully curated (but potentially misleading) screenshots. Other stakeholders that should be considered for future studies include game translators, quality assurance testers, and even in-game advertisers. Future work could provide more grounded answers to how these other kinds of users would use videogame moment retrieval systems.

Contributions

In this chapter I have identified a variety of needs that inspire some of the work described in subsequent chapters of this dissertation. The set of user needs that I have identified in these interviews is incomplete, but it is also large and diverse enough to suggest many directions that will not be explored here. By describing concrete challenges to be addressed in future technical systems, I hope that this work offers an anchor for those who will conduct technical research in this new information retrieval domain with a user-centered design approach. The research recommendations and user

requirements described in this chapter are contributions to the field of information retrieval where it intersects with game studies.

Chapter 4. Exploring Personal Gameplay Videos

Project Summary

In this chapter I describe a search tool for exploring videos of gameplay, and significant insights and prototypes generated in the course of its development. This is the first of two projects to enable search within and across games that I describe in this dissertation. The primary outcome of this project was the **Gameplay Video Explorer**[25] search tool. The tool is presented on a web site which displays these videogame moments in relation to each other, in a navigable two-dimensional space, which contains meaningful spatial arrangements generated from videogame moment vectors by a dimension reduction process. In the process of developing this tool, inspired by prior work on word vector analogies in natural language (Mikolov et al., 2013), I demonstrated that these analogical relationships also exist for videogame moments, and that they can enable a new kind of search. Following a proof of concept, I generated an interactive explainer (Hohman et al., 2020) for the concept of search-by-analogy[26] using videogame moments. Both the explainer and the search tool build on previous work that transforms videogame moments into vectors (Zhang et al., 2018), and creates visualizations of videogame moments based on those vectors (Zhang, 2018; Zhang & Smith, 2019). The search tool offers the novel ability to work from user-provided

[25] https://barrettrees.com/GameplayVideoExplorer/

[26] http://barrettrees.com/search_by_analogy_interactive/IntroductionToSearchByAnalogy.html

moments, in the form of uploaded personal gameplay videos, and to compare moments across multiple source domains (e.g., to do comparisons between games). It is extensible and flexible by being able to incorporate arbitrary embeddings to generate the videogame moment vectors that it operates on. Updating the embedding model can improve the system's game specific knowledge (or replacing it entirely, the system can be adapted for non-game visual information). Beyond the Gameplay Video Explorer site itself, the project also includes an assessment of the system using modified search evaluation metrics (e.g., measures of precision adapted from nDCG). This system directly addresses interests that were identified in the requirements analysis in Chapter 3 (above), including speedrunners' and streamers' interests working with personal gameplay data, and scholars' interests in complex forms of search. For information visualization, this work contributes both a new interaction model for video exploration and a demonstration of that model in a new domain.

Prior Work

Generating Videogame Moment Vectors

The projects described in this chapter all build on previous work done for several games for the *Super Nintendo* console (1990), which made it possible to derive videogame moment vectors from screenshots (Zhan & Smith, 2018). This process involved using machine learning techniques on a corpus of videogame moments, which in this case consisted of both screenshots and the contents of the videogame system's

memory. This process generated embedding functions which made it possible to transform a videogame screenshot into a 256-dimensional vector. Because the training included the system's memory, and not just screenshots, these embedding functions implicitly contain system-level information that is not necessarily visible on screen (See Figure 5). This information varies between games, so each of these embedding functions is specifically tuned for the game that it was trained on. The typical training set for these functions was several thousand videogame moments, generated from a combination of human play and various automatic exploration strategies. This is a miniscule fraction of the game's state-space of possible moments, and it will not include many moments reached during ordinary play. This means that a new playthrough will provide novel screenshots that were not in the corpus that the system was trained on, but for which it is still capable of generating meaningful vectors. These videogame moment vectors resemble the document vectors used in other search systems, and they can be used to facilitate comparisons between moments (e.g., the less the distance between two vectors, the more similar the moments that they represent), and videogame moment search. These vectors are also amenable to dimension-reduction based visualizations, such as t-SNE and UMAP, which can make salient relationships between moments and clusters of related moments.

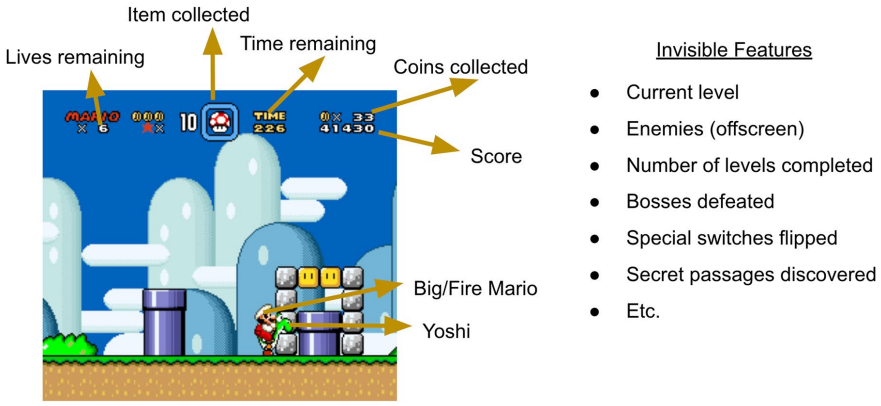

Figure 5. Hypothetical examples of the visible and invisible features of a videogame moment from Super Mario World (1990). Along with visual information, game systems information such as this is encoded in the videogame moment vector, but not necessarily these specific features.

Videogame Moment Space Visualization

Moments in nonlinear media can be conceptualized as points in an abstract space, an idea made more concrete by the existence of the videogame moment vectors described above. These vectors have a large number of dimensions (in this case 256), but this can be reduced to a human-navigable set of two or three dimensions using algorithms designed to preserve information about the relationships between points in this space. Because of this dimension reduction, some information is necessarily lost, but creating an interpretable visualization will surface other information that would not have been easily discoverable. The variety of dimension reduction techniques (e.g. t-SNE, UMAP), and the choices made when implementing them, can be thought of as various lenses which can be used to examine different aspects of the videogame moment state space. One example of how this kind of space might be navigated is

provided by the Videogame Moment Space visualization (Zhang, 2018), which uses t-SNE to map the vector embeddings of videogame moments (See Figure 6).

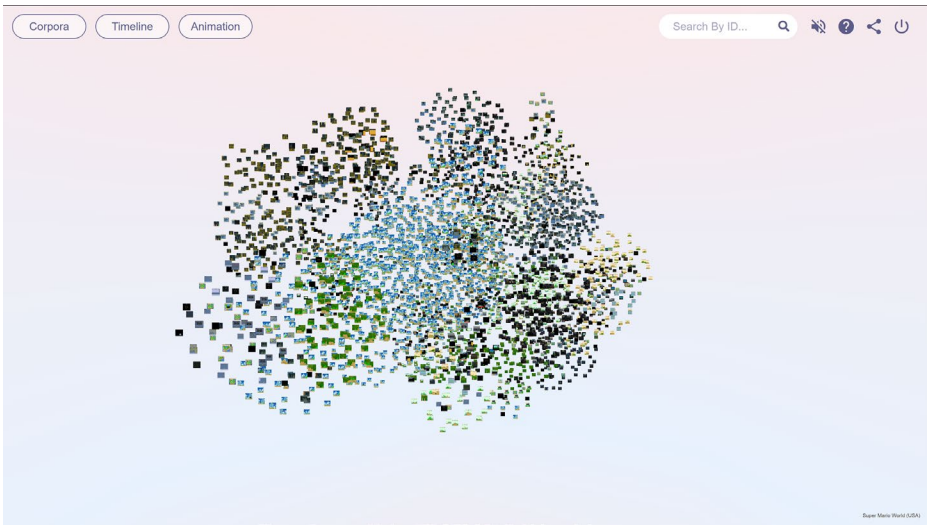

Figure 6. Videogame Moment Space visualization (Zhang, 2018) for moments from several playthroughs of the game Super Mario World (Tezuka et al., 1991).

A user can scroll, rotate and zoom around this game space, in a familiar interaction paradigm similar to that of using an online map (e.g., Google Earth, Apple Maps, etc.) but in this case in three dimensions. Observing the spatial relations between videogame moments in this space was the starting point for conceptualizing a new videogame moment search paradigm, *search-by-analogy*.

Search-by-Analogy Proof of Concept

The idea of turning videogame moments into vectors was inspired by previous research that created vector representations of words (Bengio et al., 2003; Rumelhart

et al., 1986). One way of creating these word vectors is based on the statistical co-occurrence of words in a corpus. Surprisingly, word vectors based on these statistical properties also encode semantic information. A famous result demonstrating this property is described by the equation "man - woman + king = queen." If you replace each of the words in this equation with its associated word vector representation, the equation describes a mathematical relationship between those vectors (Mikolov et al., 2013). One important subtlety is that the new vector generated by subtracting "man" from "woman" and adding the result to "king" will not be an exact match for any of the existing word vectors. However, out of all the vectors for words in the corpus it will be closest (as measured by cosine similarity) to the vector for the word "queen." This is an important detail when thinking about search-by-analogy, because we can now think of the equation as generating a synthetic search "query" vector for which "queen" was the best result from the set of actual word vectors. Positing that similar analogical relationships might hold between videogame moment vector representations, it seemed possible that this could be the basis for a new way of exploring relationships between these moments, and for a new kind of search.

Analogy in Cognition and Computation

Information retrieval, in addition to the sense described above, is also an aspect of human cognition, usually called memory (Tulving, 1972). An understanding of how retrieval functions in a psychological context provides insights that are useful for thinking about search, and particularly in considering different search paradigms. The

closest parallel to the search systems described above is probably long-term declarative memory, but there are also lessons to be learned from other aspects of human cognition. I draw some specifically from the concept of analogy, which is the ability to understand a novel context by drawing connections between it and a familiar one, mapping relationships between elements across distinct source and target domains. This capability has been called the core of cognition (Hofstadter, 2001), and there is already a history of research to build on that combines methods and knowledge from psychology and computer science.

Cognitive Modeling

Cognitive modeling is an area of research that attempts to create computational models that reflect some aspect of human cognition. These models are called cognitive architectures (Langley et al., 2009), and researchers have created several examples (e.g. Soar, ACT-R, and EPIC). These models are used to predict and validate existing behavioral data. For example, an ACT-R or EPIC model of a system of attention is generated, which is grounded in existing behavioral data. This model can then be used to generate predictions and test predictions about behavior more quickly and cheaply than actually running experiments with human participants. If subsequent behavioral experiments confirm these predictions, the model is either validated or revised. While the architectures above aspire to be generalizable, collaboration between psychologists and computer scientists has also resulted in a large variety of cognitive models of the process of analogy specifically (Gentner & Forbus, 2011), based on a variety of

approaches (connectionist, symbolic) to understanding the underlying cognitive processes (e.g. MAC/FAC, Copycat, ACME, etc.).

Cognitive Models of Analogy

This tradition led to the creation of a substantial number of models, and a complete survey is beyond the scope of this review, but I will discuss the MAC/FAC and SME models, (Forbus et al., 1995) which have the potential to be particularly relevant here, because they were specifically about the process of similarity-based retrieval. MAC/FAC built on the SME (structure mapping engine) computational model of analogy and similarity (Gentner, 1983). The core idea behind the latter was that analogy is a mapping of knowledge from one domain (the base) onto another (the target), which conveys that the system of relationships between the elements of each system are the same. The SME takes descriptions of a base and target system as input, and produces possible mappings between them as output. Thinking of the system as being asked to compare two domains, the output is interpretations of that comparison. The MAC/FAC model ("many are called but few are chosen") was an attempt to embed that process in a psychologically sound context. This began with the recognition that the majority of similarity-based remindings are mundane (e.g. seeing a bicycle reminds us of another bicycle). The goal of MAC/FAC was to generally produce these kind of mundane surface-level remindings, while still occasionally permitting the discovery of more obscure structural relationships (e.g. the link between the periodic table and octaves in music). The authors explained this as the result of conflicting computational

constraints on retrieval, and incorporated it into their model by using a two-stage process. In the first stage, a larger number of computationally cheap matching processes were carried out, which were not sensitive to structural similarity, to produce possible remindings. In the second stage, the SME is used (in parallel) on all of the candidates produced by the first stage, to find the best structural matches. Perhaps an approach based on this insight could be incorporated into future information retrieval experiments.

These models were created primarily as tools to develop a greater understanding of processes in human cognition, as in the model-experiment loop described above. Cognitive architectures were evaluated by their success in modeling the results of psychological experiments in human subjects. However, regardless of their accuracy in reflecting human cognition, they can also be considered as inspiration for non-human cognitive systemes. Attempts at computationally modeling the process of creativity, for example, can be used to inspire artifact generation systems, independently of their success at their ostensible purpose. Similarly, while cognitive systems that modeled analogical relationships were not a direct inspiration for search by analogy, there is a resemblance that it is only fair to acknowledge, and further investigation of these models and related research might suggest additional elaborations for the search by analogy concept.

Search-by-Analogy for Videogame Moments

Before videogame moment vector-based analogies could be incorporated into a new search system, it was necessary to demonstrate the technical possibility of search-by-analogy for videogame moments. I generated this proof-of-concept for two different games, chosen from the set of Super Nintendo games for which embedding functions were available (Zhang et al., 2018), and with which I had enough familiarity to generate analogies that were meaningful at the game systems level (i.e., analogies not based solely on visual properties). To generate example gameplay for the search-by-analogy proof of concept test I conducted a directed playthrough of *Super Mario World* (Tezuka et al., 1991) and *Super Metroid* (Sakamoto & Kano, 1994), using the BizHawk 2.2.2[27] emulator to record my play as a ".bk2" movie file. During these playthroughs I deliberately reached moments in each game intended to relate to each other analogically. Specifically, in my play of *Super Mario World* I reach a position in the early game level "World 1-1" without any power-ups (A), reached the same position in "World 1-1" with several power-ups (B), reached a point in the later level "World 1-2" without any power-ups (C) and reached that same position in "World 1-2" with power-ups (D). These moments (and a similar example for Super Metroid) are illustrated in Figure 7. For *Super Mario World* on the left: In moment A, Mario is in World 1-1, small, not riding Yoshi, and has no items. In moment B, Mario is in World 1-1, big, riding Yoshi, and has a mushroom. In moment C, Mario is in World 1-2, small,

[27] http://tasvideos.org/Bizhawk.html

not riding Yoshi, and has no items. In Moment D, Mario is in World 1-2, big, riding Yoshi, and has a mushroom. For *Super Metroid* on the right: In moment A, Samus is in a cave and has no upgrades. In moment B, Samus is in a cave and has a missile upgrade. In moment C, Samus is underground and has no upgrades. In Moment D Samus is underground, and has a missile upgrade.

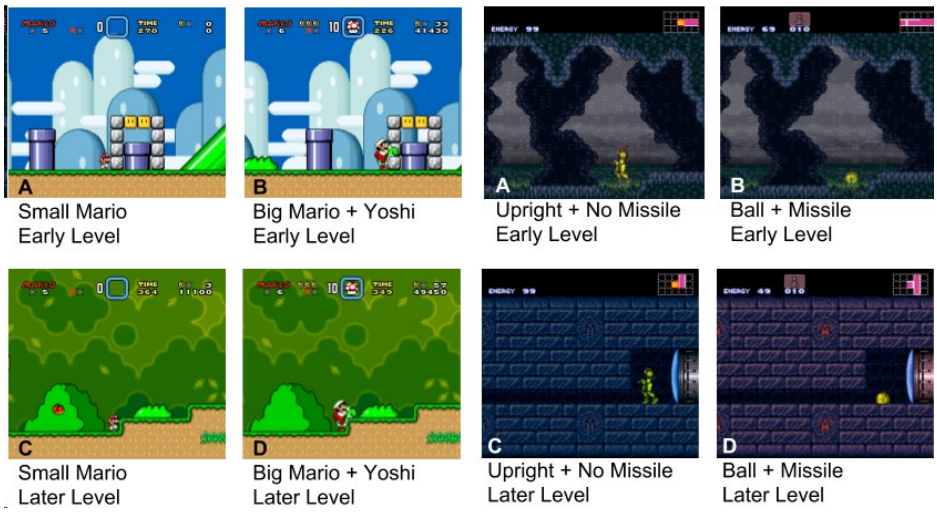

Figure 7. Example search-by-analogy query and result images for Super Mario World (left) and Super Metroid (right). Moments A, B, and C are the query moments set by the user, and D is the moment returned by the system to complete the analogy (A is to B as C is to D).

Once I had created the appropriate ".bk2" movie file, I used a Python script to rerun the gameplay in the BizHawk emulator while capturing a screenshot of each tenth frame, generating a set of 2377 screenshots from *Super Mario World* and a set of 4969 screenshots from *Super Metroid*. I then converted each screenshot into a vector representation of the moment it was derived from, using a pre-existing embedding that had been trained on a larger corpus of moments from the same game (Zhang et al.,

2018). With each moment converted to its own vector, I was able to generate a query vector, Q, which could then be compared to every other moment in the playthrough (See Figure 8).

```
def measure_similarity(embedding, a_time,b_time,c_time,test_time):
    a = embedding[a_time]
    b = embedding[b_time]
    c = embedding[c_time]
    e = embedding[test_time]
    q = c+(b-a)
    return (cosine_similarity(q,e))
```

Figure 8. Python code for finding the cosine similarity of the query vector (derived from example times) to all other moments in the corpus.

My pre-selected D moments (the ones I intentionally reached in my playthrough) were among the highest in cosine similarity, and were very close in time (within a few frames) to the highest. This result confirmed that search-by-analogy is possible for videogame moment vectors. This is a novel search paradigm in a novel domain, so a new search tool alone would likely be unfamiliar and confusing. Instead of creating such a tool directly, my next step was an effort to provide an accessible explanation of search-by-analogy for videogame moments.

Search-by-Analogy Interactive Explanation

Inspired by several modern web-based interactive essays that self-identified as explorable explanations, I set out to create a similar explainer for the concept of search-by-analogy in the context of videogame moments. The term “explorable explanation” was coined by Bret Victor (2011), as a form of text that is not content to be consumed, but rather provides an environment to think in. There are several currently popular

examples of the genre, such as the *Parable of the Polygons*, (Hart & Case, 2014) showing how small biases can lead to substantial social harm, and *Hack Your Way to Scientific Glory* (Aschwanden & King, 2015), which allows the reader to explore the scientific dangers of p-hacking. These explorable interactive systems allow an active reader to learn through play, which hopefully increases engagement and scaffolds the reader/user to a greater level of insight about the systems and relationships being modeled.

With the proof-of-concept phase described above complete, the second phase of this work began with the development of an interactive prototype. The purpose of this phase was to explore possible user interactions, in a way that could be rapidly iterated on. In order to do so I added sliders that would allow a user to change the moments chosen for the A, B, and C slots, while displaying matching screenshots for each, as well as a screenshot of the best match for moment D, and a graph of cosine similarity. This early draft is shown in Figure 9. Serendipitously, the ease of using these sliders to select moments also made it much easier to search for me to search for additional examples of analogical relationships that existed within the gameplay that I had already recorded.

The third phase was a JavaScript implementation of this interactive system, in order to enable wider distribution. This was necessary because the Python prototype could not simply be embedded in a website, and therefore would only have been accessible to expert users. The switch to JavaScript also made it possible to explore a

richer ecosystem of interactivity and visualization resources. The fourth phase was the improvement of this system by the development of a richer interactive data visualization. This work continued in JavaScript, supported by the D3.js[28] data visualization library. First, a new idealized version of the data visualization graph was conceptualized, with moments displayed as points in a space of A-B similarity (horizontal axis) and C similarity (vertical axis). Following this conceptualization, I implemented a new interactive, shown below in Figure 11. Each moment's closeness (cosine similarity) to the search vector is visually represented with both the color and size of the dot on the graph. The user selected points A, B, and C are highlighted, and the point chosen by the system as the best match (the point with the highest cosine similarity) is highlighted as point D. Mousing over any point in the graph displays a screenshot of the associated moment, and of all three of these scores.

The final deployment phase was the creation of the appropriate framing context, in narrative and images, for the wider distribution of this interactive demonstration in the context of an explorable explanation. This involved adding descriptions of the videogame moment vectors, an illustrated walkthrough of the process of search-by-analogy (See Figure 10), and contextual background information about the inspiration and contributions of this work.

[28] https://d3js.org/

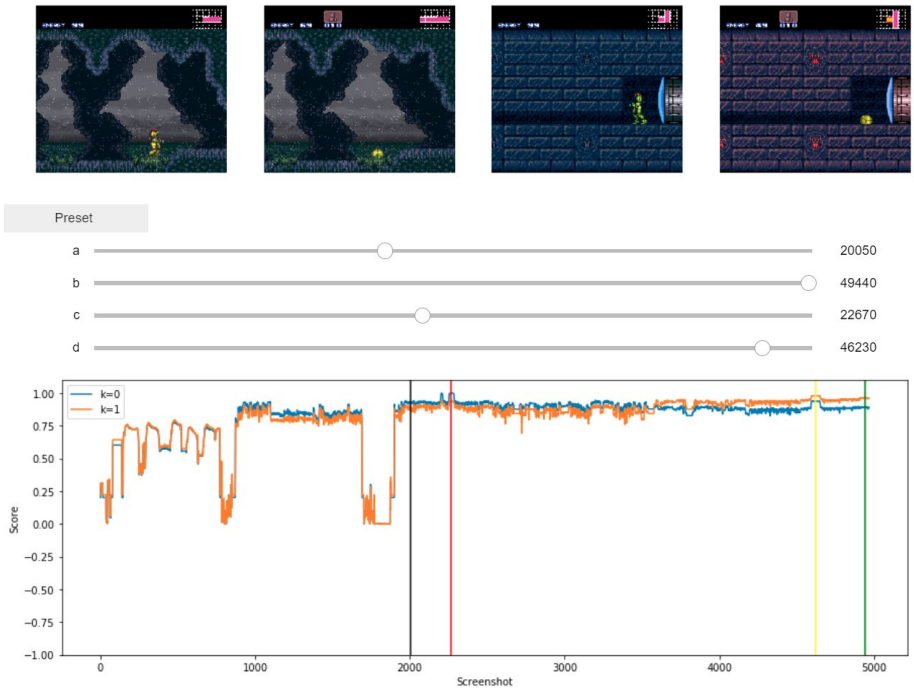

Figure 9. An early stage interactive prototype for search-by-analogy in videogame moments built in Python with Jupyter Widgets. Elements include screenshots for moments A, B, C, and D (top), sliders allowing a user to set moments (middle), and a graph displaying cosine similarity for all moments (bottom).

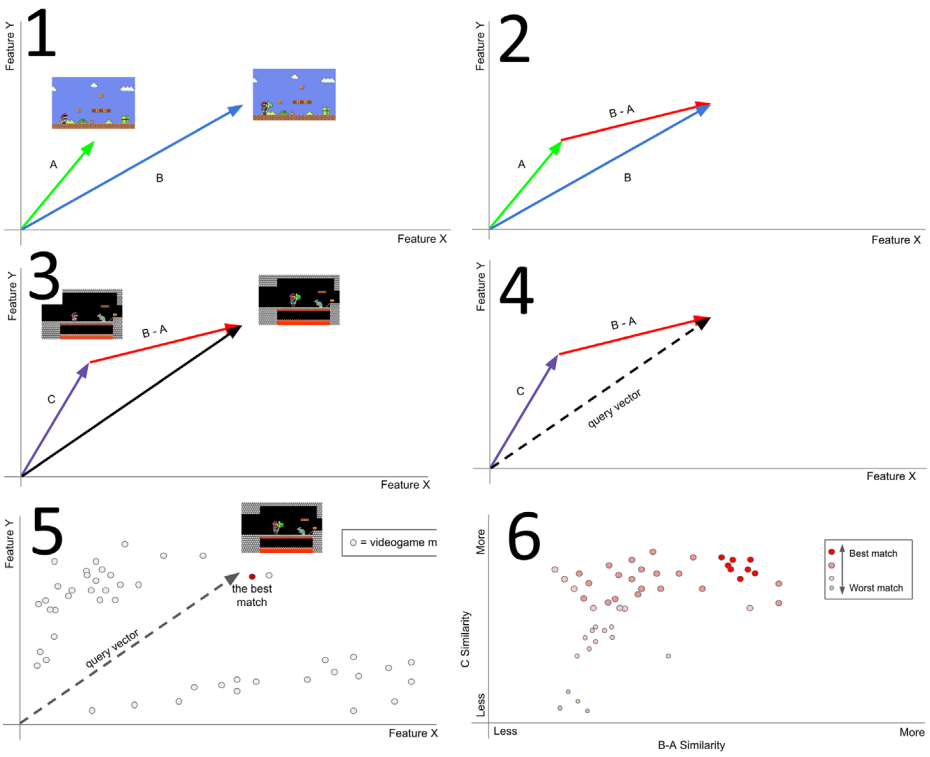

Figure 10. Step-by-step illustrations search-by-analogy for videogame moments from my explainer, preceding the interactive element. The first graph (1) depicts moment vectors for hypothetical videogame moments A and B. The next (2) depicts the B-A vector of the difference between those moments. The third graph (3) illustrates the addition of this B - A difference vector to C to form a synthetic moment vector, which then (4) becomes the query vector. The fifth graph (5) illustrates how a search system finds the best match for the query vector from the set of moments that exist. The final graph (6) exchanges the arbitrary axes for ones that provide meaningful information (similarity to moment C and to the difference vector), in a way that matches the interactive visualization.

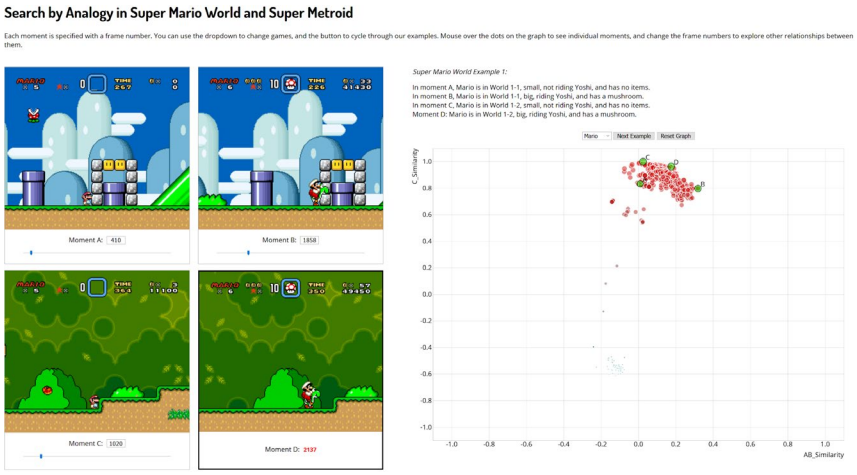

Figure 11. Search-by-analogy explorable explanation interactive graph. The user selected points A, B, and C are highlighted, and the point chosen by the system as the best search-by-analogy match (the point with the highest cosine similarity to the (B-A)+C vector) is highlighted as point D. Mouseover of the points on the graph overlays the highlighted moment's screenshot and similarity scores.

Gameplay Video Explorer System

The idea of searching for moments within a video is not unique to gameplay. Google has added the identification of key moments to their own search results, for moments like individual steps in how-to instructions (Baheti, 2019), but this feature currently requires video uploaders to provide explicitly named and timestamped segments. The search-by-analogy explainer that I described above demonstrates a novel search paradigm and demonstrates video moment search based on extracting the features of specific moments without requiring this kind of annotation. However, the interactive component of that demonstration is only capable of operating on moments that I have personally provided to the system. The next step was the creation of an

actual search tool, which makes it possible to perform this kind of search-by-analogy on arbitrary, user-provided, videogame moments. The next step is a search tool, intended to be accessible for games researchers as well as for non-data scientist experts, which allows them to search within their own gameplay data. This new tool is described below, and is currently available online.[29] Examples of user categories that might have their own gameplay data were identified in the requirements analysis above, and include scholars, streamers, speedrunners, and other hobbyists with some technical interests. Because this system is intended primarily for technically sophisticated users, and because it is also a prototype intended to be used for exploring a new kind of analysis, when there is any uncertainty, exposing settings and controls to the user is the default assumption.

Possible use patterns

The Gameplay Video Explorer system takes in a user-provided video of gameplay, extracts still images from that video, transforms those screenshots into videogame moment vectors, and allows the user to perform a search-by-analogy in the domain of those vectors. Below, I speculate on possible use patterns for the Gameplay Video Explorer system and similar future systems, including video and gameplay data sources, and analogies within and across domains.

[29] http://barrettrees.com/GameplayVideoExplorer/

Possible gameplay video data sources

There are several possible sources of gameplay video data to be considered for this search tool. The first is personal gameplay videos. These kinds of videos are possessed by speedrunners, which review their runs in the same way that athletes might review footage of their performance. YouTubers and other critics who create video essays about games, also generate gameplay footage in the process of illustrating their points. Similarly, educators often collect personal video of game moments for class demonstrations. Scholars (a category often overlapping with educators and critics) may do the same in order to find examples to support their arguments. Streaming focuses primarily on live performance, but also generates gameplay video recordings. The popular videogame streaming website Twitch[30] currently automatically generates these recordings, and if requested by the streamer stores them indefinitely. Several web tools exist for downloading these videos. Videogame streams, on any site, are often supplemented with additional non-video information, including physiological monitoring and performance data. Outside of these personal sources, many videos of gameplay exist online. Examples include well known speedruns, games from public tournaments, and footage from "Let's Play" videos and streams. Game makers promote their work with short trailers that sometimes include video of gameplay.

[30] https://www.twitch.tv/

Other data sources

Other potential sources of data include gameplay records in non-video formats, and non-game video. One example of non-video gameplay records are videogame savefiles. These files implicitly encode videogame moments, and the existence of game-specific community resources[31,32] demonstrates an existing interest in sharing these moments. Other non-video sources include replay files, such as the ".bk2" movie files mentioned above. These are more complete records of playtraces than video, including all user input, and are a non-video way of sharing speedrun records.

The system is not limited to working with gameplay video. Replacing the embedding model trained on game data with one trained on other kinds of visual information would effectively specialize it for other kinds of video. The might enable comparisons between different versions of the same film or trailer (e.g. comparing a director's cut or translations). Even without effective specialization, a user could quickly identify elements that were added or removed because similar points in different versions would overlap (Chang & Smith, 2020). This could also be applied to video examples of physical analogies, at different levels of abstraction, similar to the analogy questions that were once part of the SAT college admissions test. For example, a one embedding might allow the system to complete "a red triangle is to a red circle as a blue triangle is to a _____," while another might allow it to complete "an empty

[31] http://www.masseffectsaves.com/
[32] https://savesforgames.com/

shed is to a full shed as an empty garage is to a _____," and so on for even more abstract relationships.

Within-domain analogies

Within domain analogies are similar to those described in the search-by-analogy examples above, taking place entirely within a single game. Getting to a point in a level with a different score or different number of remaining lives would be similar to getting to the same point in another level with an identical difference in game statistics. For example, the "Mario on Yoshi in an underwater level with 8 lives" query offered by a scholar in the requirements analysis above (see Chapter 3) could be implicitly described by identifying three moments. For this query, point A would be "Mario in level 1-1", point B would be "Mario in level 1-1 with Yoshi and 8 lives," and point C would be Mario in any underwater level.

These within-domain analogies are often produced by experts in a domain in the course of their ordinary work for communicating with other experts, and generally rapidly generated, modified, and discarded (Dunbar & Blanchette, 2001). It is probable that these domain experts do not explicitly think of (or "mark") these common and close comparisons as analogies, since they are not connecting disparate contexts, but they are nonetheless structurally identical to any other kind of conceptual blend (Fauconnier & Turner, 2008). Based on these observations, these use cases are less

likely to produce the dramatic surprises of cross-domain analogies, but they are more likely to be incorporated into the actual work flow of potential users.

Cross-domain analogies

Where within-domain analogies work completely inside a given game (or other context), cross-domain analogies are about observing the same relationship across different contexts. In contrast with within-domain analogies, cross-domain analogies are rarely produced by individuals without explicit prompting. A games example might include a character with and without a specific power up in one game (e.g. Mario with and without a Fire Flower), and a character with and without a similar power up in a completely different game (e.g. Samus with and without the Super Missiles). This is not currently possible with existing embeddings, which are specific to individual games. This could even extend beyond games, for example extending the same analogy to a television show where a character acquires some sort of meaningful upgrade in power (e.g. Clark Kent becoming Superman, or Adora becoming She-Ra, Neo becoming The One). Perhaps a more sophisticated embedding might capture even more abstract differences.

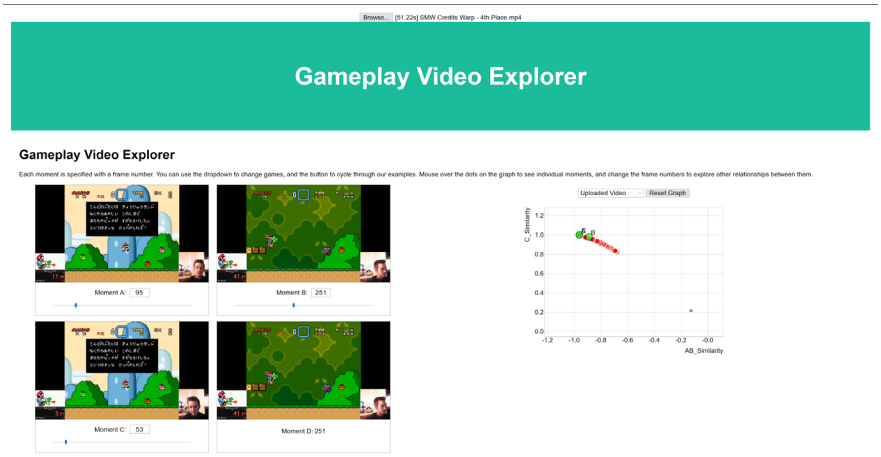

Figure 12. An early version of the Gameplay Video Explorer site with moments from a speedrun of Super Mario World (Tezuka et al., 1991) loaded, displaying a search-by-analogy style graph (B. R. Anderson et al., 2019).

The Gameplay Video Explorer site (See Figure 12) is capable of visualizing relationships between moments in the same way as the previous explainer, using axes of similarity. These results are supplemented with a dimension-reduction based visualization (t-SNE), which can identify other kinds of clusters of videogame moments and enable more exploratory uses. The site has also been extended to allow the inclusion of data from multiple video sources, in order to make comparisons across games possible. Additional datasets are be displayed as differently color-coded in the visualization (See Figure 13), allowing users to compare the effect of differences in design decisions (e.g., linear vs open-world games), specific variations in canonically different games (e.g., differently localized releases of the same title) or differences in play styles that vary between person (e.g., novice vs. expert) and purpose (e.g.,

speedrunner vs. casual play). This affordance can be used to set up cross-domain analogies, subject to the limitations of the embedding model.

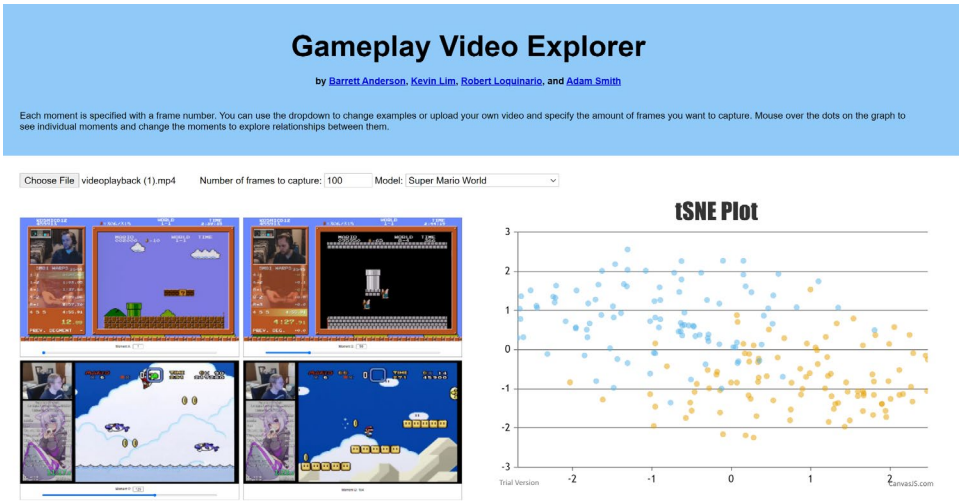

Figure 13. The Gameplay Video Explorer site, with data from two videos loaded (a speedrun of Super Mario World, and of Super Mario Bros.). Moments from each video are mapped to the same space, and color-coded by source. Mouseover reveals the specific moment. In order to handle diverse video lengths and still remain performant, a sampling approach is taken, collecting a constant number of approximately equally distributed screenshots. Previous experiments with variations in the number of displayed screenshots for the search-by-analogy explainer demonstrated that the same overall impression could be conveyed with even a small subsample.

Extending the *map* analogy for exploring videogame moment space, while users do scroll and zoom through digital representations of physical space their initial map queries are usually textual. Natural language text queries for moments in game space have been demonstrated (Zhang & Smith, 2019), see Figure 14. Incorporating a form of text search would be a good extension of this work, overcoming a weakness of another game search system (Ryan et al., 2017) and responding to one of the most

common interests expressed in my requirements analysis interviews described in Chapter 3.

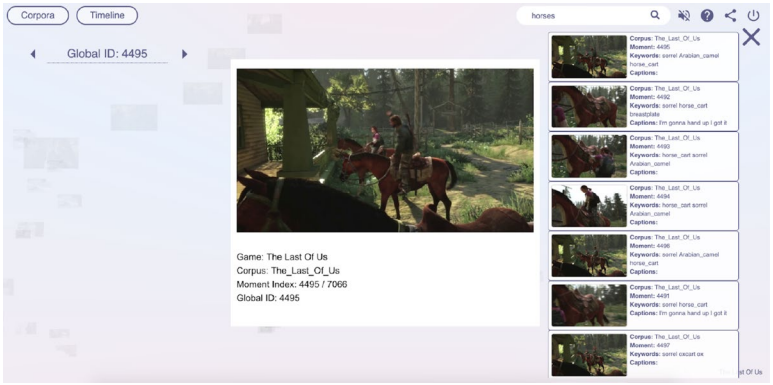

Figure 14. Natural language search for videogame moments in a modern console game. In this example the query "horses" returns moments containing horses from the game The Last of Us(McIntosh & Gregory, 2013). Reproduced from Zhang & Smith, 2019.

Quantitative Evaluation

For the purposes of this technical validation, I am focusing on games where predictive models support our search system. In order to provide a quantitative evaluation of Gameplay Video Search, it was necessary to establish a ground truth that the system's output could be compared against. This ground truth would typically take the form of a canonical set of queries and their expected results. Several queries that could be used for this purpose were provided by early test users of the analogy search system. Because the corpus of gameplay videos contained over 10,000 individual screenshots, instead of coding each one individually as correct or incorrect for each

query, I selected a set of features to check a subset of the screenshots for. These features could be visually identified from a screenshot, reflected aspects of the intended analogy implied by the queries, and partially overlapped between different games.

Query 1	A	B	C	
ID	410	1858	1020	Ideal
Screenshot				
URL	http://barrettrees.com/	http://barrettrees.com	http://barrettrees.com	N/A
Screen	Level	Level	Level	Level
World	1-1	1-1	1-2	1-2
Powers	Small	Fire	Small	Fire
Yoshi	No	Riding	No	Riding
Facing	Right	Right	Right	Right
Jumping	No	No	No	No
Item	None	Mushroom	None	Mushroom
Lives	5	6	5	6
Dragon Coins	0	3	0	3
Stars	0	10	0	10
Time	267	226	364	400
Coins	0	33	3	36
Score	0	41430	11100	52530

Query 1	A	B	C	
ID	2005	4919	2267	Ideal
Screenshot				
URL	http://barrettrees	http://barrettrees	http://barrettrees	N/A
Screen	Level	Level	Level	Level
Area	Crateria Surface	Crateria Surface	Crateria	Crateria
Morph Ball	No	No	No	No
Facing	Left	Left	Left	Left
Jumping	No	No	No	No
Missile Upgrade	No	Yes	No	Yes
Missiles		10		10
Energy	99	69	99	69

Figure 15. Selected gameplay screenshots and features for Super Mario World and Super Metroid analogy queries and ideal results. The ideal was generated by completing the A is to B as C is to __ analogy. Results were then scored compared to the ideal, with each feature in the list weighted equally.

These features also differ from what the embedding was trained on (screenshots and the contents of memory), but allow me to extrapolate an expected ideal result which the moments returned by the search system can be compared against. This criteria makes it possible to quantitatively evaluate the search system against a transparent and explainable benchmark. For the purposes of this evaluation, I compared the results of search-by-analogy against search results without the analogical component, for six queries (4 for *Super Mario World* and 2 for *Super Metroid*). The ideal result was generated by completing the analogy implied by the query screenshots, and all of the listed features were given equal weight in relevance judgments (See Figure 15). Numerical values for the ideal were generated by adding the difference from A and B to C, capping at the games' maximum and minimum values for the quantity (e.g. for

Super Mario World time ranged from 0 to 400 seconds and Dragon Coins from 0 to 3). For the relevance score difference from the ideal was generated as a percentage of this range (i.e., a result could be partially relevant depending on how many features matched).

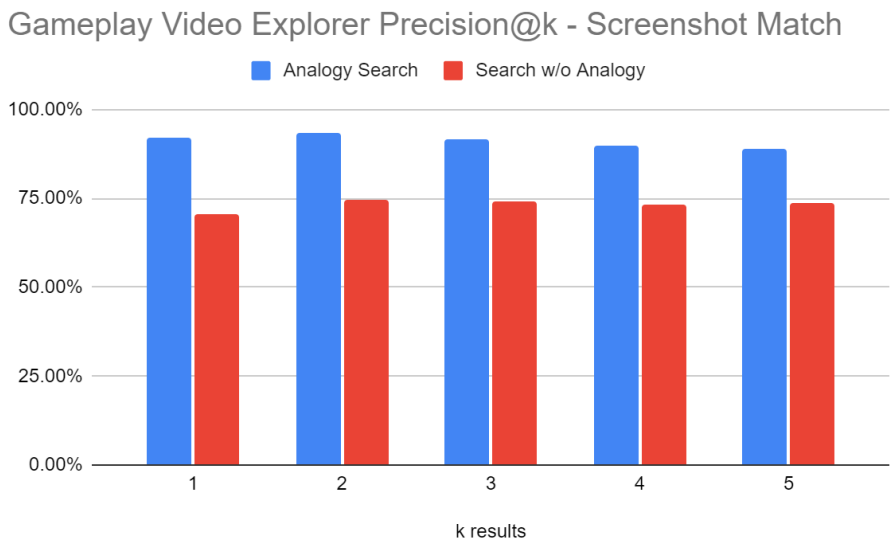

Figure 16. Precision for Gameplay Video Explorer search, comparing screenshot match for search-by-analogy to screenshot match for search without the analogy component.

Overall precision was better for search-by-analogy, compared to search without the analogical component, by around 17%. This is a quantifiable benefit of adding the analogical component to search, compared to simple moment similarity (See Figure 16).

Another way of quantifying precision as described above is normalized discounted cumulative gain (nDCG). This measure is based on a graded relevance

approach instead of binary judgement, and assumes that relevant documents are more useful earlier in the search results list (Manning et al., 2008, Chapter 12). Normalization allows for different lengths of search results to be compared, although for the present comparisons this measure is modified by capping these at the same length (5) due the human effort of making relevance judgments for each game screenshot (See Table 3. Normalized discounted cumulative gain (nDCG) for the Gameplay Video Explorer, including all example analogical search queries when the analogy component was included.). Similarly, the ideal DCG scores were generated by ordering all feature coded query responses and may have omitted better matches in the corpus. This kind of restriction to a sample for evaluation purposes would apply to any search engine evaluating a sufficiently large corpus (e.g., web search cannot be validated against the entire web). Even with this limitation, the Gameplay Video Explorer search results did not always match the ideal, indicating that there is room to improve the embedding or the search method.

	Ideal DCG	DCG (Analogy)	DCG (Non-analogy)	**nDCG (Analogy)/ Ideal**	**nDCG (Non-analogy)/Ideal**
Super Mario World Query 1	2.85	2.83	1.92	99.32%	67.46%
Super Mario World Query 2	2.86	2.86	2.67	100.00%	93.37%
Super Mario World Query 3	2.66	2.48	2.18	92.93%	81.67%
Super Mario World Query 4	2.59	2.55	2.07	98.67%	79.88%
Super Metroid Query 1	2.87	2.87	2.10	100.00%	73.06%
Super Metroid Query 2	2.44	2.23	2.05	91.61%	84.32%
Average	2.71	2.64	2.16	97.09%	79.96%

Table 3. Normalized discounted cumulative gain (nDCG) for the Gameplay Video Explorer, including all example analogical search queries when the analogy component was included.

The primary comparison here is between the average nDCG scores for the analogy search and search without the analogy component. This is the equivalent of comparing trying to fill in the "A is to B as C is to __" analogy to a search only for

moments adjacent to "C." These findings provide quantifiable confirmation that the analogical component was an important contributor to the search results, that search with analogy beats simply searching for adjacent moments. This is a demonstration of the precision for Gameplay Video Explorer search by analogy. This measure does not capture the possibility that better matches for these moments exist in the corpus that this system did not find. To do so would require a measure of recall (e.g. MRR), but providing such a measure would require characterizing every moment in the search corpus. If this task could be automated, the method of doing so itself would necessarily provide an even better search system.

Future Work

The ability to combine multiple video sources and to swap between underlying models suggests several additional possibilities. For example, the visualization might be useful for distinguishing between different styles of play for the same game (e.g., comparing two different players, or comparing a speedrun to a completionist approach). This might even surface a personal "style" unique to individual players, that persists across games. Replacing the underlying game-based embedding with an image trained embedding might enable comparisons between moments in games with photo-realistic graphics without the need for training game-specific models. Even for games that do not attempt photo-realism, an image-based embedding could be more hardware agnostic and generalizable. This avoids the expense and complexity of training a large variety of game-based embeddings, and this might even be useful for a visual search-

by-analogy or other comparisons using non-game videos. There is already work in this domain to evaluate such a contribution against (Peyre et al., 2019). An image based embedding, or possibly an embedding trained with moments from a substantial variety of different games, might be necessary to enable interesting comparisons of moments across games.

Contributions

This work makes a contribution to information visualization, in the form of an interaction model for video exploration, including applications to a new domain. The gameplay video can also be understood as a kind of personal information, connecting to the field of personal information management. For speedrunners and streamers these are recordings of personal performances, and they have interests similar to those of a sports player (or entertainer) who might record, review, evaluate, and selectively share video of their own performances. The tool itself may also be of direct use in game studies for making arguments about the differences between specific extended play sequences, or exploring relationships through cross-domain analogies.

Chapter 5. Undergraduate Games Corpus

Project Summary

In this chapter, I describe a new corpus of undergraduate student games, discuss present and possible future uses of such a resource, and document the process of its creation at a level of detail intended to support further contributions to the collection. Portions of this chapter have been adapted from work that is currently in press (B. R. Anderson & Smith, 2021). The co-author listed in this publication directed and supervised the research which forms the basis for this dissertation chapter. This builds on prior work within my own institution archiving student games and contributes documentation to promote continuity and improve the quality of the games collection process. To avoid the ethical and practical risks of pre-emptive data collection (Prabhu & Birhane, 2020), and particularly the dangers of implementing something akin to classroom surveillance by blindly collecting student data, contributors were actively informed about the nature and purpose of the game collection project and participated in contributing their own work to the corpus. By building and indexing this collection I am creating a knowledge base that is compatible with the tools I am creating, in ways that can enable some of the use cases identified in the requirements analysis reported above, and that I hope can be leveraged to go beyond these applications. This work has resulted in both a growing corpus of student games and a documented and replicable process to continue developing this collection.

Prior Work

What constitutes an archive has been a contested term in the context of digital humanities (Owens, 2014; Theimer, 2012). I have referred to this work as a student games archiving project, but since corpus is not necessarily an archive that requires some qualification. The traditional concept of an archive is essentially the permanent records of a person or institution, which continue to have value to the creating agency as well as to other potential users (Millar, 2009). In the process of developing a system capable of understanding certain types of link-based games at a systems level (see Chapter 6), one of the necessary steps was the collection of a corpus of student games, and I as I took on more of this process I also took the opportunity to attempt adding additional structured information about these games. Existing collections of student games, including ours, are similar to the archival studies concept described above, in that they contain original documents that have not been published elsewhere, and that those documents form a coherent collection by virtue of their shared link to a specific institution.

Student Writing Corpora

Outside of games research there have been various academic efforts aimed at collecting and analyzing corpora of student writing (Megyesi et al., 2016; Nesi et al., 2004), including writing collected from standardized tests, or directly from student submissions. The former explicitly describes this work as an effort in digital

humanities. It has been argued that plagiarism detection services (e.g., Turnitn) are a form of archiving as well, since they necessarily contain and preserve a substantial body of student work (Purdy, 2009). However, public access to the majority of their contents is unlikely because making this available would be counter to their competitive and commercial purpose. The public availability of archives is important not only for researchers, but for the students themselves, who should not be excluded from analyzing the work of their peers. The opportunity to do analysis at the level of an entire corpus, and to develop an understanding of the computational tools necessary to do this kind of analysis, should be part of a modern liberal arts education, and I expect that students will be more motivated to do so if these corpora are personally relevant. Corpora, including those of student writing referenced above, could also be useful for novice educators learning to provide constructive responses to diverse student voices. Games archives have the potential to serve a similar function in teacher training for games education.

Games Corpora

Games research has benefited from the contribution of several corpora. AI research efforts to build systems that can interact competently with a variety of games (Pell, 1994), have been supported by supplementary game collections beginning with the dissertation that first posed the general game playing problem (Pell, 1993). The creation of several specialized game description languages to support this work (Ebner et al., 2013; N. Love et al., 2006; Thielscher, 2010) was accompanied by collections of

game formalizations, including some inspired by classic videogames (Perez-Liebana et al., 2016). More recently, work on automatically constructing formal representations of interactive narrative (Partlan et al., 2019) was based on a collection of 20 student-authored research games (Partlan et al., 2018) built on the StudyCrafter (Harteveld et al., 2016) platform. In another sense, collections of games available commercially (e.g., through app stores, Steam, etc.) can be considered another type of games corpora, though one that is generally less relevant for research purposes due to cost and technical barriers.

Student Games Corpora

Having considered corpora of student work and corpora of games, I now turn to the smaller set of corpora that combines both. For student work specifically in the form of games, the University of California's eScholarship archive[33] has also been used for the preservation of student games since 2014. This work began in UC Santa Cruz's Center for Games and Playable Media, as an offshoot of the GAMECIP project.

Even more recently, the academic library of the University of Utah created a student game archive (Casucci, 2018). This happened in part because a student game developed in the early days of their program achieved a level of notoriety, after being featured by a popular streamer (*The Loss of "Erie": EAE Thesis Games Archive*, 2018). This led to the discovery that the game's source files were lost which served as a

[33] https://escholarship.org/uc/ucsc_games/

catalyst inspiring the development of an archive to prevent such losses in the future. Now, this archive contains many student games developed at the University of Utah as a part of their own games program. Moving focus from student writing to student games, I shift from analyzing traditional texts using digital tools to the analysis of inherently digital expressions.

While it is not entirely student work, the recently assembled corpus of games created for the Graphic Adventure Creator (GAC) also contains a large variety of examples of novice work (Aycock & Biittner, 2020). The paper reporting this collection examines the diversity of this collection from a technical perspective, which could be supplemented with additional social and cultural analysis.

Critical Information Theory

In her book *Algorithms of Oppression* (2018) Safiya Noble challenges the idea that the algorithms that guide search are somehow benign or neutral, providing many reminders that these processes are always the result of human decisions, and that they reflect the biases of their designers. Acknowledging these biases is especially important for any work that might generate a false impression of objectivity. Bringing this critical approach to search results she writes, "In practice, the higher a web page is ranked, the more it is trusted. Unlike the vetting of journalists and librarians, who are entrusted to fact check and curate information for the public according to professional codes of ethics, the legitimacy of websites' ranking and credibility is simply taken for granted.

The take-home message is that, when it comes to online commercial search engines, it is no longer enough to simply share news and education on the web; we must ask ourselves how the things we want to share are found and how the things we find have appeared." (Noble, 2020, p.155). While acknowledging that this is not primarily a technical problem, in response to requests for a technical solution the author describes an alternate way of searching across websites. She suggests throwing away the traditional ranking (sorting) paradigm of information retrieval in favor of something closer to map browsing. In this way, users are directly exposed to the diversity of available results, and there is no need to make the fraught decision about which top result is going to represent a billion potential candidates. One challenge here is that, while there are a wealth of methods for evaluating ranking-based information retrieval systems (including those used in this dissertation), no equivalent obvious technical evaluation methods exist for evaluating a map/visualization based retrieval system. It will be important to explore these possibilities in order to escape from systems that make all of the search decisions for the user (e.g., making the "I'm feeling lucky" button feel like almost all you need), and therefore guide users towards results that (subtly or not) primarily reflect the interests of the search system provider. An alternative worth imagining is offering search tools that instead of infantilizing or exploiting users, amplify existing human cognitive abilities, empowering them to skillfully perform search tasks.

Data Collection Ethics

The ethical considerations of building a dataset such as the corpus described in this chapter should not be overlooked (Leidner & Plachouras, 2017). The archive should neither misleadingly erase the labor of people who contributed to it via participation-washing (Sloane et al., 2020), nor problematically represent people or work that was never intended to be used in this kind of research (Prabhu & Birhane, 2020). The construction details of the Undergraduate Games Corpus, described in the section on corpus construction below, offers one way of navigating these constraints. It is also important to note that this responsibility to those who contributed to this work does not end with data collection, but continues into the ongoing management of this resource.

Corpus Construction

In order to generate this archival resource, I developed a process that included both elements of presentation and interaction with game design students, and a technical pipeline to prepare their work for archiving. The first part of this process was an effort to recruit student participation. This meant not only asking students to submit their own games for archival purposes, but demonstrating to them the technical games research that their submissions made possible, and inviting them to participate in that work. To do so, I gave guest lectures about technical games research to game design students. These lectures were adapted to the level of the class, but included presenting both interactive artifacts, and research opportunities.

Soliciting the games that students were creating for this course for the corpus proceeded in two phases that were incorporated into the course. In the first phase, the students were given a lecture regarding the emerging landscape of technical games research (Nelson, 2020). This lecture incorporated examples of AI projects enabled by data, such as the GameSage game recommendation engine (Ryan et al., 2016) and the GameSpace exploration system (Ryan et al., 2017), and invited them to contribute their own games towards similar research projects in the future. Students specifically saw examples of how individual games could be plucked out of a large space for inspection using simple text queries. In the second phase, the students completed a form indicating their willingness to have their own games included in the corpus. If they agreed, they were also asked to provide a game description, provide appropriate keywords/tags, and if they opted for a Creative Commons license make some decisions about how they would like their game to be distributed. Responding to this form, even only to indicate that they did not wish for their games to be added, was part of the students' participation grade for the course. Ethically it would be more optimal to make this opt-in instead of opt-out, but the trade off of doing so in a subsequent quarter (Summer 2020) resulted in no games being contributed to the corpus.

Game Collection Process Guidance

Building this archive has resulted in the production of a detailed step-by-step guide for future game educators (See Appendix - Undergraduate Game Collection Guide). This work is focused on the current incarnation of these courses, targeting the

specific kinds of games students are developing. Included in this guide is example code, currently usable without modification, to transform data obtained directly from current courseware systems and student questionnaire responses into a format that can easily be ingested by the current archiving system. All of these pieces are open to inspection, and can be modified or expanded upon as necessary to adapt archiving efforts to changes in the structure of assignments or systems.

Adding to the UC eScholarship Archive

As part of my work to develop a system for search across games, I collected student work from introductory and advanced game design classes. During the course of this process it became clear to me that this could be part of a resource for other scholars. The initial work to make this possible, including student games in the UC eScholarship archive, had already begun, which meant there was an opportunity to add the games I was collecting. I solicited student games for contribution to the UC Santa Cruz CGPM archive in the Summer quarter of 2018. In the following Winter and Spring quarters I solicited approximately 600 additional student Twine games (and some smaller number of graphical games) to be added to the archive. This work involves linking student work, permissions granted, and metadata submitted. Improvements over several iterations of the game collection process have led to the inclusion of abstracts, tags, and other relevant metadata from student game makers.

Corpus Characterization

The Undergraduate Games Corpus consists of 755 individual games. These are complete, playable games, created by student authors in introductory game design courses at UC Santa Cruz between July 2019 and July 2020. In many cases, these games reflect student's lives (e.g., their interests and personal struggles) and the trends of the cultures around them (e.g. which pre-existing games they may choose to clone). The overlap of this specific student population and these specific dates strongly shaped the themes of many of the games (e.g., relating the experience of local wildfires, power outages, labor strikes, and the global emergence of the COVID-19 pandemic). Students conveyed these themes in methods specific to interactive media, potentially employing techniques like the procedural rhetoric of failure (Treanor & Mateas, 2009), a topic covered in course lectures. Students could, for example, convey inevitability by having multiple choices within their game lead to the same undesirable outcome. The subtle topic of anxiety (unsurprisingly common in student games) could be conveyed by embedding many dead-end choices in a story, allowing the player to interactively worry along with the protagonist as they see many different outcomes play out in doom.

Many student authors provided short game descriptions, which are included for about 85% of the games in the collection. Similarly, about 70% of games were supplemented with descriptive tags. The histogram in Figure 17 visualizes the distribution of word counts in descriptions while the word cloud in Figure 18 characterizes the frequency of common descriptive tags.

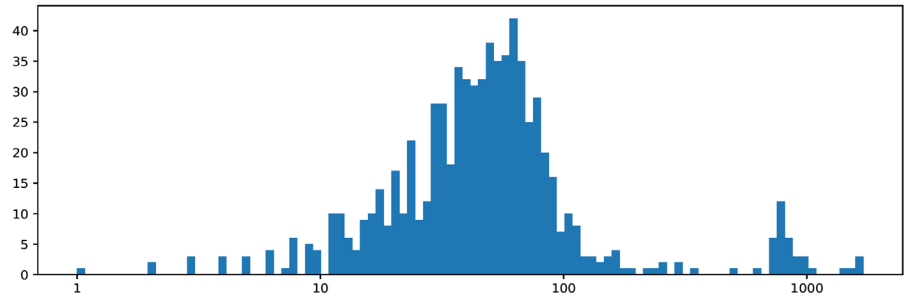

Figure 17. Game description word counts for games in the Undergraduate Games Corpus, the majority were under 100 words.

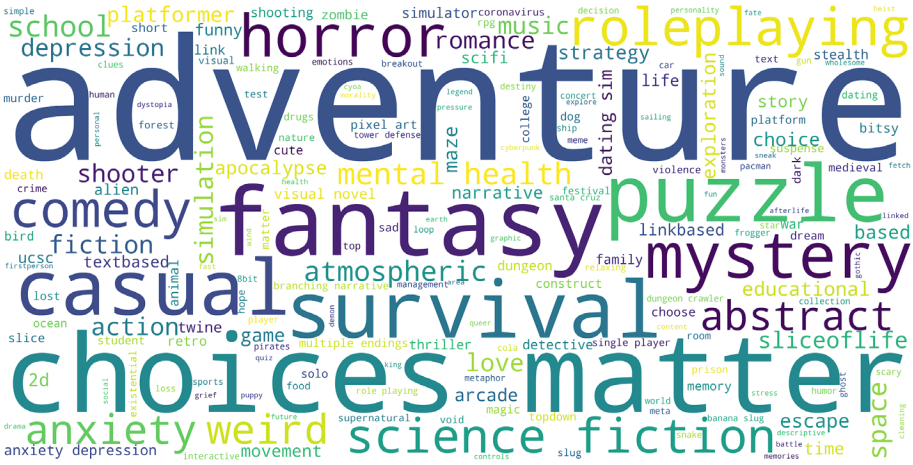

Figure 18. Game tags word cloud; size reflects frequency. Tags were manually assigned by the student authors.

Some of our games are closer to being represented in well understood formats than others. The majority of games in the corpus (about 61%) are primarily textual, created in the Twine[34] game engine. Twine games (more often called stories) are typically made up of several passages that each represent individual scenes in the story.

[34] http://www.twinery.org

While student games made extensive use of hypertext links to present the player (interactive reader) with choices, some students experimented with using the engine's scripting language to generate player choices procedurally (modeling combinatorial spaces that would be unreasonable to express with a manually constructed network of hyperlinks). Almost all Twine games make some use of dynamic scripting logic, even if it is only to slightly alter the text for passages once they have been seen by the player once (e.g., to omit the verbose description of an object once it has been introduced).

The corpus also includes graphical games. About 11% of the corpus was created using Bitsy,[35] an authoring tool for "little" games that exposes a very tightly scoped scripting language (whereas Twine authors can draw on all of JavaScript when needed). About 25% of the corpus was created with Construct 3,[36] a relatively flexible tool for making games with two-dimensional graphics that run inside web browsers. The remaining 3% primarily build on the Godot[37] game engine, an open source tool comparable to Unity[38] in terms of flexibility and support for advanced, three-dimensional graphics rendering. From smaller Twine stories to larger Godot games, the artifacts in our corpus have a wide distribution in weight and complexity. Figure 19 characterizes this diversity in terms of total project file size, which might also be read as a crude proxy for technical complexity.

[35] https://ledoux.itch.io/bitsy
[36] https://www.construct.net/en
[37] https://godotengine.org
[38] https://unity.com

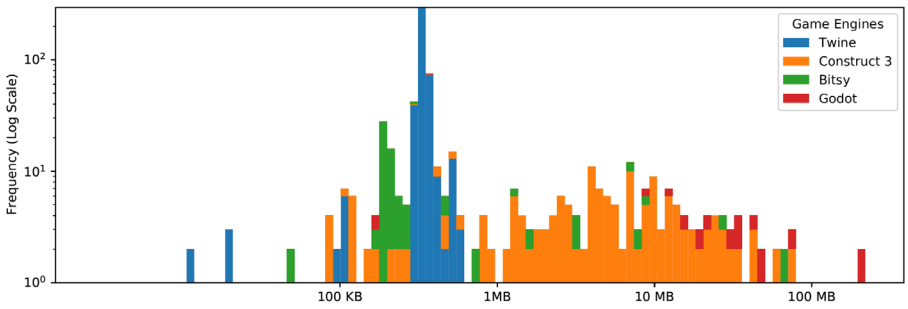

Figure 19. Distribution of game project sizes (including source code and assets), also representing the variety of game types by authoring tool (log-log plot).

Focusing specifically on the largest segment of our corpus, Twine games, Figure 20 characterizes the distribution of sizes of Twine games in terms of number of passages per story and number of words per passage within each story. Defining words simply as the number of whitespace-separated tokens, this intentionally includes words that are not directly seen by the player such as those contributing to scripting logic or to the visual formatting of the text (or even references to image files). All of these words represent effort by the author to shape the audience's perceptual experience.

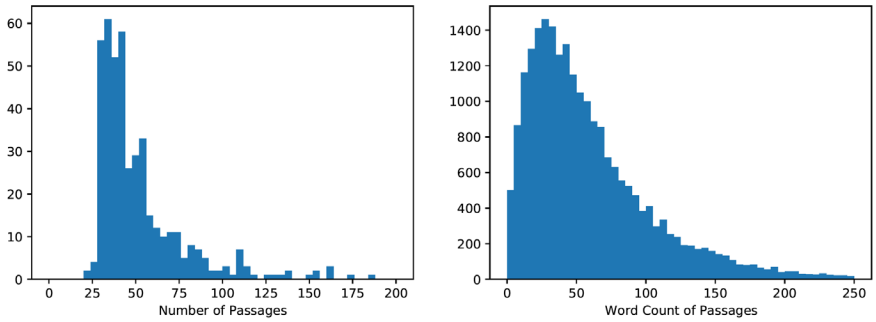

Figure 20. Scale of game for just the Twine stories, representing the number of scenes (passages) in a story and the detail within each scene (words and source code tokens per passage) for all passages in all stories.

Every game in the corpus is attributed to one or more specific student authors under their chosen names, and all are available for public consumption (while respecting the author's copyright). Further, many (about 54%) of these games are provided under a Creative Commons license that grants further rights,[39] such as to remix and transform the work, under the condition that appropriate credit is given to the original authors. Table 4 summarizes the distribution of chosen licenses.

License Type	Number of games
Copyright Only	342
CC-BY	219
CC-BY-SA	15
CC-BY-NC	73
CC-BY-NC-SA	4
CC-BY-NC-ND	2

Table 4. Use of Creative Commons licenses for game project files.

A recent paper on videogame text corpora (which would not capture how that text was directly linked to player choices) offered a list of desired characteristics for evaluating corpus quality (van Stegeren & Theune, 2020). Among their criteria was a

[39] https://creativecommons.org/licenses/

concern for *representativeness*, that the dataset represent the work of professional videogame writers and be sourced from well-known games that have a substantial user base. Our corpus instead strives for a sense of *authenticity* in the sense of having the dataset represent work intended as human experiences even if not made by professionals. Their concern for *diversity* suggested that the dataset reflect the variety of types of text occurring in videogames (e.g. dialog, tutorial, character names, etc.). By contrast, our concern for *diversity* considers the variety of subject matter, and a population of student authors that is more diverse than the population of professional game authors.

Corpus Derived Projects

The undergraduate games corpus described in this chapter has already been leveraged for several derivative projects. This includes additions and improvements to the eScholarship system, a gallery of student games (See Figure 21), and a project exploring retrieval for interactive narrative. The 755 games collected here have been added to the UCSC eScholarship archive. This brings the UC eSCholarship archive total to more than 1000 student games. This includes link-based and graphical games, and games developed in a variety of engines such as Twine, Construct, Phaser.io, Bitsy, Game Maker, Godot, and other engines. The majority of these games are now playable online, either via download or in-browser directly in the archive. Adding these games to the archive and incorporating that archive into a subsequent course also led to the improvements to the eScholarship system itself, which made online play of HTML

based games possible. A subset of these student games have also been showcased in a web gallery. This was originally created to facilitate awards selection for outstanding work, and has subsequently been modified to include only those nominated games that were also submitted to the corpus. The Twine games in this corpus also became the basis for a project exploring retrieval for interactive narrative, described in more detail in the following chapter. In addition to these derivative projects, the corpus and its metadata are available to facilitate unanticipated explorations.

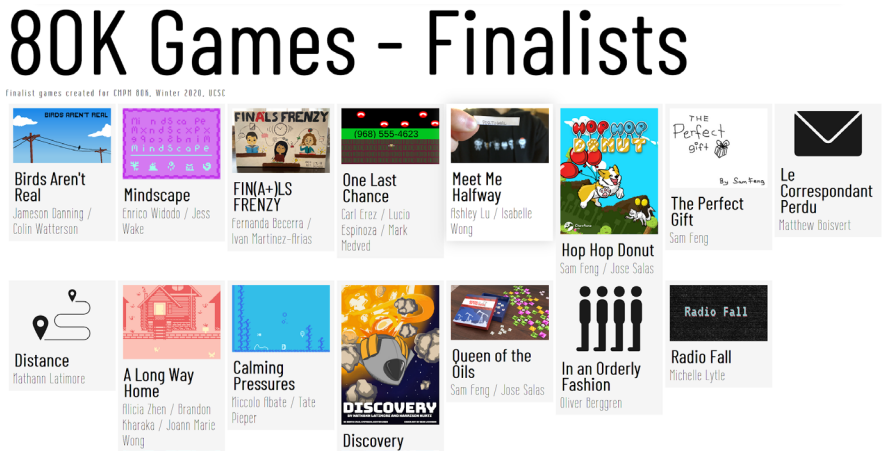

Figure 21. Archived gallery website for selected games from the CMPM 80K Winter 2020 course. All finalist Twine games (5) were sample queries for the quantitative evaluation of the Story Navigator system. The current version of this gallery includes only those games submitted to the corpus.

Future Work

A unique feature of this collection of student games is the inclusion of many early drafts alongside completed projects. This could allow for research on game development, exploring what Khald et al. (2018) refers to as the trajectory in design

space. The approach enabled by this corpus would complement theirs, offering less depth on the development of individual games in exchange for the potential to explore a broader swath of novice approaches.

One of the primary challenges of the eScholarship incarnation of this archive is searchability. While it now contains over 1000 games, the tools necessary to filter and search for any specific game of interest are very limited, because they were not created with games in mind. The archive is also not amenable to browsing in the way that resources like game stores are, because it cannot effectively make use of the tags and other metadata that student game makers have provided. This archive was created primarily to be a resource for scholarly papers, and some adaptations will be required in order for it to become an effective resource for interactive experiences, or for performing corpus level analysis. Fortunately, even if implementing these changes within the existing archive is difficult, because it is a free and open resource researchers could potentially build more effective systems outside of it. One such system, a search tool built on a subset of student interactive fiction, is described in the following chapter.

Chapter 6. Story Navigator: Retrieval Technology for Interactive Fiction

Project Summary

This chapter describes the development of a search tool that operates on the content of text-based narrative games, and the outcomes of several retrieval experiments to evaluate the effectiveness of this system, and a brief exploration of how the interactive structure of these stories could be usefully represented for search. This is an application building directly on the Twine interactive fiction subset of the student games corpus described in the previous chapter. Where the earlier work focused on searching for moments *within* games (Zhan et al., 2019; Zhang, 2019; Zhang et al., 2018), with the undergraduate games corpus described in the preceding chapter in hand we can now also explore searching for and comparing moments *across* games. The primary outcome of this work is the **Story Navigator** site, which can be used to explore a space of narrative games, and find the position of a given game (whether a well known interactive story or a student work-in-progress) relative to the others that have been mapped into the space. As in the Gameplay Video Explorer project described above, moments (here operationalized as the individual passages of a Twine story) have been converted into vectors and mapped into a navigable visual representation of an abstract space. Rather than searching across games based on their metadata or what is written about them in secondary sources, the system described in this chapter will form a representation of a game based on an aggregation of all of the moments that the game is known to contain.

Prior Work

Interactive Fiction

Interactive fiction includes works that present different texts or can be read in different ways depending on the actions of the reader (Montfort, 2005). This kind of interactivity predates and does not require digital technology. Borges' short story *The Garden of Forking Paths* (1941) containing a description of such a changing novel, and Raymond Queneau's *Yours For the Telling* (1973) provides a model for modern Choose-Your-Own-Adventure style texts (e.g. *The Cave of Time* (Packard, 1979), *Romeo and/or Juliet: A Chooseable-Path Adventure* (North, 2016), etc.). These possibilities have been explored in a variety of electronic works, from the parser-based text adventures of Infocom (e.g., *Zork* (T. Anderson et al., 1980), *Hitchiker's Guide to the Galaxy* (Adams & Meretzky, 1984), etc.), to hypertext choice-based games (e.g., *Howling Dogs* (Porpentine, 2012), *Even Cowgirls Bleed* (C. Love, 2013), etc.) that were created in the context of a recent queer art scene (Kopas et al., 2015). The focus of this project is on the latter kind of work, for technological and pedagogical reasons. From a technological point of view, focusing on primarily text-based games permits leveraging advanced text analysis technology to improve the nuance of our content analysis. Making this kind of choice-based text game is also good for beginning game design students, because it requires limited programming experience and can keep the focus of their efforts on exploring the essentials of interactivity.

Story Content Analysis

For the purpose of this system I am using story content to refer to the human-readable text of a link-based story. This is distinct from story structure, which in this context describes the properties of the story that remain when all of the actual text content has been abstracted away (e.g., a directed graph, variable states, etc.). By focusing on the content of student Twine games, I was able leverage the Universal Sentence Encoder (Cer et al., 2018), part of an existing ecosystem of tools designed for text analysis (e.g. BERT (Devlin et al., 2019) , ELMo (Peters et al., 2018), etc.), because these games are primarily text-based. This enables a simple content-based retrieval strategy that operates over the text seen by the player of the story. While it builds on a neural representation of text derived from deep machine learning, no embedding models were trained on the current games corpus for this project. I used USE v4[40] a specific USE model that had been pre-trained on multiple sources (Wikipedia, web news, web question-answer pages and discussion forums).

Story Navigator System

To begin this work, I extracted individual passage text for Twine games from the Undergraduate Games Corpus (see Chapter 5), as well as all outgoing links, variables, and passage titles. A few initial experiments with word cloud visualizations suggest that this might be a useful paradigm to incorporate into further visualizations

[40] https://tfhub.dev/google/universal-sentence-encoder/4

(See Figure 22). This requires trimming passages down to the player-visible text, and separating this from other elements (e.g. scripting, links, images, etc.).

Single Passage Entire Corpus

Figure 22. Word cloud illustrations of a single passage and of the entire student games corpus so far.

Using BERT embeddings (Peters et al., 2018) I created vectors that reflect latent semantic information (e.g., passage or story topics that are not explicitly stated in the text, See Figure 23). For a more interpretable visualization, I selected just three Twine games to embed passages from into a 3D space (See Figure 24). This made it easier to compare individual passages, drawn from different Twines. Combining text from all passages before putting them through the USE generates embeddings at the level of the entire Twine games. These whole game embeddings can be searched via a short text string, as in a typical search engine, or used to find games that are similar to each based on their entire textual content (See Table 5). In either case, the system is now dealing with the game's internal content.

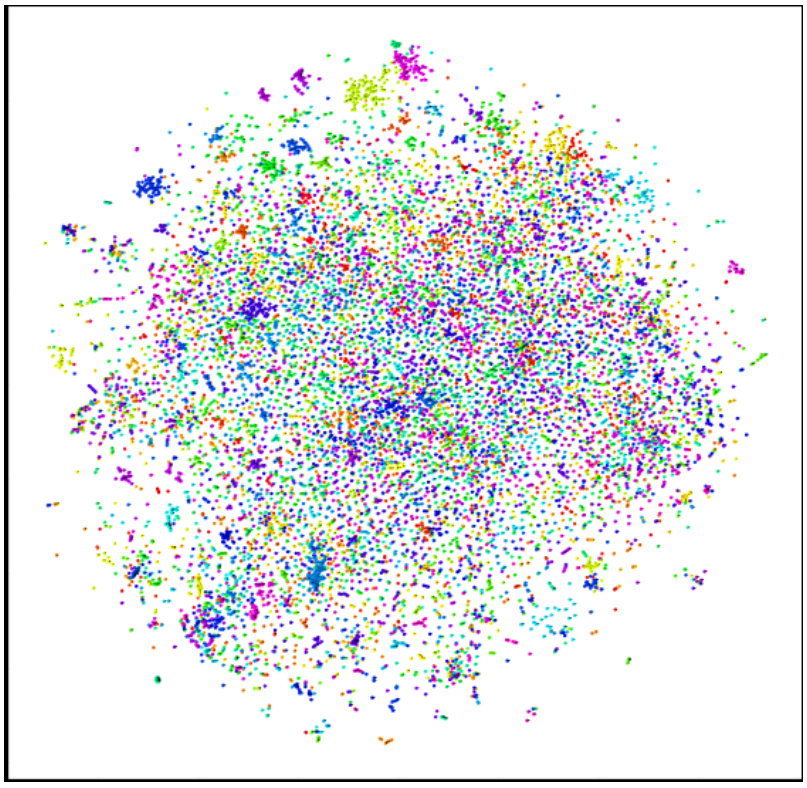

Figure 23. Passages from 317 student Twine games converted to vectors (using BERT) and plotted in two dimensions (using t-SNE). All passages from a single story are alike in color.

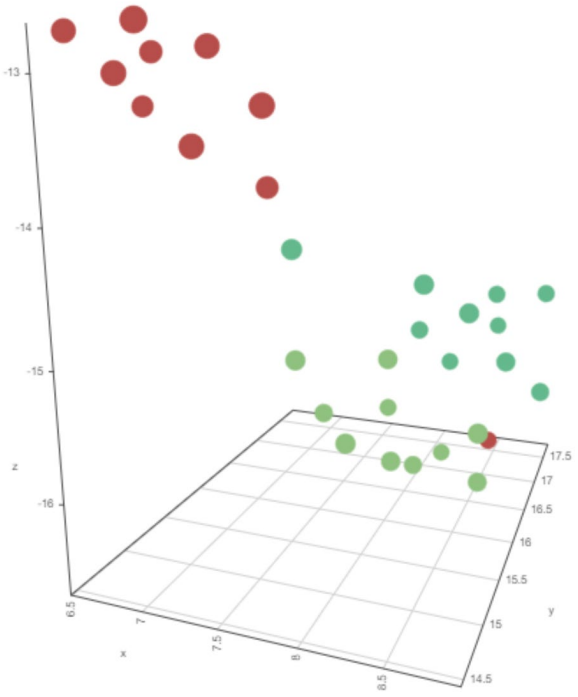

Figure 24. Passages from three Twine stories embedded via the Universal Sentence Encoder , projected into a 3D space (via t-SNE). Each point is a single passage (text revealed on mouseover), color coded by the source story.

Query Game	Procrastinate. by Jameson Danning Procrastinate is an interactive fiction game that involves the player doing things around the house to distract them from the fact they have homework. You can make food, watch TV, and many other activities that are more fun than homework. The game is just a slice of life that is made to feel relatable, and hopefully many players can identify with it.	Totally Accurate Cooking Simulator by Justin Beedle Totally Accurate Cooking Simulator challenges you to prepare the ultimate dish. Will you create a masterpiece that will reign above all? Or will you find yourself creating a disastrous kitchen nightmare? Explore your options in the kitchen and traverse through 20+ strange endings.
Results List	Cramming by Sherry Luo Cramming by Sherry Luo The game is an interactive fiction where you play as a high school student cramming for an upcoming exam. This game aims to simulate the stress of the students and the annoying thoughts that might go through their mind as they are trying to study. ---------- home by Jozi Coates home is a game about being alone. You explore your apartment and interact with the things and people around you. This game aims to show people a peek into my life in isolation and illustrate what can only happen once. ---------- CYA-40 v1.0 by Emily Applewhite CYA-40 v1.0 by Emily Applewhite CYA-40 v1.0 is a lighthearted game in which you sit down to an old machine in order to read a Choose Your Own Adventure novel. However, as you'll come to find, the machine is broken and you'll have to help fix it in order to get it back up and running. However, as you'll come to find, there's more that this thing can do than just run Choose Your Own Adventure novels...	A Trip to the Grocery Store by Barry Day A stoner-comedy adventure, in which you find yourself traveling to the grocery store after suffering from a severe case of the munchies. What seems like a simple journey quickly turns into a wacky fantasy, complete with memes and gnomes. Will you seize the best chips and the finest cereal? Or will you fail at the hands of cops, grandmas, and gnomes? Did I mention gnomes? ---------- Chef's Special by Matthew Alvarez You are the head chef in a disaster of a situation, can you save the restaurant? ---------- Rice Quest by Andre Domingo Rice Quest is a branching narrative focused on Rice. ----------

Table 5. The Story Navigator system displaying example results lists of related games, with apparent themes of stress/isolation (left) and food/cooking (right). In each case the topmost game in the list was the query, and subsequent games are ordered by proximity in the space. A student could explore this list to find games similar to their own work-in-progress.

Quantitative Evaluation

Quantitative evaluation of the Story Navigator was conducted by adapting the nDCG precision metric. The format and structure of this evaluation is similar to that described in the chapter on the Gameplay Explorer system above. It differs in how human relevance judgments are made, because complete Twine games have a different kind of structure than SNES game screenshots, or moments. In order to perform this evaluation, it was necessary to generate several queries, and to establish a ground truth by providing human relevance judgments for each of the results returned by the system. For queries, I chose to use the subset of Twine games from the corpus that had also

been selected as finalists[41] for submission to the UCSC Games Showcase[42] by the undergraduate Reader/Tutors who were part of the instructional staff of the game design course that the corpus was sourced from. This peer selection represents games that are generally of higher quality, and for the purpose of generating realistic queries it also includes games that are likely to be of interest to students.

Relevance judgements were determined by a pair of human reviewers. This determination was made by the reviewers who played each game, viewed the games in the Twine editor, assigned keywords relevant to the game's content (e.g., Magical Realism, Second Person, Dream, Isolation, etc.) and structure (e.g., Randomization, Key Items, Multiple Endings, etc.), and discussed any areas of disagreement until consensus was achieved. Ultimately each reviewer provided a binary judgment of each game's relevance as a result for the query game. The categorization component of this procedure was adapted from grounded theory methods for coding qualitative data (Charmaz, 2006; Glaser, 1992), abbreviated and focused specifically on the relevance judgments needed to evaluate a retrieval system (See Figure 25). To provide an initial quantitative evaluation of this system, relevance judgments for results returned by the Story Navigator system were compared to baseline relevance judgments for an equal number of Twine games randomly selected from the corpus.

[41] https://barrettrees.com/80k-Games-W2020/finalists.html
[42] https://sammys.soe.ucsc.edu/

	Query Story 5						
Title	Le Correspondant Perdu		Greyhound	An Extraordinary Day	Locked Phone	The Missing Roommate	Birthday Yacht Murder Case
Content Keywords	Mystery, Spy Story, Second Person, Missing Person		Content Warnings, Drama, Second Person, Mental Health, Depression	Slice of Life, Science Fiction, School, Second Person	Mystery, Second Person, Missing Person	Mystery, Second Person, Science Fiction, Missing Person	Mystery, Second Person, Murder
	Mystery		Content Warnings	Slice of Life	Mystery	Mystery	Mystery
	Spy Story		Drama	Science Fiction	Second Person	Second Person	Second Person
	Second Person		Second Person	School	Missing Person	Science Fiction	Murder
	Missing Person		Mental Health	Second Person		Missing Person	
			Depression				
Structural Keywords	Loop & Grow, Uses Variables, Multiple Endings		No Variables, Multiple Endings	Uses Variables, Multiple Endings, Branch & Bottleneck, Combat System	Uses Variables, Multiple Endings, Loop & Grow	Branch & Bottleneck, Uses Variables, Multiple Endings	Linear, Multiple Endings, Uses Visuals
	Loop & Grow		No Variables	Uses Variables	Uses Variables	Branch & Bottleneck	Linear
	Uses Variables		Multiple Endings	Multiple Endings	Multiple Endings	Uses Variables	Multiple Endings
	Multiple Endings			Branch & Bottleneck	Loop & Grow	Multiple Endings	Uses Visuals
				Combat System			
			2	2	2	2	6
Relevance Judgment (First Coder)			No	No	Yes	Yes	Yes
Relevance Judgment (Second Coder)			No	No	Yes	Yes	Yes

Figure 25. Example of coding structural and content features and generating overall relevance judgments for a subset of the results from a sample Story Navigator query.

Overall precision was better for the text-based search, compared to a random baseline, by around 25% based on the binary relevance judgment (See Figure 26). This is a quantifiable benefit of a system that can inspect the internal structure of a game, performing a search that includes the entirety of the game's textual content while ignoring potentially irrelevant or misleading details in the game's metadata.

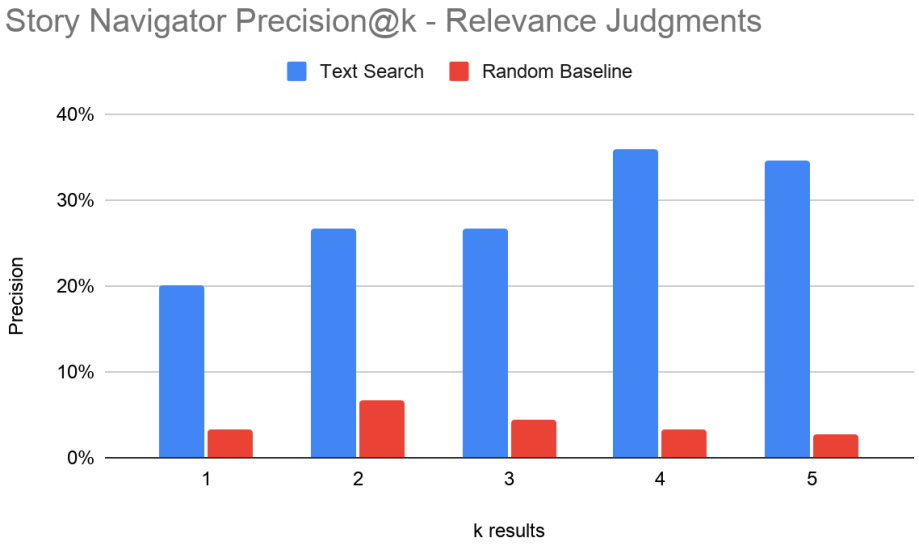

Figure 26. Precision for the Story Navigator system using text-based search compared to a random baseline, across 1-5 top results.

Precision is also quantified as normalized discounted cumulative gain (nDCG, see Chapters 2 and 4). Normalization allows for different lengths of search results to be compared, although for the present comparisons this measure is modified by capping these at the same length (5) due the human effort of making relevance judgments for each Twine game (See Table 6). Similarly, the ideal DCG scores were generated by ordering all feature coded query responses and may have omitted better matches in the corpus. This kind of restriction to a sample for evaluation purposes would apply to any search engine evaluating a sufficiently large corpus (e.g., web search cannot be validated against the entire web).

The Story Navigator search results easily outperformed the random baseline, but they did not always match the ideal ordering according to human evaluators, with

nDCG scores around 25%. This shows that there is ample room for improvement with additional research in this area. Our USE-based model establishes a stronger-than-random baseline, which future research on content-based retrieval for Twine stories can use for evaluating further improvements.

	Ideal DCG	DCG (Text Search)	DCG (Random)	**nDCG (Text Search) /Ideal**	**nDCG (Random) /Ideal**
Query 1	1.25	0.55	0.21	43.46%	16.77%
Query 2	1.37	1.10	0.17	80.04%	12.18%
Query 3	1.42	1.37	0.00	96.75%	0.00%
Query 4	1.14	0.77	0.11	67.64%	9.25%
Query 5	1.42	0.88	0.00	61.83%	0.00%
Average	1.32	0.93	0.10	69.94%	7.64%

Table 6. Normalized discounted cumulative gain (nDCG) for all example analogical search queries when the analogy component was included.

Contributions

This work makes a contribution to technical games research, and a contribution to information retrieval for a complex, culturally relevant and understudied domain. This project may have an immediate pedagogical impact, if the system developed here is brought into the classroom (See Appendix - Study Proposal: Videogame Search in the Classroom).

This work is also a technical contribution to game education, a subset of STEAM education methods (see Chapter 1). The Story Navigator is inherently a contribution to STEAM specifically (and not STEM alone) because it is a technical system designed to enhance the creation and criticism of media artifacts, and it presents the opportunity to ask students to be critical not just of those artifacts but also of the system itself. This can also be considered a contribution to the media literacy part of MESH education (Wise, 2019) (MESH stands for Media Literacy, Ethics, Sociology, and History). In this way it complements the more rounded approach that characterizes games programs which combine the development of interpretive, creative and technical skills into a more complete and modern form of liberal arts education (Brockman, 1996).

Future Work

Recently researchers developed an automated feedback system for student created interactive narrative games (Mahajan et al., 2019). This system ingests games

that were made in the StudyCrafter system, which was initially created to facilitate running psychological experiments. In keeping with this purpose, StudyCrafter games are designed to be exactly repeatable, systematically variable, and to log detailed data about player interactions. Previous work has made use of these features to develop an automated system for investigating structural features of interactive narratives (Partlan et al., 2018).

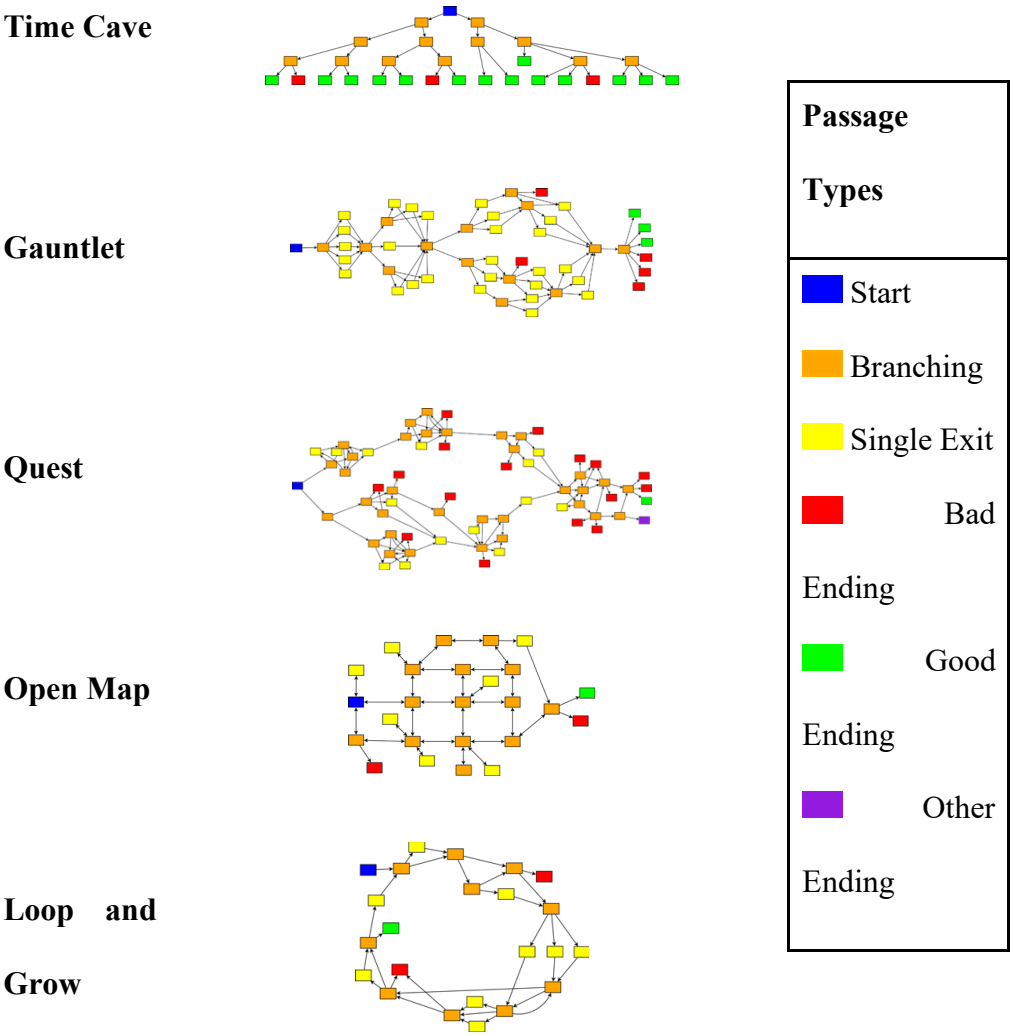

Figure 27. Examples of interactive fiction story graph categories proposed by Ashwell (2015).

An initial review of the structure of Choose-Your-Own-Adventure style stories (Ashwell, 2015) suggests several of the design patterns that future moment search

systems might be looking for (See Figure 25). Because Twine games are effectively distributed with their human-readable source code it is possible to extract structural elements such as links between passages, author created variables, and other features of the game code, enabling analysis specifically of choice-based interactivity. Choice-based interactive fiction games contain links between nodes in a way that resembles the web itself, which can assist in leveraging existing graph analysis and web crawling approaches to develop analysis systems for story structure.

Within the context of textual content alone, other text analysis systems such as LIWC (Pennebaker et al., 2001) and its derivatives like empath (Fast et al., 2016), or other topic modeling approaches like gensim (Řehůřek & Sojka, 2011), might add interesting additional nuances to the analysis of passage and story content.

Chapter 7. Conclusion

Videogames are important in part because of the influence they have on their players and in part because of their steadily increasing popularity. The critical work of analyzing this new media contributes to our understanding of its growing influence, and information gleaned from this analysis can be directly beneficial for game makers as well. Information retrieval technology designed for this media can not only automate tedious aspects of analysis, it can also be used to reveal hidden relationships, and to ask questions that would not have been possible to ask without it. This technology has a clear role to play in the growing field of game studies, for both students and scholars. However, information retrieval research on videogames lags behind that for other types of media, in part because of the inherent challenges that come along with the introduction of interactivity. This dissertation described several efforts to meet those challenges head-on, to develop information retrieval tools for videogames that meet the identified needs of game scholars and experts, and to evaluate the benefits of these tools for students and novice game makers.

Primary Contributions

The projects described in this dissertation primarily contribute to the fields of games studies, information retrieval, and game education. Games studies is a multidisciplinary field of study and learning with games and related phenomena as its subject matter (Mäyrä, 2008). Information retrieval refers to the science of searching

for documents or information within documents, including all forms of media (Manning et al., 2008). Game education is the technical and artistic training of students in the design and analysis of game systems and other interactive media, and in the connection of that media to a broader arts context.

Game Studies: The development of tools for game scholars that aid research, by facilitating both distant (searching across) and close (searching within) readings. This work contributes new technical methods for game studies.

Information Retrieval: A new domain of complex and culturally relevant data introduced for information retrieval, supported with a corpus of documents, a dataset of queries and answers, and example applications of existing metrics.

Game Education: Search and visualization tools used to analyze student game projects, to inspire future student work, and potentially to be incorporated directly into student work processes. The combination of technical, creative, and interpretive work enabled means this is inherently a STEAM contribution.

Revisiting Research Questions

At the beginning of this dissertation, I raised four research questions which have been addressed by the projects described in the chapters above. The first of these questions, **RQ1 [task definition]** asked what various categories of users might be interested in finding within works of interactive media, and how they would be interested in performing this search. This was directly addressed in the requirement

analysis interviews described in Chapter 3. Interviews with users (in the categories of scholars, hobbyists, and students), led to several technical recommendations, such as the inclusion of game data from a variety of platforms, the ability to process video data in the wild, and to comprehend personal data sources. Interview findings were also the direct inspiration for a project exploring

The second research question, **RQ2 [operationalization],** asked how to use a subset of the requirements identified above to create functional search systems for these users. This research question was primarily addressed in the descriptions of the Gameplay Video Explorer system in Chapter 4, and the Story Navigator system in Chapter 6. The former system answers some of these identified needs by allowing for operations on personal data from a variety of sources, and allowing for the construction of complex exploratory queries in the search-by-analogy mode. The latter system does so by allowing for natural language queries, and both systems do so by operating on the internal content of the media they index, rather than the metadata or secondary sources usually used to find games.

The third research question, **RQ3 [opportunity]**, asked how the systems described above could enable further actions, including supporting cases that users explicitly identified (e.g., teaching game design, improving game development, technical games research, etc.), and how the lessons learned over the course of making these systems could help future developers. My requirements analysis (Chapter 3)

offers many answers, and the Undergraduate Games Corpus (Chapter 5) is already helping to advance machine perception research (B. R. Anderson & Smith, 2021).

The fourth and final research question, **RQ4 [impact],** asked for identifying the opportunities that these systems and insights create for games education. Primarily, this is answered by the focus on student work of the Undergraduate Games Corpus (Chapter 5), and by extension the projects that it has enabled such as the Story Navigator (Chapter 6).

Future Work

Below I provide three examples of research directions, each centered on search within and across games, in the context of technical systems, expert/stakeholder evaluation, and classroom integration. For each of these directions, I suggest future projects that would contribute to the growth of videogame moments search field.

Technical Systems

The technical systems described above interact with several specific types of videogame moments, such as those generated from playing SNES ROMs, videos of the same, or from the exploration of passages and possible states in a link-based Twine game. These are each a small subset of the possibilities of game moments, and there are more ways to explore their relationships than those I have described. One obvious missing piece is the inclusion of contextual information about videogame moments, answering questions, for example, about where, when, and who.

Expert/Stakeholder Evaluation

While I was fortunate to have access to a variety of domain experts, scholars, and educators during the evaluation of the systems described above, there are several potential stakeholder profiles that remained out of reach. I would have developed additional insight into projects described above had I been able to contact exemplars of these profiles, but perhaps more importantly their different interests also suggest research questions beyond the scope of those projects. Examples that were under consideration include translators, distributors, and modders. Translators might, for example, have been interested in comparing across game versions. Distributors might be interested in search tools that enable discovery and encourage engagement. Modders might want to perform code search. These speculations could have led to some initial questions in the context of our existing systems, but even more they might have inspired the developing systems that address their specific concerns.

Classroom Integration

While the initial evaluation and planned work described below (see Appendix - Study Proposal: Videogame Search in the Classroom) are in the context of undergraduate education, game creation and evaluation take place in K-12 classes as well. These classroom contexts would presumably also benefit from educational interventions, appropriately adapted to student abilities. Students in many of these grades have demonstrated their ability to create games, and to analyze texts.

Outlook

This dissertation is founded on the acknowledgement that interactive media broadly and games specifically have importance to society, and that the ability to search across and search within games will provide social benefit by supporting game criticism and creation. I have argued that the tools and methods of information retrieval need to be upgraded in order to handle the needs of educators, scholars, developers, and game playing communities at large. My work has a strategic focus on the human makers and scholars of games, particularly in the educational settings where the next generation is being trained. This approach is human-centered, leveraging computational systems and technical games research in support of social and academic motivations. I hope that my observations and efforts described here support further work in videogame moment search, which in turn contributes to meaningful and considered media criticism and creation.

References

2018 Video Game Industry Statistics, Trends & Data—The Ultimate List. (2018, May 15). WePC.Com. https://www.wepc.com/news/video-game-statistics/

2019 Essential Facts About the Computer and Video Game Industry. (2019, May 2). *Entertainment Software Association*. https://www.theesa.com/esa-research/2019-essential-facts-about-the-computer-and-video-game-industry/

Adams, D., & Meretzky, S. (1984). *The Hitchhiker's Guide to the Galaxy*. Infocom.

Altice, N. (2015). *I Am Error: The Nintendo Family Computer / Entertainment System Platform* (N. Montfort & I. Bogost, Eds.; First Edition edition). The MIT Press.

Anderson, B. R., Lim, K., & Smith, A. M. (2019, May). *Introduction to Search by Analogy (with Videogames)*. http://barrettrees.com/search_by_analogy_interactive/IntroductionToSearchByAnalogy.html

Anderson, B. R., & Smith, A. M. (2019). Understanding User Needs in Videogame Moment Retrieval. *14th International Conference on the Foundations of Digital Games*. Foundations of Digital Games, San Luis Obispo, CA, USA.

Anderson, B. R., & Smith, A. M. (2021). The Undergraduate Games Corpus: A Dataset for Machine Perception of Interactive Media. *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*.

Anderson, T., Blank, M., Lebling, D., & Daniels, B. (1980). *Zork I*. Infocom.

André, P., Cutrell, E., Tan, D. S., & Smith, G. (2009). Designing Novel Image Search

Interfaces by Understanding Unique Characteristics and Usage. In T. Gross, J. Gulliksen, P. Kotzé, L. Oestreicher, P. Palanque, R. O. Prates, & M. Winckler (Eds.), *Human-Computer Interaction – INTERACT 2009* (pp. 340–353). Springer Berlin Heidelberg.

Aschwanden, C., & King, R. (2015, August 19). Science Isn't Broken. *FiveThirtyEight*. https://fivethirtyeight.com/features/science-isnt-broken/

Ashwell, S. K. (2015, January 27). Standard Patterns in Choice-Based Games. *These Heterogenous Tasks*. https://heterogenoustasks.wordpress.com/2015/01/26/standard-patterns-in-choice-based-games/

Aycock, J., & Biittner, K. (2020). LeGACy Code: Studying How (Amateur) Game Developers Used Graphic Adventure Creator. *International Conference on the Foundations of Digital Games*, 1–7. https://doi.org/10.1145/3402942.3402988

Baheti, P. (2019, September 17). Search helps you find key moments in videos. *The Keyword*. https://blog.google/products/search/key-moments-video-search/

Barrington, L., Chan, A., Turnbull, D., & Lanckriet, G. (2007). Audio information retrieval using semantic similarity. *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, *2*, II-725-II–728.

Bell, S., & Bala, K. (2015). Learning Visual Similarity for Product Design with Convolutional Neural Networks. *ACM Trans. Graph.*, *34*(4), 98:1–98:10. https://doi.org/10.1145/2766959

Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A Neural Probabilistic Language Model. *Journal of Machine Learning Research*, *3*(Feb), 1137–1155.

Bogers, T., Gäde, M., Koolen, M., Petras, V., & Skov, M. (2018). What was this Movie About this Chick? A Comparative Study of Relevance Aspects in Book and Movie Discovery. *Transforming Digital Worlds*, 323–334. https://doi.org/10.1007/978-3-319-78105-1_36

Borges, J. L. (1941). The Garden of Forking Paths. In N. Wardrip-Fruin & N. Montfort (Eds.), *New Media Reader*. MIT Press.

Brockman, J. (1996). *Third Culture: Beyond the Scientific Revolution* (1st Touchstone Ed edition). Simon and Schuster.

Bruckman, A. (1999). Can educational be fun. *Game Developers Conference*, *99*, 75–79. http://www.cc.gatech.edu/~asb/papers/bruckman-gdc99.pdf

Bybee, R. W. (2010). What Is STEM Education? *Science*, *329*(5995), 996–996. https://doi.org/10.1126/science.1194998

Casucci, T. (2018). *EAE Thesis Game Archive*. ULibraries Research Guides: EAE / Games: Find EAE Thesis Games. //campusguides.lib.utah.edu/eae/archive

Cer, D., Yang, Y., Kong, S., Hua, N., Limtiaco, N., John, R. S., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C., Sung, Y.-H., Strope, B., & Kurzweil, R. (2018). Universal Sentence Encoder. *ArXiv:1803.11175 [Cs]*. http://arxiv.org/abs/1803.11175

Chan, D. M., Rao, R., Huang, F., & Canny, J. F. (2018). T-SNE-CUDA: GPU-

Accelerated T-SNE and its Applications to Modern Data. *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, 330–338. https://doi.org/10.1109/CAHPC.2018.8645912

Chang, K., & Smith, A. (2020). Differentia: Visualizing Incremental Game Design Changes. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, *16*(1), 175–181.

Charmaz, K. (2006). *Constructing grounded theory: A practical guide through qualitative analysis*. Sage.

Deng, J., Dong, W., Socher, R., Li, L., Kai Li, & Li Fei-Fei. (2009). ImageNet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255. https://doi.org/10.1109/CVPR.2009.5206848

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ArXiv:1810.04805 [Cs]*. http://arxiv.org/abs/1810.04805

Dunbar, K., & Blanchette, I. (2001). The in vivo/in vitro approach to cognition: The case of analogy. *Trends in Cognitive Sciences*, *5*(8), 334–339.

Ebner, M., Levine, J., Lucas, S. M., Schaul, T., Thompson, T., & Togelius, J. (2013). *Towards a video game description language*.

Egenfeldt-Nielsen, S. (2011). What Makes a Good Learning Game?: Going beyond

edutainment. *ELearn*, *2011*(2), 2.

Ertmer, P. A., Ottenbreit-Leftwich, A. T., Sadik, O., Sendurur, E., & Sendurur, P. (2012). Teacher beliefs and technology integration practices: A critical relationship. *Computers & Education*, *59*(2), 423–435.

Fast, E., Chen, B., & Bernstein, M. S. (2016). Empath: Understanding topic signals in large-scale text. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 4647–4657.

Fauconnier, G., & Turner, M. (2008). *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books.

Feldman, A. (2015, June 16). *Why We Need to Put the Arts Into STEM Education*. Slate Magazine. https://slate.com/technology/2015/06/steam-vs-stem-why-we-need-to-put-the-arts-into-stem-education.html

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, *19*(2), 141–205. https://doi.org/10.1016/0364-0213(95)90016-0

Gee, J. P. (2013). *The Anti-Education Era: Creating Smarter Students through Digital Learning*. St. Martin's Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*(2), 155–170.

Gentner, D., & Forbus, K. D. (2011). Computational models of analogy. *Wiley Interdisciplinary Reviews. Cognitive Science*, *2*(3), 266–276.

https://doi.org/10.1002/wcs.105

Glaser, B. G. (1992). *Basics of grounded theory analysis: Emergence vs forcing*. Sociology press.

Hart, V., & Case, N. (2014). *Parable of the Polygons*. Parable of the Polygons. http://ncase.me/polygons

Harteveld, C., Stahl, A., Smith, G., Talgar, C., & Sutherland, S. C. (2016). Standing on the Shoulders of Citizens: Exploring Gameful Collaboration for Creating Social Experiments. *2016 49th Hawaii International Conference on System Sciences (HICSS)*, 74–83. https://doi.org/10.1109/HICSS.2016.18

Hearst, M. A. (2009). *Search User Interfaces* (1 edition). Cambridge University Press.

Hofstadter, D. R. (2001). Analogy as the core of cognition. *The Analogical Mind: Perspectives from Cognitive Science*, 499–538.

Hohman, F., Conlen, M., Heer, J., & Chau, D. H. (Polo). (2020). Communicating with Interactive Articles. *Distill*, *5*(9), e28. https://doi.org/10.23915/distill.00028

Hubert, J. (2013). Billions of times a day in the blink of an eye. *Google Inside Search: The Official Google Search Blog*. https://search.googleblog.com/2013/03/billions-of-times-day-in-blink-of-eye.html

Ince, S. F., Hoadley, C., & Kirschner, P. A. (2018). A Study of Search Practices in Doctoral Student Scholarly Workflows. *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*, 245–248.

https://doi.org/10.1145/3176349.3176877

Jensen, A. E., Jægerfelt, C. M., Francis, S., Larsen, B., & Bogers, T. (2018). I Just Scroll Through My Stuff Until I Find It or Give Up: A Contextual Inquiry of PIM on Private Handheld Devices. *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval*, 140–149. https://doi.org/10.1145/3176349.3176394

Jørgensen, I. K. H., & Bogers, T. (2020, September). "Kinda like The Sims... But with ghosts?": A Qualitative Analysis of Video Game Re-finding Requests on Reddit. *Proceedings of the 15th International Conference on the Foundations of Digital Games*. Buġibba, Malta.

Kaltman, E., Caldwell, C., Ryan, J., Barrett, M., & Wardrip-Fruin, N. (2014). *GAMECIP*. The Game Metadata and Citation Project. https://gamecip.soe.ucsc.edu/

Kaltman, E., Osborn, J., Wardrip-Fruin, N., & Mateas, M. (2017). Getting the GISST: A Toolkit for the Creation, Analysis and Reference of Game Studies Resources. *Proceedings of the 12th International Conference on the Foundations of Digital Games*, 16:1–16:10. https://doi.org/10.1145/3102071.3102092

Kaltman, E., Wardrip-fruin, N., Mastroni, M., Lowood, H., groat, G. D., Edwards, G., Barrett, M., & Caldwell, C. (2016). Implementing Controlled Vocabularies for Computer Game Platforms and Media Formats in SKOS. *Journal of Library Metadata*, *16*(1), 1–22. https://doi.org/10.1080/19386389.2016.1167494

Khaled, R., Lessard, J., & Barr, P. (2018). Documenting Trajectories in Design Space: A Methodology for Applied Game Design Research. *Proceedings of the 13th International Conference on the Foundations of Digital Games*, 27:1–27:10. https://doi.org/10.1145/3235765.3235767

Klimas, C. (2009). *Twine / An open-source tool for telling interactive, nonlinear stories*. http://twinery.org/

Kopas, M., Bright, B., Pratt, M., Problems, E., McHenry, T., Sampat, E., Brough, M., Bee, A., & Lake, W. (2015). *Videogames for Humans*. Instar.

Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, *10*(2), 141–160. https://doi.org/10.1016/j.cogsys.2006.07.004

Lastowka, G. (2013, May 17). All Your Nintendo Let's Plays Are Belong To Nintendo? *Gamasutra*. https://www.gamasutra.com/blogs/GregLastowka/20130517/192424/All_Your_Nintendo_Lets_Plays_Are_Belong_To_Nintendo.php

Lee, J. H., Clarke, R. I., & Perti, A. (2015). Empirical evaluation of metadata for video games and interactive media. *Journal of the Association for Information Science and Technology*, *66*(12), 2609-2625.

Leidner, J. L., & Plachouras, V. (2017). Ethical by Design: Ethics Best Practices for Natural Language Processing. *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, 30–40.

https://doi.org/10.18653/v1/W17-1604

Love, C. (2013). *Even Cowgirls Bleed*. http://scoutshonour.com/cowgirl/

Love, N., Hinrichs, T., Haley, D., Schkufza, E., & Genesereth, M. (2006). *General Game Playing: Game Description Language Specification. Rapport technique LG-2006-01*. Stanford University.

Lowther, D. L., Inan, F. A., Daniel Strahl, J., & Ross, S. M. (2008). Does technology integration "work" when key barriers are removed? *Educational Media International*, *45*(3), 195–213.

Maaten, L. van der, & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, *9*(Nov), 2579–2605.

Mahajan, S., Bunyea, L., Partlan, N., Schout, D., Harteveld, C., Matuk, C., Althoff, W., Duke, T., Sutherland, S., & Smith, G. (2019). *Toward Automated Critique for Student-Created Interactive Narrative Projects*.

Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to Information Retrieval* (1 edition). Cambridge University Press.

Mäyrä, F. (2008). *An Introduction to Game Studies: Games in Culture*. Sage.

McInnes, L., Healy, J., & Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *ArXiv:1802.03426 [Cs, Stat]*. http://arxiv.org/abs/1802.03426

McIntosh, T., & Gregory, J. (2013). *The Last of Us* [English; PlayStation 3]. Naughty Dog.

Megyesi, B., Näsman, J., & Palmér, A. (2016). The uppsala corpus of student writings: Corpus creation, annotation, and analysis. *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 3192–3199.

Melcer, E., Nguyen, T.-H. D., Chen, Z., Canossa, A., El-Nasr, M. S., & Isbister, K. (2015). Games research today: Analyzing the academic landscape 2000-2014. *Network*, *17*, 20.

Metz, C. (2016, March 16). In Two Moves, AlphaGo and Lee Sedol Redefined the Future. *Wired*. https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/

Mikolov, T., Yih, W., & Zweig, G. (2013). Linguistic regularities in continuous space word representations. *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 746–751.

Millar, L. (2009). *The Story Behind the Book: Preserving Authors' and Publishers' Archives*. Canadian Centre for Studies in Publishing Press.

Montfort, N. (2005). *Twisty Little Passages: An Approach to Interactive Fiction*. MIT Press.

Moyer, C. (2016, March 28). *How Google's AlphaGo Beat a Go World Champion*. The Atlantic. https://www.theatlantic.com/technology/archive/2016/03/the-invisible-opponent/475611/

Nelson, M. J. (2020, January). *Institutions Active in Technical Games Research*. https://www.kmjn.org/game-rankings/

Nesi, H., Sharpling, G., & Ganobcsik-Williams, L. (2004). Student papers across the curriculum: Designing and developing a corpus of British student writing. *Computers and Composition*, *21*(4), 439–450. https://doi.org/10.1016/j.compcom.2004.08.003

Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press. New York.

North, R. (2016). *Romeo and/or Juliet: A Chooseable-Path Adventure*. Riverhead Books.

Owens, T. (2014, February 27). *What Do you Mean by Archive? Genres of Usage for Digital Preservers | The Signal* [Webpage]. //blogs.loc.gov/thesignal/2014/02/what-do-you-mean-by-archive-genres-of-usage-for-digital-preservers/

Packard, E. (1979). *The Cave of Time* (1st edition). Bantam Books.

Partlan, N., Carstensdottir, E., Kleinman, E., Snodgrass, S., Harteveld, C., Smith, G., Matuk, C., Sutherland, S. C., & El-Nasr, M. S. (2019). Evaluation of an automatically-constructed graph-based representation for interactive narrative. *Proceedings of the 14th International Conference on the Foundations of Digital Games*, 1–9. https://doi.org/10.1145/3337722.3341858

Partlan, N., Carstensdottir, E., Snodgrass, S., Kleinman, E., Smith, G., Harteveld, C.,

& El-Nasr, M. S. (2018, September 25). Exploratory Automated Analysis of Structural Features of Interactive Narrative. *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*. Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference. https://www.aaai.org/ocs/index.php/AIIDE/AIIDE18/paper/view/18103

Pell, B. (1993). *Strategy generation and evaluation for meta-game playing*. Citeseer.

Pell, B. (1994). A strategic metagame player for general chess-like games. *Proceedings of the Twelfth AAAI National Conference on Artificial Intelligence*, 1378–1385.

Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates*, *71*(2001), 2001.

Perez-Liebana, D., Samothrakis, S., Togelius, J., Lucas, S. M., & Schaul, T. (2016). *General video game ai: Competition, challenges, and opportunities*.

Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. *ArXiv Preprint ArXiv:1802.05365*.

Peyre, J., Laptev, I., Schmid, C., & Sivic, J. (2019). *Detecting Unseen Visual Relations Using Analogies*. 1981–1990. http://openaccess.thecvf.com/content_ICCV_2019/html/Peyre_Detecting_Unseen_Visual_Relations_Using_Analogies_ICCV_2019_paper.html

Piaget, J. (1977). *Psychology and Epistemology: Towards a Theory of Knowledge* (Reprint edition). Penguin Books.

Plass, J. L., Homer, B. D., & Kinzer, C. K. (2015). Foundations of game-based learning. *Educational Psychologist*, *50*(4), 258–283. http://dx.doi.org/10.1080/00461520.2015.1122533

Porpentine, C. H. (2012). *Howling dogs*. http://slimedaughter.com/games/twine/howlingdogs/

Prabhu, V. U., & Birhane, A. (2020). Large image datasets: A pyrrhic win for computer vision? *ArXiv:2006.16923 [Cs, Stat]*. http://arxiv.org/abs/2006.16923

Purdy, J. P. (2009). Anxiety and the Archive: Understanding Plagiarism Detection Services as Digital Archives. *Computers and Composition*, *26*(2), 65–77. https://doi.org/10.1016/j.compcom.2008.09.002

Queneau, R. (1973). Yours For the Telling. In N. Wardrip-Fruin & N. Montfort (Eds.), & J. Crombie (Trans.), *New Media Reader*. MIT Press.

Řehůřek, R., & Sojka, P. (2011). Gensim—Statistical semantics in python. *Statistical Semantics; Gensim; Python; LDA; SVD*.

Robin, B. R. (2008). Digital Storytelling: A Powerful Technology Tool for the 21st Century Classroom. *Theory Into Practice*, *47*(3), 220–228. https://doi.org/10.1080/00405840802153916

Rockenberger, A. (2014). Video Game Framings. In N. Desrochers & D. Apollon (Eds.), *Examining Paratextual Theory and its Applications in Digital Culture* (pp. 252–286). IGI Global.

Rui, Y., Huang, T. S., & Mehrotra, S. (1997). Content-based image retrieval with

relevance feedback in MARS. *Proceedings of International Conference on Image Processing*, *2*, 815–818.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*(6088), 533–536. https://doi.org/10.1038/323533a0

Ryan, J., Kaltman, E., Fisher, A. Ma., Owen-Milner, T., Mateas, M., & Wardrip-Fruin, N. (2017). *GameSpace: An Explorable Visualization of the Videogame Medium*.

Ryan, J., Kaltman, E., Hong, T., Isbister, K., Mateas, M., & Wardrip-Fruin, N. (2016). *GameNet and GameSage: Videogame Discovery as Design Insight*. http://www.digra.org/wp-content/uploads/digital-library/paper_201.pdf

Sachdev, S., Li, H., Luan, S., Kim, S., Sen, K., & Chandra, S. (2018). Retrieval on Source Code: A Neural Code Search. *Proceedings of the 2Nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, 31–41. https://doi.org/10.1145/3211346.3211353

Sakamoto, Y., & Kano, M. (1994). *Super Metroid* (US) [Computer software]. Nintendo.

Salter, A. (2016). Educational Games. In Henry Lowood and Raiford Guins (Ed.), *Debugging Game History: A Critical Lexicon* (p. 119). MIT Press.

Salton, G., Wong, A., & Yang, C.-S. (1975). A vector space model for automatic indexing. *Communications of the ACM*, *18*(11), 613–620.

Sanderson, M., & Croft, W. B. (2012). The History of Information Retrieval Research. *Proceedings of the IEEE*, *100*(Special Centennial Issue), 1444–1451. https://doi.org/10.1109/JPROC.2012.2189916

Sherry, J. L. (2013). Formative research for STEM educational games: Lessons from the Children's Television Workshop. *Zeitschrift Für Psychologie*, *221*(2), 90–97. http://dx.doi.org/10.1027/2151-2604/a000134

Shih-Fu Chang, Chen, W., Meng, H. J., Sundaram, H., & Di Zhong. (1998). A fully automated content-based video search engine supporting spatiotemporal queries. *IEEE Transactions on Circuits and Systems for Video Technology*, *8*(5), 602–615. https://doi.org/10.1109/76.718507

Shuler, C. (2012). *What in the World Happened to Carmen Sandiego? The Edutainment Era: Debunking Myths and Sharing Lessons Learned*. http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.362.5753

Shum, H. (2012). Introducing Bing Search Quality Insights. *Bing Blogs*. https://blogs.bing.com/search/2012/03/05/introducing-bing-search-quality-insights/

Sloane, M., Moss, E., Awomolo, O., & Forlano, L. (2020). Participation is not a Design Fix for Machine Learning. *ArXiv:2007.02423 [Cs]*. http://arxiv.org/abs/2007.02423

Subrahmanyam, K., Greenfield, P., Kraut, R., & Gross, E. (2001). The impact of computer use on children's and adolescents' development. *Journal of Applied*

Developmental Psychology, *22*(1), 7–30. https://doi.org/10.1016/S0193-3973(00)00063-0

Tezuka, T., Miyamoto, S., & Nakago, T. (1991). *Super Mario World* (US) [Computer software]. Nintendo.

The loss of "Erie": EAE Thesis Games Archive. (2018). https://utah.instructure.com/courses/548092/pages/the-loss-of-erie?module_item_id=8970479

Theimer, K. (2012). Archives in Context and as Context. *Journal of Digital Humanities*. http://journalofdigitalhumanities.org/1-2/archives-in-context-and-as-context-by-kate-theimer/

Thielscher, M. (2010). A general game description language for incomplete information games. *AAAI*, *10*, 994–999.

Thomee, B., & Lew, M. S. (2012). Interactive search in image retrieval: A survey. *International Journal of Multimedia Information Retrieval*, *1*(2), 71–86. https://doi.org/10.1007/s13735-012-0014-4

Treanor, M., & Mateas, M. (2009). Newsgames-Procedural Rhetoric Meets Political Cartoons. *DiGRA Conference*.

Tulving, E. (1972). Episodic and semantic memory. In *Organization of memory* (pp. xiii, 423–xiii, 423). Academic Press.

Typke, R., Wiering, F., & Veltkamp, R. C. (2005). A survey of music information retrieval systems. *Proc. 6th International Conference on Music Information*

Retrieval, 153–160.

Vallet, D., Fernández, M., & Castells, P. (2005). An ontology-based information retrieval model. *European Semantic Web Conference*, 455–470.

Van Eck, R. (2006). Digital game-based learning: It's not just the digital natives who are restless. *EDUCAUSE Review*, *41*(2), 16.

van Stegeren, J., & Theune, M. (2020). Fantastic Strings and Where to Find Them:The Quest for High-Quality Video Game Text Corpora. *2020 Workshop in Intelligent Narrative Technologies (INT'20)*.

Victor, B. (2011, March 10). *Explorable Explanations*. http://worrydream.com/ExplorableExplanations/

Wattenberg, M., Viégas, F., & Johnson, I. (2016). How to Use t-SNE Effectively. *Distill*, *1*(10), e2. https://doi.org/10.23915/distill.00002

Welhouse, Z., Lee, J. H., & Bancroft, J. (2015). "What Am I Fighting For?": Creating a Controlled Vocabulary for Video Game Plot Metadata. *Cataloging & Classification Quarterly*, *53*(2), 157-189.

Wilks, J., Cutcher, A., & Wilks, S. (2012). Digital Technology in the Visual Arts Classroom: An [Un]Easy Partnership. *Studies in Art Education*, *54*(1), 54–65. https://doi.org/10.1080/00393541.2012.11518879

Wise, T. (2019, September 23). *Forget STEM, We Need MESH*. Medium. https://medium.com/our-human-family/forget-stem-we-need-mesh-43ab6f6273cd

Yakman, G. (2008). *STEAM Education: An overview of creating a model of integrative education*. Pupils' Attitudes Towards Technology (PATT-19) Conference: Research on Technology, Innovation, Design & Engineering Teaching, Salt Lake City, Utah, USA. https://www.academia.edu/8113795/STEAM_Education_an_overview_of_creating_a_model_of_integrative_education

Yanai, I., & Lercher, M. (2020). A hypothesis is a liability. *Genome Biology*, *21*(1), 231. https://doi.org/10.1186/s13059-020-02133-w

Yannakakis, G. N., & Togelius, J. (2018). Modeling Players. In G. N. Yannakakis & J. Togelius (Eds.), *Artificial Intelligence and Games* (pp. 203–255). Springer International Publishing. https://doi.org/10.1007/978-3-319-63519-4_5

Young, M. F., Slota, S., Cutter, A. B., Jalette, G., Mullin, G., Lai, B., Simeoni, Z., Tran, M., & Yukhymenko, M. (2012). Our princess is in another castle: A review of trends in serious gaming for education. *Review of Educational Research*, *82*(1), 61–89.

Zhan, Z., Aytemiz, B., & Smith, A. M. (2019). Taking the Scenic Route: Automatic Exploration for Videogames. *Proceedings of the 2nd Workshop on Knowledge Extraction from Games Co-Located with 33rd AAAI Conference on Artificial Intelligence*, 26–34.

Zhan, Z., & Smith, A. M. (2018). Retrieving game states with moment vectors. *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*.

Zhang, X. (2018, July 2). Visualizing Thousands of Moments in Games with Three.js – Interactions (1). *Zooey The Engineer*. https://zooeytheengineer.wordpress.com/

Zhang, X. (2019). *Navigating Moments in Videogames* [Project report for MS in Computational Media]. UC Santa Cruz.

Zhang, X., & Smith, A. M. (2019). Retrieving Videogame Moments with Natural Language Queries. *Proceedings of the 14th International Conference on the Foundations of Digital Games*, 74:1–74:7. https://doi.org/10.1145/3337722.3341867

Zhang, X., Zhan, Z., Holtz, M., & Smith, A. M. (2018). Crawling, Indexing, and Retrieving Moments in Videogames. *Proceedings of the 13th International Conference on the Foundations of Digital Games*, 16:1–16:10. https://doi.org/10.1145/3235765.3235786

Zheng, Q.-F., Wang, W.-Q., & Gao, W. (2006). Effective and Efficient Object-based Image Retrieval Using Visual Phrases. *Proceedings of the 14th ACM International Conference on Multimedia*, 77–80. https://doi.org/10.1145/1180639.1180664

Appendix - Undergraduate Game Collection Guide

Overview

This is a guide for collecting undergraduate student games for the Undergraduate Games Corpus. This was created with and is focused on collecting games from the UC Santa Cruz Computational Media course Foundations of Video Game Design (CMPM 80K), but can be adapted for other courses. The process has the following steps:

1. Course Integration
2. Game Archiving Form Data Collection
3. Linking Canvas Game Data and adding to the UGC and eScholarship

Ethics Note

The focus of this guide is on the collection of student games, and metadata about those games. Except as they choose to express themselves in these games, this process collects no information about the student authors themselves. Any change to this process that results in collecting data about students (and not just about their games) should be submitted to the IRB for review.

Course Integration

Game collection must be optional, but if the effort is going to be worthwhile the instructor will need to make responding to the archiving form a course requirement

(for credit), and the collection process needs to be presented to students. One of the options on the first page of the archiving form is to opt out. In some previous quarters, if response to the collection form was optional, no games were submitted to the collection. This contrasts with an over 80% submission rate when viewing the form, even just to opt out, was required to receive (a very small amount of) course participation credit. The collection process, and its motivations and potential benefits, should be presented to the students in class. In previous quarters this has happened as part of a guest lecture introducing technical games research, which has the added benefit of demonstrating some of the projects that a games corpus could enable. Even if a full lecture is not possible, a brief overview is likely to motivate more submissions to the corpus.

Game Archiving Forms

There should be a distinct archiving form for each game the students submit. These forms are modified slightly for each type of game. Links to examples from past quarters follow. Note that examples ask for explicit permission for archiving, and provide example game summaries and keywords specific to each assignment (e.g. Twine or graphical). Also note that a second page of each form allows students to choose how to license their work (Copyright or Creative Commons).

Twine Games Archiving Form: https://forms.gle/xPfT9GHSrFmyXeVx5

Graphical Games Archiving Form example:

https://forms.gle/awcFwaGw4RNoFtGu6

Linking to Canvas Game Data and preparing for eScholarship Upload

The data from the archiving forms and the final games that students submit on Canvas need to be combined, and the now combined data needs to be prepared for incorporation into the Undergraduate Games Corpus, as well as for upload to eScholarship. The code for doing so is available in a Jupyter Notebook that can be shared upon request. Instructions at top of the notebook describe

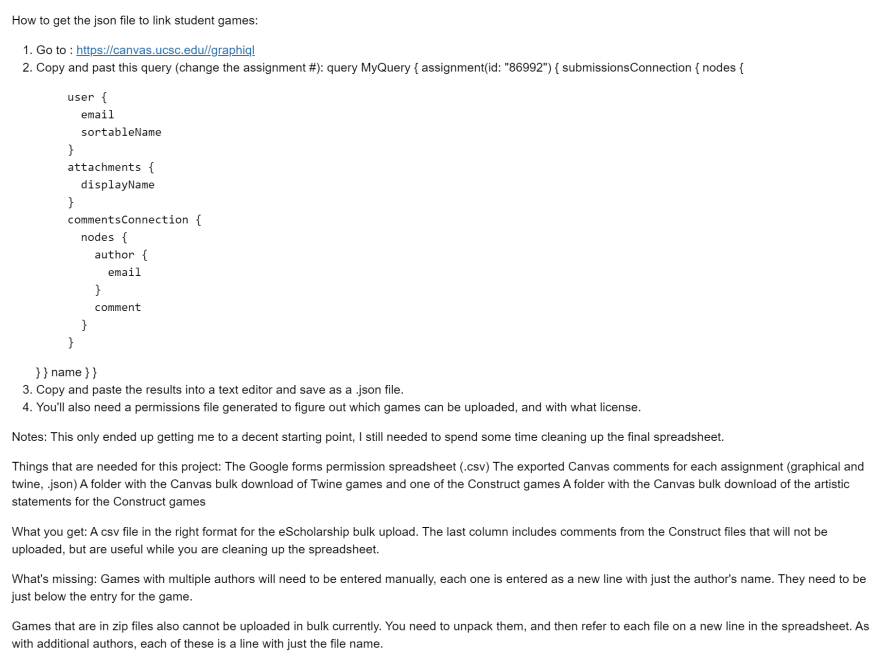

How to get the json file to link student games:

1. Go to : https://canvas.ucsc.edu//graphiql
2. Copy and past this query (change the assignment #): query MyQuery { assignment(id: "86992") { submissionsConnection { nodes {

```
user {
  email
  sortableName
}
attachments {
  displayName
}
commentsConnection {
  nodes {
    author {
      email
    }
    comment
  }
}
```

} } name } }

3. Copy and paste the results into a text editor and save as a .json file.
4. You'll also need a permissions file generated to figure out which games can be uploaded, and with what license.

Notes: This only ended up getting me to a decent starting point, I still needed to spend some time cleaning up the final spreadsheet.

Things that are needed for this project: The Google forms permission spreadsheet (.csv) The exported Canvas comments for each assignment (graphical and twine, .json) A folder with the Canvas bulk download of Twine games and one of the Construct games A folder with the Canvas bulk download of the artistic statements for the Construct games

What you get: A csv file in the right format for the eScholarship bulk upload. The last column includes comments from the Construct files that will not be uploaded, but are useful while you are cleaning up the spreadsheet.

What's missing: Games with multiple authors will need to be entered manually, each one is entered as a new line with just the author's name. They need to be just below the entry for the game.

Games that are in zip files also cannot be uploaded in bulk currently. You need to unpack them, and then refer to each file on a new line in the spreadsheet. As with additional authors, each of these is a line with just the file name.

The final step requires coordinating with a librarian and the full details are beyond the scope of this guide. Currently the best way to establish the contact necessary to begin this process is to generate a support ticket by emailing help@escholarship.org.

Appendix - Study Proposal: Videogame Search in the Classroom

Project Overview

This appendix proposes a study with two distinct but compatible goals, evaluating and improving the videogame search systems created in the previous projects (primarily the Story Navigator), and providing an intervention contributing to the education of students in a course that involves the creation and interpretation of games.

This builds on my previous collaborations with the instructors of an introductory game design course (CMPM 80K, Foundations of Videogame Design) to integrate the research project over several academic quarters. During the Winter and Summer 2019 Foundations of Videogame Design (CMPM 80K) and Game Design Experience (CMPM 120) courses, provided guest lectures on technical games research to game design students, and solicited permission to add student projects to the existing eScholarship student games archive. These interventions also included asking students to provide survey-based feedback speculating on ways they might use videogame moment search in general. In Winter 2020 I followed up by conducting a preliminary focus group study with introductory game design students.

An early effort, simply presenting the archive to introductory game design students, without any kind of scaffolding for incorporating it into their own work, led to virtually no engagement. When the eScholarship archive was presented in class,

there were errors that made downloading games from it impossible. These errors were not reported by students attempting to use the archive. In future courses, if the archive is to be of use, an assignment that directly incorporates using it will be necessary. Similarly, one quarter of requesting archive contributions without incorporating those requests into the class (as a guest lecture, or a course requirement) lead to no student participation.

Classroom Integration

Integrating technology into classroom settings is an active area of research (Ertmer et al., 2012; Lowther et al., 2008), which has led to several insights relevant to thinking about how new tools can best provide educational support. First, a specific integration plan needs to be developed, which will vary depending on specific pedagogical context and goals (Wilks et al., 2012). Information retrieval technology like search is well suited to support a constructivist approach to education, where students are encouraged and expected to create their own knowledge (Piaget, 1977). Looking to previous interventions, there are a few practices and principles that can guide my approach (Robin, 2008).

Technical Games Research Guest Lectures

To orientate and motivate students to contribute to the archives, and to further their education, I provided guest lectures discussing work on technical games research (TGR). In these lectures I described and demonstrated several technical games research

projects, including those that came out of GameCIP such as GameSpace and GameSage, as well as projects developed in the Design Reasoning Lab such as Zhang's natural language search system and my own Gameplay Video Explorer and Story Navigator prototypes. I discussed some of the technical underpinnings of this work, shared UCSC's TGR rankings, and described several ways that students might get involved in TGR over the course of their academic career. With an eye toward contributions to game moments search, I suggested relevant projects for several courses. I also suggested computational media research labs to join, and shared their undergraduate research assistant solicitations. These discussions ended with an appeal to contribute to the games archive

Focus Group Evaluation

During course sections for CMPM 80K in Winter 2020, I solicited feedback from students in a focus group study, prompting discussions about how they might use a game search system to facilitate their own creative and interpretive work. These conversations included 40 participants in 8 focus groups, and took about 40 minutes (out of a 95 minute course sections), resulting in 285 minutes of recorded conversations. Rigorous analysis of the focus group conversations is forthcoming, and may lead to several insights and suggestions that can be incorporated into the present search tools, or provide suggestions for future research. This analysis has been delayed because preliminary observations suggest that it will primarily lead to improvements in the design of the focus group study itself, and that those improvements would not be

possible to implement (at least without substantial additional design changes) in the current remote learning environment.

Classroom Study Design

Insights gleaned from both of these interventions, and from the continued development of the search tools themselves, could be used to develop a more integrated and focused user study. These studies may involve providing students with an overview of technical games research, introducing them to a novel search tool that is immediately relevant to their own course work, and observing the ways that they use such a tool. This may include elements focused on specific questions raised or not clearly answered by the previous observations, possibly including coursework based on insights derived from using the search tool.

This could include a combination of educational and user studies methods, including lecture, direct observations, group conversations, survey data collection. In some cases this may also include one-on-one interviews. I suggest that a final phase of this project would, with the instructor's consent, include a short activity that uses the tool educationally. For example, students might be asked to use the tool to find a relevant game (in relation to their own game) from a previous class, showing a screenshot of how the tool displays the relationship, and writing a brief report describing the relationship. This might be a highly similar game, a highly different, a game that is similar structurally but different in content, or the reverse. The complex analysis required by such an assignment would be educationally appropriate and

provide useful feedback about the tool itself. If it is compatible with the structure of the course, this assignment might also be incorporated into an educational experiment. In this case a subset of students would be randomly assigned to do an analysis using these tools, while a different random subset would do comparable analysis using traditional methods. In this case, it would be useful to receive an IRB exemption (Category II) so that differences in educational outcomes could be reported. The advantages of this approach will need to be weighed against the advantage of acquiring user data from all enrolled students.

Proposed Evaluation

I suggest three distinct forms of evaluation for the data collected from this project: direct observation of and interviews with individual students, collecting and coding the intent behind student queries, and the examination of short reflective student documents. The first set of observations would include talk-aloud, user observation, and survey methods. Student sessions using the tool during the evaluation process will collect queries, which will be coded and categorized for expressed intent. This coding will be informed by insights from classroom observations, guided group conversations, a small number of one-on-one interviews, and survey responses. If the intervention does include a written assignment, these documents could also be coded in ways relevant to how the students used the tool. These categories will not be generated in advance but built up from observations of the documents themselves in a grounded coding process. Insights could then be derived and compared from each of these

evaluation methods, aimed at the two broad goals of this project, of technical systems development and pedagogical intervention. The outcomes of this process may resemble those reported in the requirements analysis interviews reported above, but from a student rather than expert perspective.

Expected Findings and Contributions

It is also important to note here that this is feedback provided primarily by novice game designers. Earlier in this dissertation I described soliciting search system requirements from experts and educators, novice designers and students are important voices to add to this discussion, in part because they represent a greater diversity of perspectives (both numerically and through less rigorous self-selection at their early career stage), and in part because they have different needs than those of expert users.

Appendix – Videogame Search Requirements Analysis Interviews

This document to summarizes information from several expert interviews related to the videogame moment retrieval project, followed by notes from individual conversations.

Name	Profile	Professional Titles or Relevant Experience
Batu Aytemiz	Scholar, Designer	Graduate Student in Computational Media
Stella Mazeika	Scholar, Speedrunner	PhD Candidate in Computer Science
Alexis Ross "Protomagicalgirl"	Speedrunner	Awesome Games Done Quick announcer
Bernd “Aran Jaeger” Clemens	Speedrunner	Member of the TAS community, Author of a detailed guide for speedrunning Super Metroid
Henry Lowood	Curator, Historian, Educator	Curator for History of Science and Technology and for Film and Media at Stanford University
Elizabeth Swensen	Designer, Educator	Assistant Professor of Art and Design, Games and Playable Media, 30 under 30
Nathan Altice	Platform studies, Educator	Teaching Professor, Computational Media

Current Search Strategies

Aytemiz: Google, hoping someone else has already described the moment, mentioned gamespace and gamesage

Mazeika: Google primarily, followed by youtube/twitch, followed by manual play

Ross (Speedrunner): Speedruns.com, resources tab, telling other people to use the same

Clemens (Speedrunner): TASVideos (prefers sticking to classic games)

Swensen: Casual Internet searches, friends and acquaintances, interested in improving instruction- currently providing students with additional support primarily from personal experience

Lowood (Curator/Historian): Mentioned browsing forums (as being more valuable than searching for already known items, initially), and also the Internet Archive searchable collection of speedruns. For search education, discussed giving presentations about discoverability in the library community.

Nathan Altice (Scholar platform studies): Primarily using Youtube for Let's Plays, some Japanese video sites as well, very rarely setting up original hardware and doing capture

Summary: As expected, most individuals use Google or other modern web search engines when looking for games. The biggest exceptions were speedrunners, who use specialized websites that already curate tools and information for their needs. Our other exceptions involved individuals who have some direct connection to specific tools (e.g. helping with development, or knowing the creators).

Search Query Suggestions

Aytemiz (Designer/Scholar): Interested in searching with high level concepts (e.g. search for tutorials, search by emotional valence).

Mazeika (Scholar/Speedrunner): Interested in uploading her own runs and compare them to each other, and see similar ones.

Ross (Speedrunner): Wants to be able to find a specific speedrunning trick with a relevant screenshot. Pointed out that they would prefer to use screenshots, since if they already knew the name Google would work.

Clemens (Speedrunner): Interested in using memory to search

Swensen (Designer/Educator): Wants there to be a language for mechanics that is searchable across games. (Looking for things like "joystick text input" or games with "morality systems"

Lowood (Curator/Historian): Primarily interested in text searches for known events, e.g. a particular tournament win that used an innovative strategy, or an event that took place in a virtual world. Also interested in the possibilities of searching via code.

Nathan Altice (Scholar platform studies): Primarily interested in the ability to do "reverse lookup" from a screenshot, possibly obtained from a low quality video, or magazine scan, also interested in searching by video, by platform characteristics (e.g. show me all games between 1984-85 with a blue background and screen-to-screen movement for the NES and arcade)

Summary: Designers and scholars who are interested in specific game aspects, mechanics, or moments would like to be able to search for them by text or perhaps keyword. Speedrunners are more open to searching by image or even by memory. Henry Lowood suggested that there might be some interest in searching via code,

perhaps to find out which event or capabilities some section of code enables. A speedrunner and a scholar indicated that their primary query format would be screenshots.

Search Result Suggestions

Aytemiz (Designer/Scholar): Image would be the bare minimum, save states would be the best (until he heard about playable in-browser states, that was Ideal+).

Mazeika (Scholar/Speedrunner): She'd like to explore branching paths as a scholar, or narrow down to specific frames/moments as a speedrunner. Video with highly adjustable speed. As a scholar would like to have every citation format (no idea what will be needed).

Ross (Speedrunner): Wants to get a practice save for speedrunning. Interest in filtering results by source (e.g. no interest in seeing results from casual play when looking for a speedrunning trick). Also asked if we can connect our results to game-specific tools (e.g. Fumen for Tetris), and suggested linking search results to live play on Twitch. As a college student, is interested in our exportable citations.

Clemens (Speedrunner): Highest priority interest is in links to playable moments, or downloadable save states. (Pointed out that these might vary between emulators). Particularly interested in edge cases and weird/rare game states.

Swensen (Designer/Educator): Wants to be able to reach a moment, and hand over control to a student, so they know not just how it looks, but how it feels. Otherwise

looking for visual evidence (video), ideally with a reference for how to get to that moment (walkthrough).

Lowood (Curator/Historian): Looking for videos or screenshots of an event. For a mechanic, hoping to be able to grab a replay (e.g. GISST) or a save state. Would like to know what version of a game a screenshot if from (e.g. is this a possible game state? Is this a hack?). Suggested that it would be valuable to see the query image while looking at the results.

Nathan Altice (Scholar platform studies): would like to be able to know if the moment is possible/ reachable, to validate it came from a legit ROM or identify the hack/prototype/alternate version, also interested in obtaining music from specific points, and in getting some form of citation link

Summary: Designers were interested in save states that would be playable in the browser, and it was suggested as a pipe dream by people we spoke with before it was presented as an actual option. For speedrunners a practice save state was ideal.Scholars expressed interest in some form of ROM validation, and in shareable citation links.

Platforms

Aytemiz (Designer/Scholar): AAA, PC, Consoles and Steam

Mazeika (Scholar/Speedrunner): SNES as a speedrunner, as a scholar Twine games, art games, walking simulators

Ross(Speedrunner): Primarily interested in SNES, NES, Genesis, and other retrogames, along with the occasional indie game like Cave Story (retro-inspired?), and some arcade.

Clemens (Speedrunner): Ideally everything, but prefers platforms that are popular with speedrunners specifically the SNES and N64, and other Nintendo and Playstation platforms

Swensen (Designer/Educator): Looking primarily for games close to the present (modern platforms), but also interested in finding prior examples from the last few decades.

Lowood (Curator/Historian): Mentioned MMOs as an area of interest likely to be a challenge.

Nathan Altice (Scholar platform studies): Although he wrote a book on the NES, indicated a strong interest in more obscure/less well studied platforms (e.g. "Bally Astrocade," arcade games, MAME?)

Summary: Interest in platforms varied. Scholars, educators and designers tended to want more modern examples. In contrast, speedrunners tended to be interested in more "retro" platforms, but even they indicated a preference for variety. Several people we spoke to brought up platforms of interest that they thought we could not effectively tackle (e.g. MMOs, strategy games). Our platform studies scholar expressed interest specifically in older/more obscure and less well studied platforms.

Other Suggestions or Ideas

Aytemiz (Designer/Scholar): Would like the system to take some of his play, and give him a similar player.

Ross(Speedrunner): Suggested linking our results to external resources, the way Google does with Wikipedia boxed.

Swensen (Designer/Educator): Suggested the value of the moment searching technology in a usability study, or for bug searching during development.

Lowood (Curator/Historian): Expressed an interest in finding a way to link a result to its specific context. Also suggested that if we are searching for something mechanical, we'd be interested in seeing things that share non-obvious structural similarities. Having a persistent url link would be valuable, but also means having a repository set up around the tool.

Individual Interviews

Notes from individual interviews follow, participant responses are bolded.

Name: **Nathan Altice**

Role: **Scholar Platform Studies**

Current Search Methods (Questions)

How do you search for games right now? How do you search within specific games right now? Do you instruct others in how to search for games? To search within them?

A system like this could have saved me hours, for any kind of deep critical look at any game. If I wanted to get to a particular point you might find a save state online, but often it would mean having to play up to that point. In very rare instances to I set up the hardware and time to get to a particular point and capture for class. Primarily using Youtube for Let's Plays, sometimes NikoNiko for Japanese videos.

Why was it worth doing for certain rare instances?

Either it was something very obscure and wasn't on Youtube, or the quality of what I could find was not great, e.g. a jackass narrating over a let's play covering up the audio, or just the quality of the video. It could also be a very specific sequence of play that I cannot find footage to replicate, e.g. rowboat sequence from Silent Hill 2, a way too long sequence of manually manipulating oars to reach a lighthouse.

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system? How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

It would be useful if you could do a kind of reverse lookup from a screenshot. From a historical research publication, it would be useful to be able to take a screenshot from a magazine and link it to a particular ROM. Say Nintendo was showing off super mario world, was this a screen that actually existed in the final game? Prototype or alternate version?

One of the ironies of teaching game design is that we have to rely on secondary sources, e.g. if I want to show a level in Metroid Prime I need to go find a Let's Play somewhere. I don't have time to play through the game, set up capture.

We don't have a setup to let people just bring a Gamecube to play in class. It would be really nice to have emulation where I can look up something to demonstrate in class not through video, but through play - even if that's a few years off.

If I were studying the musicology in games it would be nice to hear the music from level 4 without having to play all the way there.

It would be nice to just have a unified citational model - a notation that can take you to a point in a game, for academic clarity. There are all sorts of variations of ROMs

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

It immediately begs the question of how you decide on emulators, we're looking at examples like the SNES that are popular and well emulated. What if I wanted to use the Bally Astrocade? This is a platform that has a poor emulator, the SNES is pretty well studied. I know I wrote a book about the NES, but now I'm more interested in weird edge cases. Could I do this with arcade games/MAME? Would this encourage more research into emulation? What about DOS games? Steam? From a teaching perspective it would also be useful to do more current things. (e.g. it would be good to search for something from Mario and find similar examples from Super Meat Boy)

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader? For each of these possible query formats, please indicate if you would use it, how important it would be for your use, and how you would go about using it.

Can it take video? Could you also do a kind of genre search. For example, say I gave it a video of Mario, could it see "jumping from a platform" and detects that mechanism or action, getting me to the protoypical platformer. Say I could do that and then also ask for it to show me all examples before 1985.

Images makes the most sense, although if I'm trying to remember something it would be useful search for something like any games that use a particular palette. What if I wanted to find "mario on yoshi with 8 lives in an underwater level" for a screenshot for class? The more ways you can search the better.

If I was able to query by platform characteristics, because a lot of early Nintendo music was not attributed I could search for ways of manipulating the audio processor. E.g. "what used this byte flag in the audio processing unit" looking at a state across snapshots, maybe figuring out what games had the same composer.

For each of these possible results: Would you use this result? (Yes/No) How useful would this be to you, compared to other result formats? How you would use it in a professional (or hobbyist) capacity?

Getting a save state would be fantastic. Being able to play in the browser would be incredible, but Nintendo probably wouldn't allow it. Just getting a save state would be incredibly valuable. From a scholarly perspective the more info the better: palette memory, screen memory, what's in RAM. I'm sure speedrunners would love that kind of stuff. The more I could dig into platform specific stuff the better.

There's not ideal/canonical save state, but it will be important to follow the standards of the community (e.g. junk headers in the NES ROMs). For other platforms, there may not be something canonical. You may be creating it.

I wrote my own citation rules because "Mario Bros 1985" was not good enough. Was it played on emulator? Original hardware? Cartridge or disk? It would have been hypocritical if I said it was better, but I hadn't done it myself. I wouldn't know how to do this for Turbografx, or for an arcade board. The work might not have a good return, needing to know so much information for each platform, and some of that may not be known. This tool could write its own format with all the headers, a standard citation style.

If I was doing historical research on mechanics, I might want to look at all games that have bosses, or all games that have vertical scrolling, or a particular kind of jump. I might want to compare screen-to-screen vs games that did scrolling in 1985, across consoles and arcade for example. It would be great to get the concrete data. Another example, what games used the color blue as a background before Mario Bros.? This could let you do an aesthetic history in games studies without needing to know a comprehensive of games, getting the "mise en scène" .

The interface is really going to be key, think over the UI/UX experience, design is something that should be thought about even from the beginning, making a beautiful/easy tool that people will be interested in using.

Name: **Batu Aytemiz**

Role: (self-identify as one or more of our listed profiles) - **Game designer/scholar**

Game-related activities: What do you do with games? **Playing games, watching, thinking about how people design**

Current Search Methods (Questions). How do you search for games right now?

Google, primary source (steam), library, youtube lets plays,

How do you search within specific games right now?

Describe moment, hoping it's already been done, not-indexed? Let's play, or play oneself

Do you instruct others in how to search for games? To search within them?

No. A little, in a game design class. Gamespace tool for games (not moments), Game sage

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system?

Entering high level concepts, emotional tags (easy/hard)

How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?)

Games that have the high-level concept. Within the game, images of moments, ideal: point in a video that can be scrubbed, optimal: let me play the game, get me to that moment, export a save file, automatic citation for that moment, academic or link

Regardless of how you think it works, how should it work to specifically help you do what you want? What could you do with this?

Vectors tutorial image matching

What couldn't you do this that you would like to be able to?

Specifying high-level concepts

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

PC, AAA, Console, Steam

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader.

Image is bare minimum, save state is best

For each of these possible query formats, please indicate if you would use it, how important it would be for your use, and how you would go about using it. Would you use this query method? How useful would this be to you, compared to other query methods? If you indicated that you would use this kind of query format, please describe how you would use it in a professional (or hobbyist) capacity:

Screenshot: Providing an image of your own, or found using another search engine.

Probably not, if I have a screenshot I have other info (myself, or where I got the screenshot)

Collection of Screenshots: As above, but searching using a combination of screenshots instead of a single one (e.g. to narrow down to an event near them in time, or to look at related events across different contexts).

Depends on additional benefits, might not be needed

Text input. How would you like this system to handle text input? Are there specific things that you know it would need to be aware of (e.g. names of characters or actions within a game, on-screen text)?

Vocabulary of tags, e.g. first level v.s level one, tutorial, first level, credits,

Abstract sketch

Drawing a schematized and/or simplified version of a game screenshot. (Should this be purely visual? Should some drawing tools have semantic content (e.g. player, enemy, platform, etc.)?) Gimmick? Would have to try. Cool if it works, might not

Other. Is there another way (besides images or text) that you could imagine inputting a query? If so, how would it be useful to you?

No. Sound? Small video?

For each of these possible results, please indicate if you would use it, how important it would be for your use, and describe how you would use it. Would you use this result? (Yes/No) How useful would this be to you, compared to other result formats? If you

indicated that you would use this kind of result, please describe how you would use it in a professional (or hobbyist) capacity: Screenshot: This result would be an individual screenshot that matches your query.

Bare minimum

Animated GIF: A brief animation centered around the most relevant result.

More helpful

Save state: A save state that can be downloaded and then loaded into your own emulator. Please indicate the format(s) that you would find useful.

Ideal

Playable Save State: A save state that can be directly played from the search results, in an emulator loaded in the browser.

Ideal+

Exportable Citation: A citation for a moment that can either be cut and pasted into a paper, or one that can be exported into your own reference management software (e.g. Zotero, Endnote). If you would use this, please indicate your preferred reference format(s).

Not clear on what that is.

Button inputs: An illustration of input happening at each moment. Presumably paired with another output format.

Not useful. Wait no ML.

Scrubbable Timeline: A scrubbable timeline which would allow you to select points in a particular trace of play, like moving the slider widget in most audio/video playback systems. Each point would be labeled with its relevance to your query.

Would be useful. How about nonlinear games?

Switchable Timeline: Similar to a scrubbable timeline above, but also allowing you to select between different playtraces. These might be identified in various ways (e.g. human-generated, AI-generated, named speedrun, order-of-game-events description, tool-assisted).

More useful. Matching specific types of things. Seeing different playtraces to explore different areas. Linear outcomes.

Game Selection: The ability to narrow relevant results to a particular game, for any queries that might have matches across multiple games.

Yes. I like these, not those.

Batch Results: The ability to download a collection of results of a type previously specified.

Yes.

Other: Are there any other types of results or result features that you can imagine? If so, how would they useful to you?

No. Play the game a bit, and then ask the system for a similar player.

Name: **Bernd Clemens Huber (aka Aran Jaeger)**

Role: (self-identify as one or more of our listed profiles)

I don't speedrun so much myself really, at least not in the usual sense, but I work on TASes and research games, check for exploits, glitches, and finding the meaning of memory addresses for a few games (mostly SNES)

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)?

regarding Super Metroid, I've gone through about all the ~450 or so hacks of it, and helped "Sniq" (some SM TASer, or well that is 1 of the nicknames he uses) for the recent Low% and Any% TASes of SM.

In the SM community I'm more known as "ED" which stands for an old maybe silly/childish name "eterniseddragon" which I did stick to.

Regarding S(uper) M(etroid), most of my contributions can be found in the wiki that we have for the game: https://wiki.supermetroid.run/Main_Page

In here it would be that section: https://wiki.supermetroid.run/ED%27s_stuff (I generally like to route TASes for SM hacks, because it is nice to find applications of exploits/glitches that have been found in/for the game)

Main Page

ED%27s_stuff

I'd say I'm an expert regarding knowing all kinds of things for this game and knowing the people that know other things about the game

Current Search Methods (Questions)

How do you search for games right now? How do you search within specific games right now?

I don't look around so much for newer games really, I'm more sticking to "classical ones" from older generations, but for these it is hard to say which particular screenshot of it or moment I would look at or remember. Usually when I find a promising name of a new game or someone tells me that I might be interested in some game, then I search for some video of it and click through different parts of the video to get an impression of it.

if it would be for a new 2D platformer game to maybe work on a TAS for it, I check if there already exists such tool-assisted speedrun/superplay, and if that doesn't exist, the next I'd search for would be speedruns by human players for the game, and if those don't exist either, there's usually some sort of longplays or casual playthroughs of them. That's also my general method of looking for new games.

Do you instruct others in how to search for games? To search within them?

Well what usually happens in the SM speedrunning community is that either a person comes up with some problem (can be a game crash or getting stuck, or very bad luck that happens rarely or other rare events that from time to time occur to players) and presents some short video or clip of it, and then in some cases I know where similar stucks/crashes or rare events (usually induced by the game's RNG) occur in other clips or videos. And the other case would be that someone is asking

for a trick or demonstration of how to do something. I think those are the most frequent ones in the speedrunning community in Discord.

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system?

How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

If someone would be searching for some moment or short time event in some video game that is related or close to a possibly existing storyline or plot in the game, then I assume it'd be easier to get good search results for such compared to if someone is just searching for some "side quest" or optional part that can but doesn't have to be done in some game, so if I'd know that what I'm searching for is close to or near some more "central event" that is featured in the game, then I probably would refer to this and then go from there.

For more "obscure" or hidden or more rare scenarios that can happen in games, it would be more complicated I think, and if the thing (trick/demonstration/...) has been given some name previously by who ever found it first or to whom ever occured this event or scenario (or in an extreme case I guess maybe some reduction of a short-time event or scenario to 1 moment part of it that basically "stands for the entire scenario"), then I think I would stick to whatever phrase it

has been given and maybe add the general area in which it happens or whatever circumstances make it special from "the usual cases" in which similar scenarios happen but maybe without the thing that is special in the given case that I'm looking for.

Capabilities. What could you do with this? What couldn't you do this that you would like to be able to?

Regarding a use of this tool, it probably would allow to find more "hidden" or usually hard to find screenshots that encompass special types of situation in games faster than with a usual search. I understand the aspect of including memory address values which change all the time in the background like hidden/mysterious numbers and for all kinds of things. I guess there'd be 2 cases of searches or search engines, depending on if one has memory address info that one can carry into a query or not (when searching for something that might be similar).

Queries and Results

What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

I mean the general most greedy answer would be "everything/all", but in a sense of trying to cover a lot while not having to put too much effort into it, keeping it reasonable.... probably a subset of the platforms/systems such that any other system isn't far off from the set of chosen ones.

I don't think I can speak for a general audience (even if one would narrow it down to speedrunners/TASers, because they focus on different platforms that are "all over the place"), but I guess one would then prioritize on which platforms or systems are most frequently used or those that gather the largest amount of interest, like some really old systems, commodore or so and the like, there might not be as much interest by so many people in these currently or anymore.

I guess there'd be 3 classes maybe into which one could sort in systems. really old ones that maybe also don't cover so many games and were not as platforms or systems so "competent" or powerful yet, then some "classical ones" in the middle, some might even exaggerate it and call it "the golden age of video gaming" probably mostly since they know these from childhood and nostalgia playing a factor and then the last case of more recent systems which of course usually are where a current interest lies in and people being interested in them for a while shortly after they appear or every time a new game appears, like some "hype" wave that then at some time goes back to a more normal state but on the other hand side, there's not as much content yet to be expected for newer systems so I'd think for the central or middle systems one might be focusing on these more so not tooooo old but also not yet very recent systems, but that is hard to say... and I haven't had so much time yet to think about these questions

I mean.. a search becomes more difficult with the amount of results one can get, so if something is really new, the number of results shouldnt be too big yet

For specific platforms, to name some... I guess the SNES, N64, hmm I'm not so familiar with systems around that same time as these were made or where the focus was on these by Nintendo, but were systems that were made by other companies. I guess some playstation system would be fitting into this time period as well, probably more or maybe instead of N64 rather gamecube, but i'm not sure how many games there are for each system.

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader?

For each of these possible query formats, please indicate if you would use it, how important it would be for your use, and how you would go about using it.

Regarding that list above and some priority list with respect to what type of content I would like or want to be able to extract from results, I'd think I'd for the moment list them as follows, with the first mentions having higher priorities than later ones: (play in browser) links, savestate, citation, button inputs, animated GIF, screenshot

info on emulator versions used because if the version differs, savestates might not work or movies desync (is something that comes to mind)

for a given game and game moments, people that research such a game probably would be more interested in the very edge cases of moments, so those that are rare and maybe harder to reproduce(edited) - maybe some list of rarest moments, those that are furthest away from the usual ones?

maybe info on how old some content is... in case one already has some overview or already did some search longer ago and would only be more interested in what new content there is

for button inputs as something to extract, these usually correspond to them having been made in some specific emulator and the same button inputs could not work on other emulators or might not lead to the wanted trace, unless the emulator and or version would be part of a movie file / button inputs

Name: **Nick Junius**

Role: (self-identify as one or more of our listed profiles)

Live analysis on streams, video essays, blog posts, modding

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)?

Textual analysis, close reading of freespace series, discussion Valhalla's worldbuilding through its systems, Developler for 3d bullethell game Enure (see itch.io), Currently streaming ME

Current Search Methods (Questions). How do you search for games right now? How do you search within specific games right now Do you instruct others in how to search for games? To search within them?

Normally to make a video, I know my game/playthrough very well, record myself playing, to get specific moments. Much more rarely finding someone else doing it, usually doing it myself (e.g. dark souls). Might feel more honest to get it yourself, to play.

For games, it's easier to focus on one example (a single game), you can go more in-depth. It's harder to get people to look at, and it's more of a balancing act.

Most of my feedback has been about reframing moments. In a couple of places I can see how it's given people new in the community an easier or faster way. The primary audience is people who are very familiar with it an still working with it, or people who have fond memories. This is my kind of work, more about recontextualizing or sharing unusual moments.

How long do I want to spend, how much do I already know, how much can I add?

Speculation

We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system? How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

Ideally give it video, up to 10 seconds? Or give it an image, or try and describe it with text. When I've used Google to try and find out if a game had a system, it's hard to figure out. It's hard to find what you want, just Googling "game that has x, y, and z." Being able to load a state from a giant database of game save states. Save files are more about inventory states or flags being tripped. Having the ability to look at more detail than what's been marked down in the database. Parts of a game, similarity, what led me here, what happens after, I don't want to scrub through a let's play. Given my gif, what led to it, what follows. An animation of a bunch of pngs of screenshots. A text document for things, other aspects of the state, buttons that are being pressed, parts of the player state that aren't obvious, finding the causes of weird glitches

If you know a game's memory can you map it to human readable

Capabilities. What could you do with this? What couldn't you do this that you would like to be able to?

Useful as a reference guide. I'm falling back more on my own knowledge. Find changes like change X alone would be useful for both analytical and development standpoint. The system can find similar moments, humans can see differences - avoid repetition. The ability to work with things that are not just text.

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

It needs to be able to handle any kind of PC game that you throw at it. One other reason to go on to PC or more PC-like, you'll still get simple things early. Getting Doom working should make Baldur's gate easier. As far as specific kinds of games, the longer it is the more useful. Witcher 3, moment searching would be incredibly helpful b/c it's open world, finding something that I might never see. On other really long things - any game you can import a save file into. One of the problems with playing through ME, exploring branching states - playing through 3 50 hour games and tweaking it. The ability to search for specific complex states outside of the game would be great.

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader?

Seeing what possible ways there are to get to the current viewed moment. On the practical side, sets of screenshots would be the most useful, or a small video or gif. Probably the most useful from an analytical content creation standpoint. I'd prefer an archive folder of pngs, but for random people a gif or video.

When I wanted to talk about an even ⅔ or ¾ through ME3, and I didn't have a save file and couldn't find one that I wanted. It took me about 15-20 hours or play to get 3-4 minutes of footage. In this specific case, I wanted the ME3 main story

mission on the planet with the boss fight with kie lang where he throws you in a hole. Give me moments from ME3 priority mission Thessia, and I could pare it down more to boss fight, pre-boss-fight conversation.

From a dev perspective, there are a lot of old PC games that are hard to get running or not accessible. Having the ability to flip through a game would be helpful. Some people on the team for a game I worked on (Inure) did not have as much of reference pool, could have been useful to get people on the same page.

ME Save states - a model of how people explore reachable paths, share their exploration, could our tools participate in these ecosystems, share and contribute

How long is a moment? Exploring every quarter second, all 60 frames, but so little is changing

Name: **Ted Lim**

Role: (self-identify as one or more of our listed profiles**) designer at netease, mobile game developer, feature design documents, user tests, player feedback, interaction**

Player of multiplayer games, rpgs, pokemon, fire emblem, f2p mobile games,

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)? - **yugioh competitive, decklists, testing**

Current Search Methods (Questions)

How do you search for games right now? How do you search within specific games right now?

Do you instruct others in how to search for games? To search within them?

Looking at the most popular games, based on word of mouth, twitch, following specific brands

Within games, looking for moments to review mistakes (e.g. starcraft replay systems), but I have a different perspective as a dev. What we lack is getting to a specific point. For a lot of mobile egames, figuring out change in player lifecycle is important, a fresh or long term player will have a different ideas of what is and isn't a good deal. Finding moments would be more useful for making a better experience. In a MOBA you are looking for a specific moment in a specific game. You are trying to optimize for a specific in f2p, planning for or easing pain points. Player behavior and mentality is different when someone is a new player. The experience at each level of the game is different. What would be useful would be being able to skip to a moment to understand a player's experience. What can we offer e.g. as a promotional item.

Speculation

We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system? How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

One of the most important things is making the game fun. If we had user testing data and audio, could we find a moment with a shouted exclamation linked to the

moment in the game. More mechanically, once your player reaches a certain level, we're maybe only interested in actions beyond this point. Want to observe a player's reactions and experience. This could be by level, currency, strength of a character. If I was looking at a less familiar or competitor's game e.g. fortnite, images, streams, promos, if I want to know what makes it fun I'd want to know view counts, user watch count, favorites, looking at streamer revenue, images only capture what's in the game but we want to know if it's actually relevant for the person playing it. Using image as a search tool might not get the data, might only find a type of branding or art style, e.g. silly vs. serious but not capturing the player reaction. You want to know that you are hitting the right audience. For results what's most useful is behavior going forward from a specific moment, the primary pillar is always going to be is the game fun. Any part of the game that isn't fun needs to be cut/improved, and that's based on player reaction (e.g. balancing is subjective). You are always looking for the most balanced gameplay, but fun is more important. For f2p this is for the largest possible audience, maybe less important for other kinds of games.

Capabilities. What could you do with this? What couldn't you do this that you would like to be able to?

What it sounds like is finding a moment in gameplay, and getting to that moment. If I could drop an image of a screenshot Clash Royal, would I need the game installed somewhere. This might be useful to get into what the competitor

company and audience is thinking. How I would imagine this working for research on competitors would be comparing screenshots from the app store screenshots, to find a comparison between those "best" images and the state of the game they are capturing to other screenshots that are available (youtube, twitch). Comparing the competitor's perception of their own game to the crowdsourced reaction. For a live game, looking at gameplay to compare beliefs of the competitor to the audience reactions. This might give a better insight, more useful for a live game than one in development.

You would try to find moments that people are enjoying, then you would need human analysis in order to look at these moments to try to mimic it in a different game (e.g. "fast follows" just copy what they see visually).

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

In the mobile space it would be useful for mobile games, mobile apps, twitch streams, even chat reactions, youtube highlights.

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful?

Ideally a bunch of information about the player, reactions, and behavior, soundbites if possible. It would also be useful to have game states for testing reactions to pain points and paywalls.

Name: **Henry Lowood**

Role: (self-identify as one or more of our listed profiles) **Historian Curator**

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)?

Look at what players do in games, replays, superior athletic play, historical/documented, youtube, forum to another forum, links to videos, warcraft machinima

Curatorial side - internet archive, adding description, tagging, metadata

Tool for capturing replays

Current Search Methods (Questions). How do you search for games right now? How do you search within specific games right now?

Context dependent discovery, different communities will look for different things (speedrun/e-sports/let's play) they talk about things differently

For example: speedruns at internet archive, which are searchable now

There hasn't been a let's play archive similar to the speedrun one in terms of the amount of care and attention that's gone into curating it.

Do you instruct others in how to search for games? To search within them?

A little bit. Some presentations within the library community (more than in the context of instructor/student). Discovery - if you do a search, you wouldn't know from the result that the doom doc was posted in a forum, embedded in a discussion, looks like a book on a shelf, decontextualized. I'm interested less in a known item search, more about browsing through forums, and how can you get more of the context

Search is not a neutral tool, but we often talk about it as if it were. Knowing the context is more important when you are first introduced to something. Once you already know it, then it's more ok that you don't get it from the search.

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system?

Several - for me, most interested in description of known event (e.g. in a virtual world, the "gnome demonstration") In a competitive game, the moment when Grubby when used x-move to win the world cybergames. How do you get this? There's normally no replay file - although in this case I have it, because I was the referee. Others might want to look for "the rocket jump" or the "Super Mario zero [minus] world glitch." Either a known event, or a particular game mechanic.

I could thing of some far-fetched things with code. 99% of search would be text though.

How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?)

Video or screenshots for an event. For a mechanic, it might be nice to have a replay format/memory state to put yourself in the game. A GISST replay.

Regardless of how you think it works, how should it work to specifically help you do what you want?

It would be interesting if you could somehow go from source to rendering from code. Then we could ask: What does this code do? What does it enable? What about if you change the code while the game is running?

Capabilities. What could you do with this? What couldn't you do this that you would like to be able to?

What happens when you have multiple versions of a game?

Right away - I have big corpus - given a screenshot of someone in zero world, but I don't know much about how it's actually happened - what would the matches around the screenshot show? Having a save state is also possible, more generally - "here's this event" that's well known, but I don't know what game it was actually done in, or if this screenshot is of something that could have actually happened? Is this a real screenshot that matches what's possible in the game? It could also show that something exists only in a hacked version.

It would be useful to see your query image while you are looking at the results.

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

MMOs would be challenging. Theoretically the screenshot could be matched to the location, but no memory states/map. Could match to sequences in a replay file, but that might be very game specific. Maybe taking the rocket jump as a query, by dropping in the rocket jump from memory in one game. Maybe you find something that's similar in the system, but didn't just get things that you already knew were rocket jumps.

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader?

GISST, a persistent url link, which means having a repository around the tool. Using the tool, being able to create some sequences that can take people back into the tool to play with it. Screenshots and videos are how things used to be documented, but you might need to authenticate it. Save states, replays are better than video, but not available for many types of games. Talk to Eric Kaltman, if you haven't already, I'm sure he'd have useful suggestions.

Name: **Stacey Mason**

Role: **Scholar, (academic) Streamer, Developer**

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)?

Person on the Internet - Social media user, blogger, (doesn't like this term, but..) "influencer"

Current Search Methods (Questions). How do you search for games right now? How do you search within specific games right now? Do you instruct others in how to search for games? To search within them?

I would play through a game to get it to a state I want, and have saves ready to go. There is no good academic citation system except for Eric's tool which is very narrow. Other kinds of citations might mean citing a collection of screenshots or a recording on youtube.

Text would be the most obvious way to search, also for mentally keeping track of what I'm looking for. For text I don't need to have a mental map of the game. If I need to search for where a particular item is, search is most useful if I don't already know. In that case text search seems like the best bet. For a game like Detroit, a plot map (a tree map of the plot) would be useful. The easier and dumber you could make it the better.

The way people search through games is going to vary between narrative games like Detroit, and simulation games like Sims, Starcraft. The data will have to organized differently for these different types of games.

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system? How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

As a developer I would use this the least, compared to as a streamer or academic. Maybe in QA, I could see a quick capture and sending of state could be useful for sharing an already found bug.

The most common use of a moment for me would be as an academic, I need this moment to support an argument. As a streamer, my work is an extension of my use as an academic. When I am streaming I am [using examples from a game] and making some claims (unusual for a streamer). That's probably for true for [my] use cases on twitter or other social media. As a streamer, I might want to "Julia Childs" a playthrough (i.e. have it ready to go). On the other hand, viewers would want to see how we got to the point. But in some cases we wanted to jump from level 2 to level 40 and for practical reasons it might be better to use something like this in a separate stream.

Capabilities. What could you do with this? What couldn't you do this that you would like to be able to?

Searching by location in a moment in the game, including details like level, spatial and chronological locations, then adding parameters. Could be useful as a way to

play through variations, particularly for story games to look at variations on a particular moment.

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

I do narrative games for all of my things. For games like narrative games like Detroit, or Dark Souls. My character at a point in the game with particular set of items or with a particular set of bosses killed, etc.

I think the ultimate win condition for the system is a complicated game with lots of features. If you can do this for something like Overwatch, there are so many variation and metameaning in those variation. It would be useful to capture that.

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful?

What I'd most want is a parameterized save state, which is easier to get than by playing myself (also interested in a short video/gif). A super ideal cases would be having some integration with the game itself. If I could hit a button in the game to pull up and share a moment. As a consumer streamer, I would want this, but as a developer of narrative/story games I would hate this thing existing. I don't want my players to manipulate games to an exact configuration. This could be a problem for more progression based games as well.

Name: **Stella Mazeika**

Role: (self-identify as one or more of our listed profiles) **Scholar Speedrunner**

Game-related activities: What do you do with games?

Scholarsplay, practicing speedrun techniques

Current Search Methods (Questions). How do you search for games right now?

google

How do you search within specific games right now?

youtube , twitch, manual search (playing)

Do you instruct others in how to search for games? To search within them?

no

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system?

Love to explore branching paths as a scholar, see variations on choice, John Murray's telltale analysis, narrowing down points in speedrun for glitches, slow step by step, frame by frame input loop

Input clarification: look at games in parallel, metaphorical, speedrun search by glitch/game name, specifying section within game, e.g "ice palace from a link to the past bomb jump", free text, dropdowns

Capabilities. What could you do with this? What couldn't you do this that you would like to be able to?

Upload multiple runs, see differences moment to moment in one game vs. another, extract from videos on twitch, compare between different runs

Adding text search for name of the game

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

SNES already good (for rpg speedrun gauntlet practice, if it came out tomorrow), as an academic twine games, walking sims, other artistic value games, what about strategy? (civ/starcraft)

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader.

Controllable video - save state, augment animation/video with button presses, ability to slowdown or speed up video, might need ten seconds but all 600 frames

Academic - def a citation, as many formats as possible (no idea what's needed most)

Controllable save state, I probably already know the moments I want - e.g. branching choice screenshot, show both outcomes

Name: **Alexis "Protomagicalgirl" Ross**

Role: **Speedrunner**

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)?

Speedrunning and event production (e.g. looking into what a runner is offering).

Current Search Methods (Questions). How do you search for games right now? How do you search within specific games right now? Do you instruct others in how to search for games? To search within them?

Start with Speedrun.com, for route notes, walkthrough, look at the resources tab, or just looking at a run. Telling other people to use the same page.

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system? How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

Search by image would be helpful. (e.g. frame perfect trick to kill Castlevania 1 boss in one hit), I'd like to just drop in a screen shot. Searching by text is also useful (e.g. speedrun route for "Paper Mario Thousand Year Door speedrun route") If I wanted to know "paper mario thousand year door teleporter room early" it would be cool to see clips or a video of that trick.

It would be nice to have this work like Google does when you search for wikipedia, e.g. a little box for the trick, and the rest of the results below that.

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you?

SNES, NES, Genesis, retrogames, the occasional indie game like Cave Story, Arcade stuff if possible (e.g. "Tetris the Grandmaster")

How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader?

There are some game specific tools, e.g. Fumen for tetris. Example Fumen for the named move "": http://harddrop.com/fumen/?v115@+gC8GeA8AeA8GeA8CeB8CeI8KeWWevhCOmf2qBAAA

Could we take tetris screenshots and "Fumenize" them? What about PuyoPuyo or Tetris Attack. The Fumen above is the move named "Babe come to column five" aka the best tetris move

For each of these possible query formats, please indicate if you would use it, how important it would be for your use, and how you would go about using it.

Searching by screenshots would be a high priority, over text , because if I know what something is called Google has probably already got it. What I think I would

get from this is it being smarter about how videogames work, to take that leap from screenshot to a specific moment in the game in the game's own language.

Just being able to look up a trick and find a practice save would be cool. As a college student who takes every excuse she can get to write about nerdy stuff having citations would also be cool. It would also be useful to be able to filter the results based on the source, e.g. if I want a speedrunning trick I do not want results from a casual playthrough or a review of the game.

If you're searching for a particular game, it would also be nice to get some Twitch directory results, since it would be live right now.

Name: **Elizabeth Swensen**

Role: (self-identify as one or more of our listed profiles) **Educator Designer**

Game-related activities (open-ended question): What do you do with games (professionally, or as a hobbyist)?

I design games for my own artistic practices, for clients, and I teach game design. I am also a player.

Current Search Methods (Questions). How do you search for games right now? How do you search within specific games right now? Do you instruct others in how to search for games? To search within them?

Casual Internet searches. I've gotten good at using search terms even if there isn't an existing code for what that features or mechanic or system dynamic is. It is a

little sloppy. When I feel particularly stuck I will reach out to my social network of game people. From casual acquaintances who are more clued into Youtube/Let's Play culture, to people who have been playing games for 50+ years. These are in-person conversations, I don't send blasts via social media, and I usually already know who I need to talk to.

Research as a part of brainstorming is a part of what we teach in intro classes, but we don't spend as much time with it as we probably should. They probably use a basic search engine. I do guide people beyond the basics, but that's usually based on my own knowledge base.

Speculation. We are making a videogame moment search engine which indexes the content inside of a game. How do you think people provide the search query in this kind of system? How do you think the results are displayed? (What's in a result listing? How are they sorted? What happens when you click on one?) Regardless of how you think it works, how should it work to specifically help you do what you want?

I can give you an example of a time I needed to show a student a moment, what the current systems allowed me to do and what I wish it could, that I don't believe it could. I wanted to have a conversation about doing text-based input with a controller, related to some specific moments in the game "Beyond Good and Evil." They have a wonderful satisfying text-input for joysticks. What I do now is Youtube is desperately, and luckily there are people as obsessed with it as I am - or I can find a longer video and scrub for the moment. Unfortunately, what's

amazing there is not what you see but what you feel, and how seamless it is. In a perfect world, I could show them the moment in a way that I could hand them the controller, in a dream world. You can't just show something, sometimes you need people to feel it.

I could imagine there was a language created about mechanics that could be searchable across games. If a student wanted to know about text-input with a controller, and then could get visual evidence, ideally video. It would be nice to have a reference of how to get to that moment (since timestamps won't work), but would that mean that every reference would need to come with a walkthrough?

Demo (Priming) -Tech demo segment and Back-end explainer

It seems like this technology could be very useful in usability. If I'm developing the next Overwatch, I could take a screenshot of a bug and then find all the screenshots that match.

Queries and Results. What specific platforms (or even specific games for those platforms) would you need the search engine to handle to be useful for you? How much to you need to be able to extract from the results (a link, a screenshot, a downloadable save state from which you can continue play, a GISST format citation, another citation format, an animated GIF, button inputs, scrubbable/switchable game timeline) in order for the system to be useful? What do you need in order to share the result with a friend/colleague/reader?

I often find myself reaching for particular mechanics and looking for priors in that space. I create more things than I publish, so I'm less concerned with citations. I'm interested in games that for example, have an evolving morality system. I'd like to find more outside of my social circle or outside of my own experience. I'd want to know how they work, how they were received, and how it compares to what I (or a student) is considering. This doesn't have to be historical, it is more about a constraint of personal attention. I only know what I see in my glimpse of media, I'd like to see a wider scope, I'd like to find things closer to the present and then move back to the 80s/90s. I'm often lecturing or thinking about specific mechanics (e.g morality system, a particular jump, an action drafting system, etc.). I don't just want xyz mechanics, but also the moment where that mechanics exist, it would be good to find history and modern examples.