

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Accelerated Fuel Qualification of Uranium Mononitride: Mechanistic Modeling, Bayesian Calibration, and Regime-Stratified Validation

Permalink

<https://escholarship.org/uc/item/5t55p5kx>

ISBN

9798276020013

Author

Miller, Zachary Aaron

Publication Date

2025-12-06

Peer reviewed|Thesis/dissertation

Accelerated Fuel Qualification of Uranium Mononitride: Mechanistic Modeling, Bayesian
Calibration, and Regime-Stratified Validation

by

Zachary Aaron Miller

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Nuclear Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Massimiliano Fratoni, Chair

Professor Peter Hosemann

Dr. Alex Levinsky

Fall 2025

Accelerated Fuel Qualification of Uranium Mononitride: Mechanistic Modeling, Bayesian
Calibration, and Regime-Stratified Validation

Copyright 2025
by
Zachary Aaron Miller

Abstract

Accelerated Fuel Qualification of Uranium Mononitride: Mechanistic Modeling, Bayesian Calibration, and Regime-Stratified Validation

by

Zachary Aaron Miller

Doctor of Philosophy in Engineering – Nuclear Engineering

University of California, Berkeley

Professor Massimiliano Fratoni, Chair

Accelerated fuel qualification for advanced nuclear systems requires mechanistic understanding coupled with rigorous uncertainty quantification to compress traditional timelines. This dissertation investigates whether mechanistic models grounded in first-principles physics, systematically validated against experimental data across multiple operating regimes, and calibrated through Bayesian inference can substantially reduce epistemic uncertainty in fuel performance predictions while establishing defensible operational boundaries. The work applies this hypothesis to uranium mononitride (UN) fuel for micro-reactor deployment, addressing the challenge that historical UN irradiation data exhibit significant scatter and gaps at high burnup and elevated temperatures relevant to modern micro-reactor designs.

A mechanistic model for fission gas release and swelling model was investigated in the BISON fuel performance code. The model explicitly captures temperature-dependent diffusion mechanisms through decomposition of effective diffusivity into thermal equilibrium, defect-assisted, and thermal components. Multi-population bubble tracking differentiates nucleation pathways, and grain boundary coalescence provides mechanistic basis for understanding saturation transitions and release. Validation against various historical irradiation test cases revealed regime-dependent model performance with robust agreement in low-temperature, low-burnup conditions and interpretable limitations at regime boundaries.

The dissertation establishes that mechanistic modeling with rigorous uncertainty quantification may successfully enable compressed accelerated fuel qualification timelines. Broader methodological contributions include regime-stratified validation identifying mechanistic regime boundaries based on physics rather than arbitrary thresholds, surrogate-accelerated Bayesian inference achieving significant computational speedup while maintaining statistical rigor through hyperparameter optimization and empirical calibration inflation factors, and mechanistically guided prior specification improving parameter convergence by 30 to

50% while maintaining Bayesian validity. These approaches are generalizable to other high-dimensional inverse uncertainty quantification problems for expensive mechanistic models.

The surrogate-accelerated Bayesian Markov chain Monte Carlo (MCMC), using latin hypercube sampling (LHS)-sampled input parameters to train Gaussian process emulators, achieved robust parameter within the primary regime for low temperature and low burnup. Parameter uncertainty was reduced by 80 to 90% through Bayesian inference, with posterior distributions full convergence (Gelman-Rubin $\hat{R} < 1.01$). Forward propagation of calibrated posteriors produced predictions with 86.4% experimental coverage, 0.30 vol% mean absolute error for swelling, and 0.05% mean absolute error for fission gas release. Comparison of prior (uninformed) versus posterior predictions quantified information gain, with swelling prediction accuracy improving 78.6% and fission gas release accuracy improving 92.6%.

To my friends and family—the journey began and will end with you

Acknowledgments

This journey began with a decision born equally of spontaneity and fate in my final year at the University of Arizona, when stalled job applications led me to a US Navy recruiting booth. After raising my right hand and committing to serve as a submarine officer in the world's greatest Navy, the next decade was set in motion which would take me places I never imagined, both literally (the bottom of the ocean) and figuratively (a PhD in nuclear engineering).

The submarine force trained me, challenged me, broke me and eventually shaped me into the engineer I am today. It sent me on deployments across the globe, put me through countless hours of watch rotations, and introduced me to some of the most intelligent, resourceful, and hilariously sarcastic people I've ever known. Somewhere between monitoring reactor parameters at 2 AM and the unique camaraderie that comes from living in a steel tube underwater, I discovered something unexpected: I wanted more. Not the submarines, necessarily, but the nuclear technology. I wanted to go back to school and contribute to advancing the field.

After leaving the Navy, I brought that passion to UC Berkeley's Master of Engineering program, where I had the great fortune of meeting Professor Massimiliano Fratoni. When he asked if I would be interested in joining his PhD group and conducting research at Los Alamos National Laboratory, I said yes without hesitation. Three years later, I'm reflecting on this improbable path from that career fair booth to this dissertation, and all the extraordinary people who helped me along the way.

I want to thank you, Max, for believing in me and giving me this opportunity. Your mentorship and expertise challenged me to be the best engineer I can be. I would not have this opportunity without you. To my dissertation committee, Professor Peter Hosemann and Dr. Alex Levinsky, thank you for your wisdom, feedback and overall support.

Thank you to Los Alamos for the funding that supported this work. To the scientists at Los Alamos that supported me and mentored me along the way, thank you. Alex and Galen, thank you for your persistent belief in me, to push me to be the best version of myself, will have a lasting impact. Both of you have challenged me personally, professionally and academically to become the person I am today. I will always be grateful for your continued mentorship, even after my time at Los Alamos technically was completed.

To all the friends I have met along the way, thank you. My time in the Bay Area has been nothing short of amazing. Friends have always been my backbone and you all have been there to support me, through the ups and downs. I would not have been able to get to this point without you.

To my family, thank you. You are all my biggest cheerleaders. To mom and dad, thank you for always supporting me and shaping me into the person I am today. To you Tay, YDID. To Ben and Ashley, thank you for sharing your home and transforming me through love and compassion. To Justin, you're a brother to me and are always there to give me real and honest support. To the Winslows, thank you for believing in me and showering me with love. And to the Alwins, your love, support and belief in me has been nothing short of a blessing. Thank you to my family, I love you all.

And to you, my love. I will forever be grateful for your unconditional love and support every step of the way. Your overwhelming passion and interest in my work, even though you may only understand 1% of what I do, means the world to me. Your excitement to brag about me and the pride you take in being part of this journey has never gone unnoticed. Thank you for always believing in me, standing beside me through the late nights, the setbacks, and the moments of doubt. Your steady presence made all the difference.

To everyone who believed in this journey, thank you.

Table of Contents

Table of Contents	iv
List of Figures	xi
List of Tables	xiv
List of Acronyms	xv
1 Introduction	1
1.1 Motivation	1
1.2 Accelerated Fuel Qualification Framework	2
1.2.1 Historical Context and Necessity	2
1.2.2 Core Philosophy and Principles	2
1.2.3 Multiscale Modeling and Simulation	3
1.2.4 Uncertainty Quantification and Bayesian Inference	4
1.2.5 Targeted Experimental Validation	4
1.2.6 Application to Uranium Mononitride Fuel	5
1.3 Uranium Mononitride as an Advanced Reactor Fuel	5
1.3.1 Material Characteristics	5
1.3.2 Historical Performance Data	6
1.3.3 Data Gaps and Knowledge Limitations	7
1.4 Fission Gas Release and Swelling as Performance-Limiting Phenomena	7
1.4.1 Fission Gas Generation and Release	8
1.4.2 Fuel Swelling	8
1.5 Dissertation Outline	9
1.5.1 Dissertation Objectives	9
1.5.2 Dissertation Structure	9
2 UN Material Properties and Correlations	11
2.1 Methodology: Material Property Compilation and Correlation Development	11
2.1.1 Overview and Computational Context	11
2.1.2 Experimental Data Collection and Compilation	12
2.1.3 Correlation Development Approach	13

2.2	Mechanical Properties	14
2.2.1	Overview and Computational Role in Fuel Performance	14
2.2.2	Temperature-Dependent Elastic Properties	14
2.2.3	Young's Modulus	15
2.2.4	Poisson's Ratio	17
2.2.5	Shear Modulus	18
2.2.6	Bulk Modulus	19
2.2.7	Thermal Creep	20
2.2.8	Coefficient of Thermal Expansion	22
2.2.9	Uranium and Nitrogen Self-Diffusivity	24
2.2.9.1	Nitrogen Self-Diffusivity	25
2.2.9.2	Uranium Self-Diffusivity	25
2.3	Thermal Properties	26
2.3.1	Overview and Computational Role	26
2.3.2	Thermal Conductivity	27
2.3.2.1	Thermal Transport Physics in UN: Electronic vs. Phononic Contributions	27
2.3.2.2	Data Compilation and Porosity Effects	28
2.3.2.3	Thermal Conductivity Correlation	30
2.3.3	Specific Heat Capacity	31
2.3.3.1	Heat Capacity Theory and Temperature Dependence	31
2.3.3.2	Experimental Data and Computational Validation	32
2.3.3.3	Heat Capacity Correlation	33
2.4	Summary and Integration with Irradiation Physics	33
3	Fuel Performance Under Neutron Irradiation	35
3.1	Overview and Computational Role	35
3.2	Fundamental Mechanisms Governing Irradiation-Induced Fuel Performance	36
3.2.1	Fission Gas Generation and Transport	36
3.2.2	Fuel Swelling Mechanisms and Saturation Phenomena	36
3.3	Historical Irradiation Testing Programs and Key Findings	37
3.3.1	BR-10: Fast Reactor Testing	37
3.3.2	SNAP-50: Space Reactor Program	39
3.3.3	NASA Advanced Power Reactor Program	41
3.3.4	Battelle Research Reactor High-Temperature Program	41
3.3.4.1	Saturation Swelling Discovery	42
3.3.4.2	Parametric Investigation of Saturation Mechanisms	42
3.3.4.3	Battelle Synthesis	44
3.3.5	SP-100 Program	44
3.4	Dominant Trends Across the Historical Database	45
3.4.1	Temperature and Power Density Dependence	45
3.4.2	Program-Specific Trends and Microstructural Effects	46

3.5	Empirical Correlations, Limitations, and Mechanistic Necessity	48
3.5.1	Ross Swelling Correlation	48
3.5.2	Machine Learning Swelling Models	50
3.5.3	Impact of Mechanistic Modeling	51
4	UN SIFGRS Model	52
4.1	Introduction and Mechanistic Model Motivation	52
4.2	Mechanistic Model Framework: SIFGRS Overview and Bubble Population Tracking	53
4.2.1	Fuel Performance Code	53
4.2.2	Conceptual Framework	54
4.2.3	Evolution Equations and Rate Processes	55
4.3	Fission Gas Diffusion: Three Mechanistic Regimes	55
4.3.1	Thermal Equilibrium Diffusivity: D_1	55
4.3.2	Irradiation-Enhanced Diffusivity: D_2	56
4.3.3	Irradiation-Induced Athermal Diffusivity: D_3	57
4.3.4	Total Diffusivity and Transition Regimes	57
4.4	Bubble Nucleation and Growth Kinetics	58
4.4.1	Homogeneous Nucleation: Bulk Bubbles	58
4.4.2	Heterogeneous Nucleation: Dislocation Bubbles	59
4.4.3	Bubble Coalescence and Mean Size Evolution	60
4.5	Resolution of Fission Gas from Bubbles	60
4.6	Grain Boundary Bubble Evolution and Interconnection-Driven Release	61
4.6.1	Migration Toward and Accumulation at Grain Boundaries	61
4.6.2	Saturation, Interconnection, and Gas Release	62
4.7	Model Parameters and Uncertainty Sources	63
4.7.1	Lower-Length-Scale Calculations	63
4.7.2	Empirical Parameters	64
4.7.3	Primary Uncertainty Sources	64
4.8	Model Limitations and Validity Domain	65
4.9	Transition to Model Validation Against Experimental Data	67
4.10	Chapter Summary	68
5	UN SIFGRS Model Validation	70
5.1	Introduction and Validation Strategy	70
5.2	SNAP-50 Experimental Context and Data	71
5.2.1	Historical Overview and Motivation	71
5.2.2	BISON Modeled Test Cases and Irradiation Conditions	71
5.2.2.1	Operating Envelope	72
5.2.2.2	Fuel Fabrication and Physical Properties	72
5.2.2.3	Experimental Measurements and Uncertainties	73
5.3	BISON Model Implementation	73

5.3.1	Fuel Pin Geometry and Mesh Design	73
5.3.2	Thermal Boundary Conditions	74
5.3.3	Thermal Model and Material Properties	75
5.3.4	Power History and Fission Rate Definition	76
5.3.5	Summary: Model Validation Strategy	78
5.4	Baseline Model Performance and Gas Partitioning Analysis	78
5.4.1	Collective Model-Experiment Comparison	78
5.4.2	Strategic Specimen Selection for Parametric Investigation	79
5.4.3	Gas Partitioning Diagnosis: Baseline Model Limitations	80
5.5	Systematic Parametric Investigation	83
5.5.1	Saturation Coverage Reduction	83
5.5.2	Grain Boundary Vacancy Diffusion Enhancement	85
5.5.3	Dislocation Density Reduction	86
5.5.4	Intragranular Diffusion Enhancement	88
5.5.5	Trapping and Resolution Rate Adjustment	90
5.5.6	Combined Parameter Investigation	92
5.5.7	Multi-Objective Validation	95
5.6	Regime Identification and Model Limitation Analysis	97
5.7	Summary and Transition to Sensitivity Analysis	101
6	Global Sensitivity Analysis	103
6.1	Morris Elementary Effects Methodology	103
6.2	Fission Gas Release Sensitivity	104
6.3	Volumetric Swelling Sensitivity	106
6.4	Parameter Selection for Bayesian Calibration	107
6.5	Summary	108
7	UQ and Bayesian Calibration	110
7.1	Introduction and Scientific Motivation	110
7.2	Bayesian Inference and MCMC Theory	112
7.3	Surrogate-Accelerated Bayesian Inference Framework	114
7.3.1	Motivation and Computational Challenge	114
7.4	Surrogate Model Development	115
7.4.1	Design of Experiments and Training Data	115
7.4.2	Gaussian Process Surrogate Specification	116
7.4.3	Surrogate Training and Accuracy Validation	117
7.4.3.1	Empirical Uncertainty Calibration	117
7.5	MCMC Calibration: Methodology, Convergence, and Parameter Uncertainty	118
7.5.1	Surrogate-Based Likelihood and MCMC Architecture	118
7.5.2	No-U-Turn Sampler Selection and High-Dimensional Efficiency	119
7.5.3	Prior Distribution Specification and Experimental Uncertainty	119
7.5.4	Convergence Diagnostics Validating MCMC Stationarity	120

7.6	Prior Refinement Through Mechanistic Guidance and Regime-Stratified Calibration	120
7.6.1	Prior Specification Strategy: From Uninformed to Mechanistically-Guided	120
7.6.1.1	Regime Definitions and Experimental Coverage	121
7.6.2	Impact of Prior Refinement on MCMC Convergence	122
7.6.3	Enhanced Parameter Identification Through Informed Priors	124
7.6.4	Systematic MCMC Calibration Across Experimental Regimes	125
7.6.4.1	Convergence and Information Gain by Regime	126
7.6.4.2	Selection of Informed Priors and Multi-Regime Analysis Strategy	126
7.7	Posterior Parameter Analysis and Regime-Dependent Model Validation	127
7.7.1	Multi-Regime Posterior Comparison: Revealing Data Sufficiency Across Operating Conditions	128
7.7.1.1	Individual Regime Posteriors	128
7.7.1.2	Direct Visual Evidence of Regime Dependence	131
7.7.1.3	Posterior Statistics and Quantitative Regime Comparison	132
7.7.1.4	Data Insufficiency Assessment and Experimental Implications	133
7.7.2	Mechanistic Parameter Interpretation for Primary Calibration Domain	134
7.7.2.1	Grain Radius and Dislocation Density	134
7.7.2.2	Fractional Yield	135
7.7.2.3	Surface Energy and Saturation Coverage	135
7.7.2.4	Scale Factors and Their Physical Roles	136
7.7.2.5	Trapping and Resolution Scale Factors	137
7.7.2.6	Resolution Scale Factors	138
7.7.2.7	Nucleation Rate and Defect-Limited Regimes	138
7.7.2.8	Summary: Parameter Landscape and Mechanistic Pathways	138
7.7.3	Validation and Forward Pathway	140
8	Posterior Prediction and Forward UQ	142
8.1	Introduction	142
8.2	Methods	143
8.2.1	Validation Framework and Success Criteria	143
8.2.2	Posterior Predictive Distributions	144
8.2.3	Uncertainty Decomposition	144
8.2.4	Prior Versus Posterior Comparison	145
8.2.5	Validation Metrics	145
8.3	Results	146
8.3.1	Low-Temperature and Low-Burnup Posterior Predictions	146
8.3.2	BISON Validation Against Direct Simulations	147
8.3.3	Out-of-Sample Validation: Held-Out Experiment 613	151
8.3.4	Prior Versus Posterior Comparison: Quantifying Information Gain	153

8.3.4.1	Uncertainty Reduction	153
8.3.4.2	Accuracy Improvement	154
8.3.4.3	Calibration Journey: Representative Example	154
8.3.4.4	Information Gain Summary	155
8.3.5	Cross-Regime Validation Summary	155
8.4	Conclusions: Validation and Mechanistic Constraints	157
9	Conclusions and Future Directions	160
9.1	Central Hypothesis and Research Objectives	160
9.2	Key Findings and Hypothesis Validation	161
9.2.1	Material Property Foundation	161
9.2.2	Historical Experimental Database Synthesis	161
9.2.3	Mechanistic Model Implementation	162
9.2.4	Validation and Regime Identification	162
9.2.5	Global Sensitivity Analysis	163
9.2.6	Bayesian Parameter Calibration and Information Gain	164
9.2.7	Posterior Predictions and Forward Uncertainty Quantification	165
9.3	Validation of Central Hypothesis	166
9.4	Summary and Future Directions	167
9.4.1	Computational Contributions	167
9.4.2	What Was Achieved	167
9.4.3	What Remains Uncertain	168
9.4.4	Priorities for Future Investigation	168
	References	169
	Bibliography	175
A	INL MCMC Implementation	179
1.1	Overview and Significance	179
1.2	Affine Invariant Ensemble Sampler with Multi-Variance Support	180
1.2.1	Foundational Sampler	180
1.2.2	Multi-Variance Architecture	180
1.3	Multi-Observable Extension and Variable Grouping	181
1.3.1	Multi-Variable Gaussian Likelihood with Grouped Noise	181
1.3.1.1	Sample-Major Data Organization	181
1.3.1.2	Multi-Variable Reporter and Data Aggregation Pipeline	182
1.3.2	Multi-Variance State Management and Prior Specification	182
1.3.3	Multi-Variable Decision Logic with Convergence Masking	183
1.3.3.1	Convergence-Based Automatic Rejection	184
1.3.4	State Reversion and Output Reorganization	184
1.4	Computational Performance and Scaling Analysis	185

1.4.1	Deployment Configuration	185
1.4.2	Computational Cost Analysis and Dimensionality Scaling	185
1.4.3	Fundamental Barriers to Direct 16+2 Dimensional Inference	186
1.5	Model Complexity and Nonlinear Solver Challenges	186
1.5.1	Physical Nature of Computational Barriers	186
1.6	Transition to Surrogate Acceleration and Knowledge Transfer	187
1.7	Summary and Scientific Significance	188

List of Figures

1.1	Accelerated fuel qualification framework developed at Los Alamos National Laboratory.	3
1.2	Crystal structure of Uranium Mononitride.	6
1.3	Crystal structure of Uranium Dioxide.	7
2.1	Temperature-dependent elastic properties of UN from ab initio molecular dynamics simulations.	16
2.2	UN experimental Young's modulus data versus percent theoretical density. . . .	17
2.3	UN experimental Poisson's ratio data versus percent theoretical density.	18
2.4	UN experimental shear modulus data versus percent theoretical density.	20
2.5	UN experimental bulk modulus data versus percent theoretical density.	21
2.6	Typical creep deformation curve showing strain versus time at constant stress. .	23
2.7	UN experimental thermal expansion coefficient data versus temperature.	24
2.8	UN experimental nitrogen self-diffusion data for UN versus temperature.	26
2.9	UN experimental uranium self-diffusion data for UN versus temperature.	27
2.10	Comparison of thermal conductivity for UN and UO_2 versus temperature. . . .	29
2.11	UN experimental thermal conductivity data versus temperature at various porosities.	30
2.12	Experimental thermal conductivity data for UN versus temperature at experimental porosities.	31
2.13	Experimental thermal conductivity data for UN versus temperature corrected to 100% TD using porosity correction factor.	32
2.14	UN experimental heat capacity data versus temperature.	34
3.1	Volumetric swelling versus burnup color-coded by irradiation temperature and power density across all historical UN testing programs.	46
3.2	Fission gas release versus burnup color-coded by irradiation temperature and power density across all historical UN testing programs.	47
3.3	Volumetric swelling versus burnup for all historical UN irradiation programs, color-coded by test program. Program-specific clustering reflects irradiation temperature and power density effects.	48
3.4	Fission gas release versus burnup for all historical UN irradiation programs, color-coded by test program. Different symbol shapes denote distinct testing programs.	49

5.1	Single fuel pin mesh utilizing BISON fuel pin mesh generator.	74
5.2	Three fuel pin configuration showing the variation in fuel temperature due to power influence and cladding boundary condition.	74
5.3	Representative ABAQUS-fitted temperature profile for BISON cladding boundary condition.	76
5.4	Baseline SIFGRS model predictions compared to experimental measurements for all SNAP-50 test cases.	79
5.5	Baseline model gas partitioning for six representative SNAP-50 specimens showing fraction of total produced fission gas in each microstructural location versus burnup.	81
5.6	Effect of grain boundary saturation coverage reduction on fission gas release predictions for three specimen.	84
5.7	Effect of grain boundary vacancy diffusion enhancement on on fission gas release predictions for three specimen.	85
5.8	Effect of dislocation density reduction on fission gas release predictions for three specimen.	87
5.9	Gas partitioning evolution for dislocation density reduction across three temperature regimes.	88
5.10	Effect of intragranular diffusion enhancement on fission gas release.	89
5.11	Gas partitioning for moderate trapping and resolution rate adjustments (trapping = $0.5\times$, resolution = $5\times$) across three temperature regimes.	91
5.12	Effect of trapping and resolution on fission gas release predictions for three specimen.	92
5.13	Gas partitioning comparison for aggressive combined parameter optimization (trapping = 0.1, resolution = 10, diffusivity = 10) across three temperature regimes.	93
5.14	Effect of trapping, resolution and diffusivity on fission gas release predictions for three specimen.	94
5.15	Predicted versus experimental scatter plots for volumetric swelling and fission gas release across all parameter variations.	96
5.16	Regime analysis of SNAP-50 experimental data and baseline model performance.	98
5.17	Fission gas release diagnostics for the standard regime.	99
5.18	Swelling diagnostics for the standard regime.	100
6.1	Fission gas release parameter importance ranked by μ^* . Results colored by temperature, power density, and burnup.	105
6.2	Volumetric swelling parameter importance ranked by μ^* . Results colored by temperature (left), power density, and burnup.	108
7.1	Convergence diagnostics across five calibration regimes comparing uninformed priors and informed priors.	123
7.2	Parameter information gain with variance reduction and precision gain across regimes	125

7.3	Posterior distributions for all 14 mechanistic parameters in low-temperature, low-burnup regime.	128
7.4	Posterior distributions for all 14 mechanistic parameters in high-temperature regime.	129
7.5	Posterior distributions for all 14 mechanistic parameters in high-burnup regime.	130
7.6	Posterior distributions for all 14 mechanistic parameters overlaid to directly compare regime-dependent identifiability.	131
8.1	Posterior predictive distributions for low-temperature and low-burnup regime showing swelling and fission gas release across all calibration experiments.	147
8.2	BISON validation of posterior predictions across low-temperature and low-burnup regime and Test 613 (hold-out validation).	149
8.3	BISON calibration performance showing predicted versus observed for low-temperature and low-burnup regime.	150
8.4	Calibration for experiment 603 middle ($T = 1450$ K).	154

List of Tables

5.1	SNAP-50 Operating Envelope and Measurement Ranges	72
5.2	Specimen selection for parametric investigation including operating conditions. .	80
5.3	Fission gas release predictions across parameter variations.	95
5.4	Volumetric swelling predictions across parameter variations.	96
6.1	Summary of top parameters for fission gas release based on mean and maximum values of μ^* and σ	106
6.2	Summary of top parameters for swelling based on mean and maximum values of μ^* and σ	107
7.1	MCMC calibration regimes with experimental coverage and operating ranges. .	121
7.2	Marginal posterior statistics for all 14 mechanistic parameters in low-temperature and low-burnup.	132
7.3	Posterior means for six mechanistically critical parameters across three calibration regimes.	133
7.4	Marginal posterior statistics for all 14 mechanistic parameters in low-temperature and low-burnup in physical space.	134
8.1	Regime posterior prediction metrics by experiment and position.	148
8.2	Surrogate fidelity metrics comparing Gaussian Process predictions to direct BISON simulations for identical posterior parameter samples.	151
8.3	Out-of-sample validation results for held-out experiment 613 not included in MCMC training.	152
8.4	Cross-regime validation metrics for posterior predictions.	156

List of Acronyms

AFM antiferromagnetic.

AFQ accelerated fuel qualification.

AIE affine invariant ensemble.

AIMD ab-initio molecular dynamics.

ANL Argonne National Laboratory.

BCA binary collision approximation.

BMI Battelle Memorial Institute.

CANEL Connecticut Advanced Nuclear Engineering Lab.

CD cluster dynamics.

CTE coefficient of thermal expansion.

DFT density functional theory.

EBR-II Experimental Breeder Reactor II.

FAST Fission Accelerated Steady-state Testing.

FCMI fuel-cladding mechanical interaction.

FFTF Fast Flux Test Facility.

FGR fission gas release.

FIMA fissions per initial metal atom.

FM ferromagnetic.

INL Idaho National Laboratory.

LANL Los Alamos National Laboratory.

LDA+DMFT local density approximation + dynamical mean field theory.

LHS latin hypercube sampling.

LLS lower-length-scale.

MAE mean absolute error.

MCMC Markov chain Monte Carlo.

MD molecular dynamics.

ML machine learning.

MOOSE Multiphysics Object-Oriented Simulation Environment.

NEAMS Nuclear Energy Advanced Modeling and Simulation.

NUTS No-U-Turn Sampler.

ORNL Oak Ridge National Laboratory.

PIE post-irradiation examination.

QE quantum espresso.

RMSE root mean squared error.

SIFGRS Simple Integrated Fission Gas Release and Swelling.

TD theoretical density.

TEM transmission electron microscopy.

UN uranium mononitride.

UO2 uranium dioxide.

UQ uncertainty quantification.

Chapter 1

Introduction

This dissertation investigates whether mechanistic models which are systematically validated against experimental data across multiple operating regimes and calibrated through rigorous Bayesian uncertainty quantification (UQ), can substantially reduce epistemic (knowledge-based) uncertainty in nuclear fuel performance predictions and establish defensible boundaries for qualified operational domains and suggests necessary experiments. The work applies this framework specifically to UN fuel for micro-reactor systems, demonstrating how accelerated fuel qualification (AFQ) principles can compress qualification timelines while maintaining rigorous scientific and regulatory standards.

1.1 Motivation

Advanced reactor concepts are emerging as critical technologies for addressing global energy demands while meeting stringent safety, non-proliferation, and environmental requirements. Among these concepts, micro-reactors represent a novel class of compact nuclear systems designed for deployment in remote locations, industrial heat applications, and distributed energy generation scenarios. These systems operate under conditions fundamentally different from conventional light-water reactors, often featuring higher power densities, elevated operating temperatures, and long operational periods without refueling. To enable the commercial viability and regulatory acceptance of micro-reactor technology, the nuclear fuel must be carefully selected and rigorously validated for safe and reliable performance across the anticipated operational envelope.

UN has emerged as a candidate fuel for micro-reactor applications due to several advantageous material properties compared to conventional uranium dioxide (UO₂). Most notably, UN possesses a thermal conductivity of approximately 20.8 W/(m·K) at 1000 K, significantly higher than UO₂'s degraded conductivity over temperature. Combined with its high fissionable uranium density (13.5 g/cm³ for UN versus 9.7 g/cm³ for UO₂) and elevated melting temperature (3120 K), UN enables compact fuel pin designs and improved operational margins for micro-reactor systems [1]. However, substantial data gaps remain

regarding UN behavior under irradiation, particularly with respect to fission gas release (FGR) and fuel swelling at the high burnups anticipated for advanced reactor deployments.

The experimental irradiation database for UN, while valuable, is largely historical, dating from space power reactor programs of the 1960s through 1980s. The wide scatter in available data, combined with differences in testing conditions, fuel specifications, and measurement methodologies, creates substantial uncertainty in fuel performance predictions for modern micro-reactor designs. Additionally, practical irradiation testing timelines extend over multiple years due to lengthy irradiation testing, creating a significant bottleneck for fuel qualification. This confluence of factors establishes the foundation for developing an AFQ program tailored specifically to UN that couples mechanistic modeling with targeted experimental validation.

AFQ methodologies employ mechanistic modeling, advanced statistical techniques, and targeted experimental validation to compress the fuel qualification timeline while maintaining regulatory defensibility. By coupling mechanistic understanding of fuel physics with UQ frameworks, AFQ programs can identify critical knowledge gaps and prioritize experimental campaigns to maximize information gain. For UN in particular, this approach enables systematic investigation of how fission gas release and swelling phenomena respond to variations in temperature, power density, burnup, and fuel microstructure, ultimately establishing defensible operational boundaries for micro-reactor deployment.

1.2 Accelerated Fuel Qualification Framework

1.2.1 Historical Context and Necessity

Historically, nuclear fuel qualification has followed a sequential, data-intensive approach that relies heavily on empirical models derived from extensive irradiation testing. This methodology, while rigorous, typically requires 20 or more years to accumulate sufficient data for regulatory approval of a new fuel system [2]. For advanced reactor technologies such as micro-reactors, this traditional timeline represents a significant barrier to commercial viability and deployment. The AFQ methodology addresses this challenge by integrating mechanistic fuel performance modeling, advanced experimental validation, and data science techniques to compress the qualification timeline to as few as five years while maintaining the same high safety standards [2].

1.2.2 Core Philosophy and Principles

The fundamental principle underlying AFQ is that mechanistic, physics-based models can effectively interpolate between sparse experimental data and identify critical knowledge gaps, thereby reducing the overall number of costly integral irradiation tests required for qualification [2]. Rather than accumulating decades of empirical data and then fitting mathematical correlations after, the AFQ approach couples predictive modeling with

targeted, theory-informed experiments performed in parallel. This synergy between modeling and experiment enables efficient resource allocation and accelerates the path to regulatory confidence. The AFQ framework used to guide this work is shown in Figure 1.1.

AFQ integrates four complementary elements working in concert [2]: (1) high-fidelity, physics-based mechanistic models developed using multiscale and multiphysics approaches to provide predictive capability grounded in fundamental physics rather than empirical fitting (2) data science and UQ methods to identify dominant phenomena and prioritize experiments to maximize information gain from each experimental campaign (3) targeted out-of-pile and in-pile experiments to validate models and reduce parameter uncertainty through strategic data collection (4) integration of the validated models into engineering-scale fuel performance codes suitable for reactor licensing and design optimization.




Figure 1.1: Accelerated fuel qualification framework developed at Los Alamos National Laboratory.

1.2.3 Multiscale Modeling and Simulation

Mechanistic fuel performance models must bridge multiple length and time scales to capture the physical processes governing fuel behavior [2]. At the atomic scale, density functional theory (DFT) and machine learning (ML) interatomic potentials enable calculation of defect formation and migration energies, which feed into cluster dynamics and

molecular dynamics (MD) simulations. These lower-length-scale methods predict microscale phenomena such as fission gas bubble evolution and dislocation dynamics. The results are parameterized into engineering-scale fuel performance codes (e.g., BISON) that solve coupled thermo-mechanical and species diffusion equations for entire fuel pins under reactor conditions.

For UN, this multiscale approach is particularly valuable because the limited irradiation database and scatter in historical data make purely empirical correlations domain-limited. By developing mechanistic models designed in first-principles calculations and validated against targeted experiments, the AFQ framework enables confident predictions of fuel behavior under micro-reactor operating conditions not well-represented in historical datasets.

1.2.4 Uncertainty Quantification and Bayesian Inference

Modern data science techniques, including ML and Bayesian inference, enable efficient extraction of information from sparse datasets and systematic propagation of uncertainties through mechanistic models. Reduced-order models developed via data-driven methods can capture the complex dependencies predicted by high-fidelity simulations at computational costs suitable for engineering-scale analysis. Sensitivity analysis within this framework identifies which material properties, model parameters, or operating conditions most strongly influence fuel performance predictions, thereby guiding prioritization of experimental validation efforts.

UQ extends beyond traditional error analysis to quantify epistemic uncertainties arising from gaps in data and mechanistic understanding. By propagating these uncertainties through coupled models to final safety margins (e.g., maximum allowable burnup or temperature), the AFQ methodology provides an objective, defensible basis for regulatory decisions regarding acceptable margins and identifies the specific experiments or simulations most likely to reduce the controlling uncertainties.

1.2.5 Targeted Experimental Validation

Experiments within accelerated qualification programs are designed using model predictions and UQ results to maximize information content per experiment. Separate-effects tests (out-of-pile and in-pile) target specific phenomena identified as critical to fuel performance, such as fission gas diffusion, swelling rates or thermal conductivity degradation. Post-irradiation examination (PIE) coupled with advanced characterization techniques (scanning electron microscopy, transmission electron microscopy, X-ray diffraction, etc.) provides microstructural data to bridge length scales between atomistic models and engineering-scale predictions.

Accelerated irradiation techniques such as Fission Accelerated Steady-state Testing (FAST) and ion irradiation enable rapid accumulation of relevant data while maintaining prototypic fuel temperatures. These methods, combined with mechanistic modeling, allow

the AFQ program to achieve qualification-relevant burnups and doses on timescales of months to years rather than decades.

1.2.6 Application to Uranium Mononitride Fuel

For UN in micro-reactors, the AFQ methodology addresses specific challenges: (1) fission gas release behavior at high burnup has not been extensively characterized at micro-reactor temperatures and power densities; (2) swelling mechanisms, particularly breakaway swelling at high damage levels, remain incompletely understood; and (3) the limited historical irradiation database (primarily from 1960s through 1980s space power programs) shows significant scatter and may not adequately represent modern fuel fabrication practices or micro-reactor operating regimes. By combining mechanistic models of fission gas generation, migration, and release with targeted experiments and rigorous UQ, this dissertation develops the theoretical, computational, and experimental foundation necessary for confident qualification of UN fuel under accelerated timelines suitable for micro-reactor deployment.

1.3 Uranium Mononitride as an Advanced Reactor Fuel

1.3.1 Material Characteristics

UN is a ceramic nuclear fuel comprised of a combination of the following family of ceramic materials: Uranium Mononitride (UN), Uranium Sesquinitride (U_2N_3) and Uranium Dinitride (UN_2). During the fabrication process of UN, all three ceramic materials will be present with the majority being UN. UN crystallizes in a lattice of NaCl type, which is a face-centered cubic (fcc) crystal structure [1]. If the lattice constant and theoretical mass of the crystal structure is known, the theoretical density may be calculated utilizing the following equation:

$$\text{Theoretical Density} \left(\frac{\text{g}}{\text{cm}^3} \right) = \frac{\text{mass}_{\text{cell}}}{\text{volume}_{\text{cell}}} \quad (1.1)$$

At room temperature (≈ 300 K), a UN unit cell has a mass of 1.6741×10^{-21} g with a lattice constant of 4.899 ± 0.001 Å [1], yielding a unit cell volume of 116.93×10^{-24} cm³. Utilizing Equation (1.1), the theoretical density (TD) of UN at room temperature is 14.317 g/cm³.

Comparing UN to other nuclear fuels, such as UO_2 , both fuels have cubic structures however, UN has a rock salt (NaCl) crystal structure, as seen in Figure 1.2, while UO_2 has a fluorite crystal structure, as shown in Figure 1.3. What makes UN stand out above other fuel types, in particular UO_2 , is its high thermal conductivity (≈ 20.8 W/(m·K)) at 1000 K), heavy metal (fissionable) density 13.5 g/cm³ for UN versus 9.7 g/cm³ for UO_2), and




Figure 1.2: Crystal structure of Uranium Mononitride.

high melting temperature (3120 ± 30 K) [1]. Due to these properties, UN is increasingly considered for use in micro-reactors and other high-temperature reactor systems.

1.3.2 Historical Performance Data

While UN has demonstrated promise in historical irradiation programs from the 1960s through 1980s space power reactor era (SNAP-50, SP-100), the experimental database exhibits significant scatter and may not adequately represent modern fuel fabrication practices or micro-reactor operating regimes [1]. Moreover, many correlations developed from this historical data lack quantified uncertainty bounds, and gaps remain in understanding UN behavior at the high burnup and elevated power densities and temperatures required for micro-reactor operations. The SNAP-50 program at Connecticut Advanced Nuclear Engineering Lab (CANEL) and Oak Ridge National Laboratory (ORNL), BR-10 testing in Russian fast reactors, SP-100 irradiation campaigns at Experimental Breeder Reactor II (EBR-II) and Fast Flux Test Facility (FFTF), NASA Advanced Power Reactor testing at Plum Brook Reactor Facility, and Battelle high-temperature research reactor programs collectively accumulated substantial irradiation experience spanning temperatures from 1200 K to 1900 K and burnups up to 8 %FIMA.

These programs demonstrated that UN fuel could sustain significant burnup at elevated temperatures, establishing proof-of-concept for advanced reactor deployment. However, the experimental database exhibits considerable scatter, with volumetric swelling varying by 30 to 50 percent around trend lines and fission gas release varying by factors of 2 to 3 at nominally identical irradiation conditions. This scatter reflects differences in fuel fabrication practices, specimen geometry, reactor thermal conditions, and post-irradiation examination methodologies employed across different research groups and time periods.




Figure 1.3: Crystal structure of Uranium Dioxide.

1.3.3 Data Gaps and Knowledge Limitations

The historical UN database exhibits limitations that prevent confident extrapolation to modern micro-reactor operating regimes. High-temperature data above 1600 K are limited to two SNAP-50 specimens, insufficient for establishing temperature-dependent behavior with confidence. High-burnup data above 3 % are sparse, where burnup is given in units of fissions per initial metal atom (FIMA), limiting understanding of saturation and potentially breakaway swelling mechanisms.

Modern nuclear fuel fabrication now includes stringent controls on oxygen and carbon content that differ from historical specimens, potentially affecting bubble nucleation and gas migration kinetics. Furthermore, the limited dataset at high burnups (greater than 3 %FIMA) and the absence of data representative of the sustained high power densities and elevated temperatures characteristic of micro-reactor operation create critical gaps in the empirical knowledge base.

This confluence of superior material properties combined with data gaps and quantified uncertainties in existing knowledge creates both technical opportunity and practical necessity for advanced mechanistic modeling and rigorous UQ to enable confident UN fuel qualification.

1.4 Fission Gas Release and Swelling as Performance-Limiting Phenomena

Under irradiation in a nuclear reactor, UN experiences complex physical and chemical transformations driven by the intense neutron flux and the heat generated by fission. Two dominant phenomena govern UN fuel performance and establish the operational boundaries for safe reactor operation: FGR and fuel swelling. Both phenomena are interconnected through the generation, migration, and accumulation of fission products within the fuel

matrix, and both have profound implications for fuel-cladding mechanical interaction (FCMI), cladding integrity, and ultimate fuel qualification [1].

1.4.1 Fission Gas Generation and Release

When uranium nuclei undergo fission, they split into fission fragments with significant kinetic energy. These fragments decay into stable isotopes known as fission products, which are classified into three categories: solid, volatile, and gaseous. Among the gaseous products, xenon (Xe) and krypton (Kr) are the dominant contributors to fuel performance phenomena due to their insolubility in the fuel matrix and their tendency to remain gaseous at reactor operating temperatures [1]. Within the fuel, fission gas atoms accumulate through two competing pathways: retention as dispersed nano-scale bubbles within the grain structure, or migration to grain boundaries where bubbles coalesce and grow. The balance between these two mechanisms, controlled by temperature, fuel microstructure, and the rate of fission product generation, determines the total FGR to the fuel plenum.

Historical irradiation data for UN shows that fission gas release remains minimal (typically less than 1 percent) at burnups below approximately 3 % burnup. However, at elevated temperatures and high power densities, FGR increases significantly with burnup, reaching values of 10 to 30 percent or higher at high burnups [1]. This nonlinear behavior reflects a transition from the retention regime (where fission gases remain trapped in bubbles and interconnected porosity) to the release regime (where grain boundary coalescence creates pathways for gas escape). The mechanisms governing this transition, particularly the conditions that initiate and sustain rapid gas release, remain incompletely characterized, representing a critical knowledge gap for high-burnup micro-reactor operation.

1.4.2 Fuel Swelling

Irradiation-induced swelling of UN arises from two primary mechanisms: solid fission product swelling and gaseous swelling. Solid fission products contribute approximately 1 percent volumetric swelling per 1 %FIMA burnup, a relatively modest contribution due to the low thermal gradients in UN that discourage significant migration of solid fission atoms [1]. The dominant swelling mechanism is gaseous swelling, driven by the accumulation and expansion of fission gas bubbles constrained within the fuel. In the low-temperature, low-burnup regime, swelling proceeds at a relatively constant rate dominated by the elastic accommodation of gas bubbles. However, as temperature increases or as fission gas bubbles coalesce into larger voids at grain boundaries, the creep strength of the fuel becomes the limiting factor controlling swelling. When the internal gas pressure and thermal stresses exceed the fuel's creep strength, accelerated swelling occurs, a phenomenon known as breakaway swelling, leading to rapid volumetric expansion that can result in cladding contact and potential mechanical failure.

Historical UN irradiation data demonstrates complex swelling behavior strongly dependent on temperature and power density [1]. Some test campaigns showed progressive linear

swelling up to 8 percent volumetric swelling over 6 to 8 %FIMA burnup (saturation swelling), while others exhibited breakaway swelling at moderate burnups (2 to 4 %FIMA). The physical origins of these competing behaviors and their dependence on fuel microstructure, fabrication history, and operating conditions are not fully understood from the historical database. This variability underscores the necessity for mechanistic modeling validated against targeted experiments to predict swelling behavior reliably under micro-reactor operating conditions.

1.5 Dissertation Outline

1.5.1 Dissertation Objectives

This dissertation investigates UN fuel performance within the AFQ framework. The central hypothesis is that mechanistic, first-principles models, rigorously validated and calibrated using Bayesian UQ, can substantially reduce epistemic uncertainty in fuel performance predictions while establishing defensible operational boundaries.

Three research goals guide the work:

1. **Validate the UN mechanistic model** against historical irradiation data, identifying performance boundaries and regimes where mechanistic understanding is complete or incomplete.
2. **Perform sensitivity analysis** to identify high-impact input parameters for calibration.
3. **Calibrate the model via Bayesian inference**, quantifying and reducing remaining uncertainty for regulatory guidance and prioritization of experimental needs.

The research progresses logically from characterization of UN behavior to model development, validation, and uncertainty quantification.

1.5.2 Dissertation Structure

Chapter 2 establishes the material property foundation for subsequent modeling, compiling thermal and mechanical properties of UN from historical literature. Temperature-dependent correlations are developed and uncertainty bounds quantified to define baseline epistemic uncertainty for downstream models.

Chapter 3 synthesizes historical irradiation data, analyzing fission gas release and swelling observations. This chapter identifies data gaps, characterizes variability, and establishes validation targets for mechanistic modeling.

Chapters 4–5 define and validate a mechanistic model for fission gas release and swelling. The model incorporates fission product generation, bubble kinetics, atomic diffusion across multiple regimes, and grain boundary behavior. Chapter 5 rigorously

compares model predictions to experimental data, defining regime-dependent operational boundaries and extracting key physical insights.

Chapters 6–8 focus on uncertainty quantification. Chapter 6 conducts global sensitivity analysis to identify influential parameters and quantify prediction uncertainty. Chapter 7 applies surrogate-assisted Bayesian inference to calibrate model parameters, reducing epistemic uncertainty and producing posterior distributions for forward prediction. Chapter 8 evaluates predictive performance, performing coverage analysis, out-of-sample validation, and uncertainty decomposition to assess model generalization and reliability.

Chapter 9 synthesizes findings, validating the central hypothesis and discussing implications for accelerated qualification of UN fuel in advanced reactor environments. The dissertation demonstrates that integrated mechanistic modeling and UQ provide a robust framework for fuel performance prediction and informed experimental prioritization.

Chapter 2

Uranium Mononitride Material Properties and Correlations

Understanding the temperature-dependent material properties of nuclear fuel is essential for accurate reactor design, safety analysis, and fuel performance modeling. For UN, comprehensive characterization of mechanical and thermal properties establishes the foundation for mechanistic fuel performance predictions under micro-reactor operating conditions. Accurate material property correlations are critical inputs to multiphysics fuel performance codes such as BISON, which depend on property data to resolve coupled stress-strain and temperature fields, simulate mechanical feedback on thermal behavior, and evaluate long-term fuel performance under design and operational conditions [3]. This chapter compiles experimental data from historical UN characterization campaigns spanning 1960–2020, develops temperature-dependent correlations for critical material properties, and quantifies uncertainties in these correlations to support subsequent modeling efforts in Chapter 5. The properties addressed are limited to unirradiated UN fuel at beginning-of-life conditions; irradiation-induced changes in material behavior, fission gas release, and fuel swelling phenomena are treated separately in Chapter 3.

2.1 Methodology: Material Property Compilation and Correlation Development

2.1.1 Overview and Computational Context

To construct reliable empirical correlations for material properties suitable for mechanistic fuel performance modeling and uncertainty quantification, a consistent and systematic approach was employed. The methodology is designed to be broadly applicable across properties and material systems, ensuring that correlations capture relevant trends in experimental data while maintaining rigorous statistical validation. Critically, this work explicitly quantifies epistemic uncertainties arising from experimental scatter and data gaps, enabling these

uncertainties to propagate probabilistically through the model to determine the uncertainty in key fuel performance metrics.

The material properties developed here serve as direct inputs to engineering-scale fuel performance codes such as BISON [3], which require material model functions of the form $k(T, \phi)$ for thermal conductivity, $c_p(T)$ for specific heat, $\alpha(T)$ for thermal expansion, $E(T, \phi)$ for elastic modulus, $\dot{\epsilon}_{cr}(T, \sigma, \dot{F})$ for creep rate, and $D(T)$ for diffusion coefficients. Section 2 details the methodology used to develop these correlations; subsequent sections (3–4) present the correlations themselves and connect each property to its computational role in mechanistic models.

2.1.2 Experimental Data Collection and Compilation

A comprehensive literature search spanning six decades of UN research was conducted to identify all published experimental measurements of material properties. Experimental data were collected from multiple source categories specifically peer-reviewed journal articles in materials science and nuclear engineering journals and technical reports from national laboratories (Los Alamos National Laboratory (LANL), ORNL, Argonne National Laboratory (ANL), Battelle Memorial Institute (BMI)). The intent was to maximize data accessibility while maintaining scientific credibility through established publication venues and quality assurance practices at national laboratories.

Data extraction methodology was rigorous to preserve information quality. Tabulated property values were transcribed directly with associated measurement conditions. Where only graphical data were available, digitization was performed using tools to extract numerical values from published figures. For each data point, the following metadata were recorded: temperature (K), porosity (percent theoretical density), measurement methodology, fuel fabrication technique (sintering, arc melting, etc.), and reported experimental uncertainty where available. Cross-referencing with source publications ensured consistency and identified any apparent transcription errors.

The experimental database encompasses measurements from the following data sources:

- **Elastic Properties:** Dynamic resonant frequency methods were employed by Honda and Kikuchi [4], Padel and de Novion [5], and Whaley et al [6]. Ultrasonic pulse measurement techniques were used by Hall [7], Muta et al. [8], Speidel and Keller [9], Taylor and McMurtry [10], Guinan and Cline [11], and Adachi et al [12]. This methodological diversity introduces systematic differences that are explicitly accounted for in uncertainty quantification.
- **Thermal Expansion:** Dilatometer measurements were conducted by Speidel and Keller [9], Taylor and McMurtry (under vacuum to minimize oxidation) [10], DeCrescente et al. [13], and Carlsen and Harms [14]. The vacuum testing by Taylor and McMurtry contrasts with air-atmosphere tests by others; oxidation effects at elevated temperatures are documented and considered in uncertainty assessment.

- **Thermal Conductivity:** Steady-state heat flow measurements were performed by Kollie and Moore [15], Moore et al. [16], Speidel and Keller [9], Takahashi et al. [17], Hayes and DeCrescente [18], Endebrook et al. [19], and Muta et al [8]. These investigations span a broad temperature range (298–1923 K) and multiple measurement geometries (cylindrical, planar samples).
- **Heat Capacity:** Differential scanning calorimetry and drop-calorimetry techniques were employed by Westrum and Barber [20], Counsell et al. [21], Affortit [22], Harrington [23], Takahashi et al. [17], Oetting and Leitnaker [24], Cordfunke and Muis [25], and Muta et al. [8].
- **Diffusion Coefficients:** Nitrogen self-diffusion data come from Holt and Almassy [26], DeCrescente et al. [13], and Matzke [27]. Uranium self-diffusion data are from Reimann et al. [28] and Matzke [27].

2.1.3 Correlation Development Approach

For each material property, experimental data was systematically fit to candidate functional forms: linear, exponential, power-law, and polynomial models. The best-fit expression was selected based on coefficient of determination (R^2) combined with physical plausibility of the fitted form and extrapolability characteristics. For properties spanning multiple decades in magnitude (e.g., diffusion coefficients), power-law fits typically provided the best statistical performance. Uncertainties in developed correlations arise primarily from statistical uncertainty. Scatter in experimental data around fitted correlations was characterized by calculating the standard deviation of residuals (difference between measured values and correlation predictions):

$$\sigma_{\text{residual}} = \sqrt{\frac{\sum_{i=1}^n (P_{\text{measured},i} - \hat{P}_{\text{correlation},i})^2}{n - k}} \quad (2.1)$$

where n is the number of data points and k is the number of fitted parameters. This standard deviation quantifies aleatoric uncertainty arising from inherent measurement variability and material heterogeneity not explicitly parameterized by temperature and porosity.

For each correlation, confidence intervals are reported using the framework:

$$\text{CI}_{95\%} = \hat{P}(T, p) \pm 2 \cdot \sigma_{\text{residual}} \quad (2.2)$$

where $\hat{P}(T, p)$ is the correlation prediction as a function of porosity and/or temperature, and σ_{residual} is the standard deviation of residuals from Equation 2.1. Standard deviations were calculated using appropriate degrees of freedom for the chosen functional form (typically $k = 2$ for linear and power-law fits, $k = 3$ for quadratic forms).

Model validity was verified through residual analysis. Residuals were plotted versus predicted values and versus independent variables (temperature, porosity) to check for non-

random patterns, heteroscedasticity (non-constant residual variance), or other signs of systematic model inadequacy. Correlations exhibiting residual patterns suggesting model inadequacy were refined (e.g., by switching functional forms) until residuals displayed appropriate random characteristics. This rigorous statistical validation ensures that confidence intervals accurately capture correlation uncertainty rather than reflecting inadequate model specification.

2.2 Mechanical Properties

2.2.1 Overview and Computational Role in Fuel Performance

Mechanical properties are fundamental to understanding the structural behavior of UN under both steady-state and transient reactor conditions. Parameters such as Young’s modulus, Poisson’s ratio, creep resistance, and thermal expansion govern the mechanical response of the fuel to thermal gradients, fission-induced swelling, and externally applied constraints. These properties not only define the material’s stiffness and deformation behavior but also control the development of internal stresses that lead to fracture, fuel pellet relocation, or fuel-cladding mechanical interaction [3].

From a computational perspective, the mechanical properties developed in this section are direct inputs to the stress-strain solver in mechanistic fuel performance codes (Chapter 4), where they parameterize the constitutive relationships governing elastic deformation, plastic creep, and mechanical accommodation of thermal stress and internal fission gas pressure. Elastic properties control the rate of elastic stress buildup in response to thermal gradients; creep properties determine how readily plastic deformation relaxes stresses; thermal expansion coefficients govern the magnitude of thermal stress generation itself.

2.2.2 Temperature-Dependent Elastic Properties

Elastic properties are inherently dependent on porosity, which reduces the effective stiffness of the bulk material by introducing compliance-reducing discontinuities (pores). However, temperature also plays a significant role in the elastic response through temperature-induced changes in atomic bonding, lattice constant evolution, and lattice dynamics. While porosity often dominates experimental scatter in elastic property measurements, measurable reductions in elastic moduli with increasing temperature reflect anharmonic effects in the phonon spectrum and thermal softening of interatomic potentials. Direct experimental characterization of temperature dependence remains challenging due to difficulties in high-temperature mechanical testing and maintaining phase stability in UN samples at elevated temperatures.

In the available literature, only Padel and de Novion [5] have performed temperature measurements of UN Young’s modulus, providing the only direct experimental data for temperature dependence. Historically, this temperature dependence was utilized across

all elastic properties. To overcome limitations imposed by this sparse high-temperature experimental data, temperature-dependent scaling was incorporated using normalized fits from ab-initio molecular dynamics (AIMD) simulations by Kocevski et al. [29]. By combining the accuracy of room-temperature experimental measurements with physics-based temperature scaling from AIMD, correlations extend beyond the experimental database while remaining grounded in verified data.

The AIMD data span 300 K to 1800 K under both antiferromagnetic (AFM) and ferromagnetic (FM) magnetic configurations; normalized scaling factors were applied following:

$$P_{\text{corr}}(T) = P_{\text{exp}}(T_{\text{ref}}) \cdot \frac{P_{\text{AIMD}}(T)}{P_{\text{AIMD}}(T_{\text{ref}})} \quad (2.3)$$

where $P_{\text{exp}}(T_{\text{ref}})$ is the experimental value at reference temperature (300 K), and $P_{\text{AIMD}}(T)$ and $P_{\text{AIMD}}(T_{\text{ref}})$ are the AIMD-computed properties at the target and reference temperatures, respectively. This scaling preserves the experimental baseline accuracy while leveraging AIMD's strength in capturing temperature-dependent physical trends. Resultant temperature-dependent elastic property correlations are presented as:

$$E(T) = 1.0728 - 2.4269 \times 10^{-4}T \quad (2.4)$$

$$K(T) = 1.0536 - 1.7869 \times 10^{-4}T \quad (2.5)$$

$$G(T) = 1.0702 - 2.3401 \times 10^{-4}T \quad (2.6)$$

$$\nu(T) = 0.9949 + 5.3524 \times 10^{-6}T + 3.9272 \times 10^{-8}T^2 \quad (2.7)$$

where E is Young's modulus, K is bulk modulus, G is shear modulus, ν is Poisson's ratio, and T is temperature in Kelvin. These correlations apply over the range 300–1800 K. The underlying AIMD results are shown in Figure 2.1, which displays normalized elastic properties versus temperature under both AFM and FM configurations; mechanistic fuel performance modeling (Chapter 4) primarily utilizes AFM predictions as the ground-state configuration relevant for reactor conditions at low magnetic fields.

2.2.3 Young's Modulus

Young's modulus (or modulus of elasticity) characterizes a material's stiffness and its ability to resist uniaxial tensile or compressive stress without permanent deformation. Mathematically, Young's modulus is defined as the ratio of stress to strain within the elastic region of the stress-strain curve (Hooke's Law):

$$E = \frac{\sigma}{\epsilon} \quad (2.8)$$

where σ is the applied stress and ϵ is the induced strain. A material with high modulus exhibits less elastic strain under a specified applied stress compared to a material with lower




Figure 2.1: Temperature-dependent elastic properties of UN from ab initio molecular dynamics simulations [29] normalized to 300 K experimental reference values. Both antiferromagnetic (AFM) and ferromagnetic (FM) magnetic configurations are shown.

modulus. As temperature increases, atomic vibration within the crystal structure intensifies, effectively increasing interatomic distances and weakening atomic bonding, which reduces Young's modulus. This temperature softening is captured in Equation 2.4.

Experimental data for UN Young's modulus span nine independent investigations employing either dynamic resonant frequency methods [4–6] or ultrasonic pulse measurement techniques [7–12]. All room-temperature measurements were performed at approximately 300 K, with porosity ranging from theoretical density (i.e. 0% porosity) to approximately 30% porosity. Using a power-law fit combined with normalized AIMD temperature scaling (Equation 2.3), the Young's modulus correlation becomes:

$$E(T, D) = 2.479 \times 10^{-4} (D)^{3.013} [1.0728 - 2.4269 \times 10^{-4} T] \quad [\text{GPa}] \quad (2.9)$$

where D is the percent theoretical density and T is temperature in Kelvin. This correlation applies over 300–1800 K and 70–100% TD. The uncertainty envelope encompassing the correlation is ± 28.1 GPa ($\pm 2\sigma$), with all experimental data within this band. For high-density UN samples ($>95\%$ TD), predictions fall within 10% of measured values, validating

applicability to micro-reactor fuel specifications where density typically exceeds 95% TD.




Figure 2.2: Experimental Young's modulus data for UN versus percent theoretical density at room temperature (≈ 300 K) for data from nine independent research groups [4–12]. Uncertainty band represents $\pm 2\sigma$ residual scatter (± 28.1 GPa).

2.2.4 Poisson's Ratio

Poisson's ratio describes the lateral (transverse) strain that develops when a material experiences axial strain under uniaxial loading. It is defined as the negative ratio of transverse to axial strains:

$$\nu = -\frac{\epsilon_{\text{transverse}}}{\epsilon_{\text{axial}}} \quad (2.10)$$

The negative sign ensures Poisson's ratio is always positive. For isotropic materials, Poisson's ratio typically ranges between 0 and 0.5. Poisson's ratio is coupled to Young's modulus and shear modulus through elastic theory, allowing derivation of one quantity if the others are known experimentally.

Limited experimental data exist for UN's Poisson's ratio. Only two studies [5, 12] directly measured or inferred this quantity at room temperature. Augmenting sparse experimental data with AIMD temperature scaling (Equation 2.3), the Poisson's ratio correlation is:

$$\nu(T, D) = 7.575 \times 10^{-4} (D)^{1.286} [0.9949 + 5.3524 \times 10^{-6} T + 3.9272 \times 10^{-8} T^2] \quad (2.11)$$

The polynomial temperature dependence (visible in Equation 2.7) from AIMD captures a subtle increase in Poisson's ratio with temperature, reflecting changes in acoustic wave propagation velocities and competition between shear and normal stiffness evolution with temperature. Uncertainty bounds are ± 0.029 ($\pm 2\sigma$), with experimental data scattered within 21% of the fitted trend and average differences of 5%. For high-density UN, predictions show agreement within 10% of measured values.




Figure 2.3: Experimental Poisson's ratio data for UN versus percent theoretical density at room temperature (300 K) for data from two independent research groups [5, 12]. Uncertainty bands represents $\pm 2\sigma$ residual scatter (± 0.029).

2.2.5 Shear Modulus

Shear modulus (or rigidity modulus) quantifies a material's resistance to shear deformation. Unlike Young's modulus, which describes response to normal stress, shear modulus governs material stiffness under shear loading. Shear modulus is defined as:

$$G = \frac{\tau}{\gamma} \quad (2.12)$$

where τ is the applied shear stress and γ is the shear strain. The three independent elastic constants for isotropic materials, Young's modulus, shear modulus, and Poisson's ratio), are interrelated through:

$$G = \frac{E}{2(1 + \nu)} \quad (2.13)$$

For UN, shear modulus exhibits similar porosity and temperature trends to Young's modulus, though with distinct sensitivity coefficients reflecting differences in how shear and normal deformations couple to the material microstructure. Integrating experimental data from six ultrasonic measurement campaigns [5, 6, 9–12] with normalized AIMD temperature scaling, the shear modulus correlation is:

$$G(T, D) = 2.173 \times 10^{-5}(D)^{3.353} [1.0702 - 2.3401 \times 10^{-4}T] \quad [\text{GPa}] \quad (2.14)$$

The higher porosity exponent (3.353 versus 3.013 for Young's modulus) indicates that shear modulus is more sensitive to porosity variations, likely reflecting the pronounced effect of pores as stress concentrators under shear loading. Uncertainty bounds are ± 7.96 GPa ($\pm 2\sigma$), with 95% of experimental data within this envelope and average differences of 5%. For high-density fuel, shear modulus predictions fall within 10% of measured values except for Taylor and McMurtry data, which shows agreement within 13%.

2.2.6 Bulk Modulus

Bulk modulus (or incompressibility) represents the material's resistance to volumetric compression and is directly related to Young's modulus and Poisson's ratio through:

$$K = \frac{E}{3(1 - 2\nu)} \quad (2.15)$$

Physically, bulk modulus characterizes how hydrostatic pressure (or uniform stress) changes material density. For UN under reactor conditions where fuel pellets experience modest internal pressure from fission gas accumulation, bulk modulus influences the coupling between internal pressure and density changes, controlling the swelling response to fission gas bubble pressurization. The bulk modulus is defined formally as:

$$K = -P \frac{\partial V}{\partial P} \quad (2.16)$$

where P is applied pressure and V is volume.

Experimental data for bulk modulus are limited; only four independent studies provided sufficient data for correlation development. Combining these measurements [5, 6, 11, 12] with normalized AIMD predictions, the bulk modulus correlation is:

$$K(T, D) = 4.425 \times 10^{-8}(D)^{4.841} [1.0536 - 1.7869 \times 10^{-4}T] \quad [\text{GPa}] \quad (2.17)$$




Figure 2.4: Experimental shear modulus data for UN versus percent theoretical density at room temperature (300 K) for data from six independent research groups [5, 6, 9–12]. Uncertainty band represents $\pm 2\sigma$ residual scatter (± 7.96 GPa).

Bulk modulus exhibits the largest relative uncertainty among elastic moduli: ± 19.0 GPa ($\pm 2\sigma$). This elevated uncertainty stems from both experimental scatter and sensitivity to subtle differences in sample preparation and pore geometry. High-density UN samples show bulk modulus values within 10% of predictions (Taylor and McMurtry data at 15%), supporting model applicability for near-theoretical-density micro-reactor fuel. Bulk modulus enters mechanistic models where it controls the relationship between internal gas bubble pressure and density change, directly affecting the coupling between fission gas accumulation and fuel swelling.

2.2.7 Thermal Creep

Creep is the time-dependent permanent deformation of a material subjected to constant stress at elevated temperature. As illustrated in the classical creep curve, creep deformation progresses through three distinct stages: (1) primary creep exhibits high plastic deformation rate that decreases over time as work hardening accumulates; (2) secondary creep shows a relatively constant strain rate (steady-state), where work hardening is balanced by recovery processes; (3) tertiary creep exhibits a rapid strain rate increase, typically leading to material




Figure 2.5: Experimental bulk modulus data for UN versus percent theoretical density at room temperature (300 K) for data from four independent research groups [5, 6, 11, 12]. Uncertainty band represents $\pm 2\sigma$ residual scatter (± 19.0 GPa).

failure. For ceramics such as UN, creep becomes mechanistically important at temperatures above approximately 50% of the melting temperature. Since the melting point of UN is approximately 3120 K, significant creep rates are anticipated above approximately 1560 K.

Nuclear fuels subjected to radial temperature gradients experience substantial thermal stress in the fuel. Thermal creep can cause permanent dimensional changes; when creep occurs and the applied load is removed, material relaxation does not follow the original elastic pathway, resulting in permanent deformation. A practical consequence is a permanent increase in fuel pellet radius, potentially reducing clearance with the cladding and eventually leading to fuel-cladding mechanical interaction, a concern particularly important in compact micro-reactor core designs where initial pellet-cladding gap is small.

High-temperature, steady-state creep rate is described by the power-law relation:

$$\dot{\epsilon} = Ad^{-m}\sigma^n \exp\left(-\frac{Q}{RT}\right) \quad (2.18)$$

where $\dot{\epsilon}$ is the creep rate, d is grain size, m is the grain size exponent, σ is applied stress, n is the stress exponent, Q is activation energy, R is the gas constant, and T is temperature. The empirical constants A , m , n , and Q must be determined experimentally.

Hayes et al. [30] performed a comprehensive creep analysis using UN data from Fassler et al. [31], Vandervoort et al. [32], and Massaaki and Michio [33] across temperatures 1770–2083 K and stresses 20–34 MPa. A stress exponent of 4.5 was determined, indicating that UN follows a dislocation climb mechanism at high temperatures (stress exponents 2.5–5.0 characterize dislocation climb [34]). Since UN creep is dominated by dislocation climb, it is independent of grain size ($m = 0$). Based on porosity sensitivity observations between hypo-stoichiometric porous and stoichiometric dense samples, porosity significantly impacts creep rate; however, due to limited experimental data spanning the full stress-temperature space, the primary correlation assumes theoretically dense UN:

$$\dot{\epsilon} = 2.054 \times 10^{-34.5} \sigma^{4.5} \exp\left(-\frac{39369.5}{T}\right) \quad [1/\text{s}] \quad (2.19)$$

where σ is applied stress in MPa and T is temperature in Kelvin. This relation is specific to the experimental testing conditions specified above (1770–2083 K, 20–34 MPa). The correlation yields reasonable values down to approximately 1350 K. For porous UN, a porosity correction factor applies:

$$f_p = 0.9871 - \frac{P}{27.6 - 8.65P} \quad (2.20)$$

where P is the fractional porosity. The corrected creep rate is obtained by multiplying Equation 2.19 by Equation 2.20.

The temperature dependence of creep (exponential in Equation 2.19) creates strong coupling. Modest increases in centerline temperature due to higher power or reduced thermal conductivity can double or triple creep rates, dramatically affecting swelling accommodation mechanisms.

2.2.8 Coefficient of Thermal Expansion

The coefficient of thermal expansion (CTE) describes the fractional change in length per unit temperature change:

$$\alpha = \frac{1}{l_0} \frac{\partial l}{\partial T} \quad (2.21)$$

where l_0 is the original length and T is temperature. Thermal expansion plays a critical role in fuel performance. As temperature varies radially within a fuel pellet, differential expansion induces complex internal stress fields. Radial temperature gradients cause the center to expand more than the surface, generating tensile stresses on the pellet surface that can lead to radial cracking and mechanical fracture. Understanding fuel pellet expansion is essential for predicting pellet-cladding mechanical interaction, associated cladding stresses, and thermal fatigue from operational cycling.

Experimental data for UN thermal expansion span four independent investigations [9, 10, 13, 14] using dilatometer methods over room temperature to approximately 2500 K. Taylor




Figure 2.6: Typical creep deformation curve showing strain versus time at constant stress. Three distinct stages are visible: primary creep (high strain rate, decreasing over time as work hardening accumulates), secondary creep (relatively constant strain rate, steady-state where hardening balances recovery), and tertiary creep (accelerating strain rate to failure).

and McMurtry [10] conducted experiments under vacuum to minimize oxidation while others performed tests in air with noted oxidation at elevated temperatures. Data were collected at various porosities: DeCrescente et al. [13] at 94% TD, Carlsen and Harms [14] at 91.5% TD, Taylor and McMurtry [10] at 89.0% TD, and Speidel and Keller [9] at 97.6% TD.

Using a power-law fit that best captures the experimental data distribution, the thermal expansion coefficient correlation is:

$$\alpha(T) = 2.089 \times 10^{-6} T^{0.209} \quad [1/K] \quad (2.22)$$

valid over 298–2523 K and 0–20% porosity. This power-law form reduces the average spread between experimental data and theoretical predictions by approximately a factor of three compared to alternative fitting forms. The uncertainty envelope is $\pm 6.453 \times 10^{-7} 1/K$ ($\pm 2\sigma$), with all experimental data within 13% of the trendline and average differences of 3%. Hayes et al. [35] developed a linear correlation using lattice parameter changes that is also applicable, however, the power-law form presented here minimizes residual variance across

the full experimental database and is therefore preferred for fuel performance modeling. The subtle temperature exponent (0.209) reflects the gradual saturation of thermal expansion at elevated temperatures as lattice anharmonicity effects plateau.




Figure 2.7: Experimental thermal expansion coefficient data for UN versus temperature for data from four independent research groups [9, 10, 13, 14]. Uncertainty band represents $\pm 2\sigma$ residual scatter ($\pm 6.453 \times 10^{-7}$ 1/K).

In mechanistic models, thermal expansion couples to the thermal stress generation via the thermal strain induced by temperature gradients:

$$\varepsilon_{\text{thermal}} = \int_{\text{ref}}^T \alpha(T') dT' \quad (2.23)$$

The variation of α with temperature affects the radial stress profile within the fuel and is coupled to elastic properties (Young's modulus) to predict stress magnitudes. Higher thermal expansion coefficients generate larger thermal strains and thus larger stresses for a given temperature gradient.

2.2.9 Uranium and Nitrogen Self-Diffusivity

Diffusion in UN is less extensively studied compared to more established fuels such as UO_2 , yet experimental data exist from multiple sources combined with recent computational

approaches. Diffusion properties are critical for understanding fission gas behavior, fuel swelling mechanisms, and long-term performance under irradiation. Self-diffusion coefficients parameterize the Fick’s law constitutive equation governing mass transport of uranium and nitrogen atoms through the lattice, which in turn affects the rate of fission gas atom migration, fission product precipitation, and bubble nucleation kinetics.

2.2.9.1 Nitrogen Self-Diffusivity

Experimental nitrogen diffusion data come from several sources employing distinct measurement techniques: Holt and Almassy [26] used alpha particle activation analysis; DeCrescente et al. [13] performed direct measurements (noting that unit conversions from Celsius to Kelvin reconcile apparent high values with other sources); Matzke [27] reported values across a wide temperature range. Multiple computational approaches augment experimental data: developed cluster dynamics models [36], applied machine learning optimization [37], and developed neural network surrogates [38].

The nitrogen self-diffusion correlation developed for this work, valid over 1400–2400 K, is:

$$D_N(T) = 1.255 \times 10^{-65} T^{15.294} \quad [\text{m}^2/\text{s}] \quad (2.24)$$

Uncertainty bounds are $\pm 1.199 \times 10^{-14} \text{ m}^2/\text{s}$ ($\pm 2\sigma$), with experimental data within 144% of the trendline and average differences of 32%, indicating substantial scatter. This large uncertainty reflects both the experimental challenges inherent in measuring diffusion at high temperatures and the limited temperature range (1400–2400 K) of available data.

2.2.9.2 Uranium Self-Diffusivity

Uranium diffusion data are more limited. Reimann et al. [28] used alpha energy degradation methods; results qualitatively agree with defect chemistry calculations and show dependence on nitrogen partial pressure. Matzke [27] generated U diffusion coefficients at two nitrogen partial pressures.

The uranium self-diffusion correlation developed for this work, valid over 1900–2200 K (combining data from both nitrogen pressures), is:

$$D_U(T) = 1.604 \times 10^{-106} T^{26.264} \quad [\text{m}^2/\text{s}] \quad (2.25)$$

Uncertainty bounds are $\pm 2.267 \times 10^{-19} \text{ m}^2/\text{s}$ ($\pm 2\sigma$), with experimental data within 175% of the trendline and average differences of 65%, indicating very significant scatter. This exceptionally large uncertainty reflects both the experimental challenges in measuring uranium diffusion and the very limited temperature range (1900–2200 K) of available data.




Figure 2.8: Experimental nitrogen self-diffusion data for UN versus temperature (1400–2400 K) from multiple sources [13, 26, 27]. Uncertainty band represents $\pm 2\sigma$ residual scatter.

2.3 Thermal Properties

2.3.1 Overview and Computational Role

Material selection for nuclear reactors is strongly influenced by thermal response to temperature changes, thermal gradients, and elevated temperatures. Thermal conductivity of nuclear fuel directly impacts the maximum centerline temperature that develops during operation. Higher thermal conductivity facilitates more efficient heat transfer through the fuel, reducing the temperature difference between center and surface. Understanding thermal properties enables accurate prediction of fuel temperature distributions, which drive mechanical stress fields, fission gas release kinetics, and long-term dimensional stability. The thermal properties presented in this section, thermal conductivity and specific heat capacity, are direct inputs to the transient thermal equation that the fuel performance codes solve iteratively coupled with mechanical deformation and fission gas transport.

From a computational standpoint, thermal conductivity enters the heat conduction equation:

$$\rho C_p \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) + \dot{q} \quad (2.26)$$

where ρ is density, C_p is specific heat capacity, T is temperature, k is thermal conductivity,




Figure 2.9: Experimental uranium self-diffusion data for UN versus temperature (1900–2200 K) (both nitrogen pressure regimes combined) [27, 28]. Uncertainty band represents $\pm 2\sigma$ residual scatter.

and $\dot{q}$ is volumetric heat generation rate (from fission). Thermal conductivity multiplies the temperature gradient term, directly controlling how efficiently heat diffuses radially from the high heat generation center toward the cooler surface. Higher thermal conductivity reduces the temperature difference required to conduct a given heat flux, thereby reducing centerline temperatures and thermal gradients. Thus, thermal conductivity uncertainty propagates to all downstream fuel performance predictions.

2.3.2 Thermal Conductivity

2.3.2.1 Thermal Transport Physics in UN: Electronic vs. Phononic Contributions

Thermal conduction is the process by which heat transfers across a material due to a temperature gradient, with energy flowing from high to low temperature regions. This process is quantified by thermal conductivity, defined through Fourier's law:

$$\frac{dq}{dt} = -k \frac{dT}{dx} \quad (2.27)$$

where $\frac{dq}{dt}$ is the rate of heat flow per unit time, k is thermal conductivity (W/(m·K)), and $\frac{dT}{dx}$ is the temperature gradient across the material. Heat transfer in solids occurs through two primary physical mechanisms: lattice vibrations (phonons) and free electron movement. When both mechanisms are present, the total thermal conductivity is the sum of phonon and electronic contributions:

$$k_{\text{total}} = k_{\text{phonon}} + k_{\text{electronic}} \quad (2.28)$$

For most ceramic materials, thermal energy transport is dominated by lattice vibrations. Phonons carry thermal energy from regions of high temperature to low temperature through crystal lattice interactions. However, phonon mobility is intrinsically lower than free electron mobility because phonons require scattering events with the lattice rather than direct transport [39]. Consequently, ceramic thermal conductivities are typically lower than those of metals and exhibit decreasing thermal conductivity with increasing temperature due to enhanced phonon-phonon scattering. For example, UO_2 follows this behavior with thermal conductivity decreasing monotonically with temperature.

UN exhibits fundamentally different thermal transport physics. Computational studies using multiple methods, such as local density approximation + dynamical mean field theory (LDA+DMFT) [40], MD [41], quantum espresso (QE) [42, 43], and AIMD [29], consistently demonstrate that UN thermal conductivity is dominated by free electron movement rather than phonon transport. Yin et al. [40] reported that at 1000 K, the phonon contribution to thermal conductivity is only 4.4 W/(m·K), while the electronic contribution is 12.1 W/(m·K), yielding a total of 16.5 W/(m·K). This electronic dominance arises from UN's rock salt (NaCl) crystal structure and metallic-like electronic band structure, which supports substantial free carrier density. As a consequence, UN thermal conductivity increases with temperature (Figure 2.10), opposite to the behavior of UO_2 and most ceramics. This temperature-increasing trend reflects enhanced electronic mobility at elevated temperatures, partially offset by increased phonon scattering. This distinctive physics makes UN more suitable to UO_2 for high-temperature applications where maintaining high thermal conductivity is advantageous.

Phonon transport in UN, though subdominant at unirradiated baseline, is heavily influenced by defects that form under irradiation, including fission products on the lattice, point defects, and dislocations. These defects act as phonon scattering centers, reducing phonon conductivity. Fission-induced microstructure evolution therefore affects thermal conductivity through phonon scattering degradation, even though electronic contributions dominate the baseline.

2.3.2.2 Data Compilation and Porosity Effects

Porosity significantly impacts UN thermal conductivity through two mechanisms: (1) air-filled pores having low thermal conductivity, directly reducing effective thermal conductivity through volume averaging, and (2) pore surfaces act as phonon scattering sites, further reducing phonon transport [39]. Experimental data systematically demonstrate that




Figure 2.10: Comparison of thermal conductivity for UN and UO₂ versus temperature.

thermal conductivity decreases on average by 12.5% for every 5% reduction in density [44]. Experimental thermal conductivity data for UN span seven independent investigations [8, 9, 15–19] conducted between room temperature and approximately 1923 K across porosities from theoretical density to 30% porosity.

Experimental data exhibit scatter beyond porosity effects. Cross-referencing fabrication provenance data reveals that oxygen concentration in UN samples systematically lowers thermal conductivity. Takahashi et al.’s 89.9% TD sample with the highest oxygen concentration (0.085 wt.%) exhibits the lowest thermal conductivity among comparable-density samples [17]. Similarly, comparing the two Speidel and Keller samples, the specimen with higher oxygen concentration (0.047 wt.%) yields lower thermal conductivity than the oxygen-free sample (0.0 wt.%) [9]. Oxygen forms oxide phases (primarily UO₂ and U₂N₃) at grain boundaries, introducing phonon scattering sites and disrupting electronic conduction paths. Additional microstructural variables influencing thermal conductivity include grain size, grain boundary density, dislocation density, and impurity content.

Figure 2.12 shows experimental data at measured porosities, while Figure 2.13 shows data corrected to 100% TD using the porosity correction factor:

$$k_p = k_{100} - 2.14P \quad (2.29)$$

where k_p is the thermal conductivity of the porous sample, k_{100} is the thermal conductivity

at 100% TD, and P is the volume fraction porosity.




Figure 2.11: Experimental thermal conductivity data for UN versus temperature at various porosities [44].

2.3.2.3 Thermal Conductivity Correlation

Using a power-law fit to the compiled experimental database combined with the porosity correction, the thermal conductivity correlation developed for this work is:

$$k(T, P) = (2.123 - 2.14P) T^{0.338} \quad [\text{W}/(\text{m}\cdot\text{K})] \quad (2.30)$$

where T is temperature in Kelvin and P is volume fraction porosity. This correlation applies over 298–1923 K and 0–20% porosity. The temperature exponent (0.338) reflects the balance between increasing electronic mobility (which increases k) and phonon scattering saturation at elevated temperatures (which decreases k). Uncertainty bounds are $\pm 5.03 \text{ W}/(\text{m}\cdot\text{K})$ ($\pm 2\sigma$), with experimental data within 39% of the trendline and average differences of 11%. For high-density UN, thermal conductivity predictions agree within 15% of measured values across the full temperature range, validating applicability to micro-reactor fuel specifications where density targets exceed 95% TD.

The thermal conductivity correlation enters fuel performance models as a critical input controlling the fuel temperature distribution. Uncertainties in thermal conductivity directly

propagate to centerline temperature predictions, which in turn affect creep rate, fission gas diffusion kinetics, and swelling rates through exponential temperature dependencies.




Figure 2.12: Experimental thermal conductivity data for UN versus temperature at experimental porosities from seven independent research groups[8, 9, 15–19]. AIMD [29], LDA+DMFT [40], and MD [41] computational predictions overlay experimental data.

2.3.3 Specific Heat Capacity

2.3.3.1 Heat Capacity Theory and Temperature Dependence

Heat capacity quantifies the amount of energy required to raise the temperature of a material by 1 K, defined as:

$$C = \frac{dQ}{dT} \quad (2.31)$$

where C is heat capacity (J/(mol·K)) and dQ is the energy required to produce a temperature change dT . Specific heat capacity extends this definition to a per-unit-mass basis (J/(g·K)). Heat capacity can be measured under two thermodynamic constraints: constant volume (C_V) or constant pressure (C_P). These are related by the Nernst-Einstein equation:

$$C_P = C_V + \kappa T \quad (2.32)$$




Figure 2.13: Experimental thermal conductivity data for UN versus temperature corrected to 100% TD using porosity correction factor (Equation 2.29). Uncertainty band represents $\pm 2\sigma$ residual scatter (± 5.03 W/(m·K)).

where κ is an empirically determined constant and T is temperature.

Einstein developed a theoretical framework for heat capacity based on phonon contributions, treating lattice vibrations as independent quantum harmonic oscillators. The Einstein model for constant-volume heat capacity is:

$$C_V = 3R \left(\frac{\Theta_E}{T} \right)^2 \frac{\exp(\Theta_E/T)}{[\exp(\Theta_E/T) - 1]^2} \quad (2.33)$$

where R is the universal gas constant, Θ_E is the Einstein temperature (K), and T is material temperature (K). The Einstein temperature represents the characteristic vibrational frequency of the lattice; for UN, Θ_E was empirically determined to be 365.7 K through fitting to experimental data.

2.3.3.2 Experimental Data and Computational Validation

Experimental heat capacity data for UN span eight independent investigations conducted between 298 K and 2900 K [8, 17, 20–25]. Measured heat capacities were obtained at constant

pressure. Hayes et al. [45] developed a best-fit correlation for UN heat capacity by empirically determining the constants in Equations 2.32 and 2.33.

2.3.3.3 Heat Capacity Correlation

The resulting correlation is:

$$C_p(T) = 51.14 + \frac{(\Theta_E)^2}{[\exp(\Theta_E/T) - 1]^2} \exp(\Theta_E/T) + 9.941 \times 10^{-3}T \quad [\text{J}/(\text{mol}\cdot\text{K})] \quad (2.34)$$

where $\Theta_E = 365.7$ K and T is temperature in Kelvin. This correlation applies over 298–2900 K. The first term (51.14 J/(mol·K)) represents the high-temperature (classical) limit (Dulong-Petit law); the second term captures phonon contributions via the Einstein model; the third term accounts for electronic heat capacity contributions (linear in temperature), reflecting the metallic-like electronic structure of UN. Uncertainty bounds are ± 5.34 J/(mol·K) ($\pm 2\sigma$), with experimental data within 14% of the trendline and average differences of 3%.

Computational validation was performed using AIMD (Kocevski et al. [29]) and MD (Kurosaki et al. [41]). AIMD calculations show good agreement with experimental data across the full temperature range, capturing both phonon and electronic contributions. Molecular dynamics results systematically underestimate heat capacity because MD considers only atomic motion (phonon contributions) without electronic contributions. Kurosaki et al. [41] noted that the difference between MD predictions and experimental data reflects precisely the electronic heat capacity term ($9.941 \times 10^{-3}T$ in Equation 2.34), providing computational validation of the electronic contribution to UN’s thermophysical properties and confirming the correctness of the correlation structure.

2.4 Summary and Integration with Irradiation Physics

This chapter has established the comprehensive material property foundation required for predicting UN fuel performance under micro-reactor operating conditions. The correlations developed here, spanning elastic moduli, thermal expansion, creep, thermal conductivity, heat capacity, and diffusion coefficients, represent a systematic compilation and analysis of experimental data spanning six decades (1960–2020) of UN characterization campaigns.

Under irradiation, these baseline material properties become starting points for understanding how fission products and displacement defects modify UN’s thermomechanical response. Fission gas atoms (xenon and krypton) generate within the fuel at rates controlled by the fission rate, and their subsequent migration through the lattice and accumulation in bubbles depends critically on uranium and nitrogen self-diffusion coefficients. The competing mechanisms of gas bubble nucleation and growth within the grain structure versus




Figure 2.14: Experimental heat capacity data for UN versus temperature from eight independent research groups (298–2900 K) [8, 17, 20–25]. Uncertainty band represents $\pm 2\sigma$ residual scatter (± 5.34 J/(mol·K)).

coalescence at grain boundaries, processes central to fission gas release and fuel swelling, depend on material properties including surface energy, bulk modulus, and shear modulus. Irradiation-induced hardening and embrittlement, while not quantified in the unirradiated baseline, arise from accumulation of point defects and dislocations that interact with existing microstructure, effectively modifying the creep strength that controls swelling rates.

Chapter 3 examines the historical UN irradiation database, synthesizing experimental observations of fission gas release and fuel swelling across multiple test programs and operating conditions. The material properties established in the present chapter serve as essential inputs to the mechanistic models developed in later chapters, which predict how fission gas generation, diffusion, and release rates respond to the coupled thermal-mechanical fields within an irradiated fuel pin. By combining the baseline material properties here with an understanding of irradiation-induced microstructural evolution, mechanistic modeling of fission gas release and swelling, and rigorous uncertainty quantification, this dissertation enables confident fuel performance predictions across the operational envelope anticipated for advanced micro-reactor systems.

Chapter 3

Fuel Performance Under Neutron Irradiation

3.1 Overview and Computational Role

Understanding uranium mononitride fuel behavior under neutron irradiation is essential for fuel performance prediction in micro-reactor environments. Neutron irradiation introduces two fundamental changes to fuel that directly affect all thermomechanical properties developed in the preceding chapter. Fission gas generation, where neutron-induced fission creates fission product atoms (primarily xenon and krypton) that must be accommodated within the fuel matrix, and irradiation-induced microstructural evolution, where displacement cascades create vacancies, interstitials, dislocations, and defect clusters that accumulate with dose. These irradiation-induced defects degrade mechanical properties, enhance fission gas atom diffusion, and promote fission gas bubble nucleation and growth at grain boundaries and intragranular sites.

This chapter synthesizes historical UN irradiation testing data spanning the 1960s to the early 2000s, including the SNAP-50 programs at CANEL and ORNL [13, 46], BR-10 testing in the Russian fast reactor [47, 48], NASA Advanced Power Reactor testing at Plum Brook Reactor Facility [49, 50], Battelle high-temperature testing at the Battelle Research Reactor [51, 52], and SP-100 irradiation campaigns at EBR-II and FFTF [53]. Together, these programs provide quantitative relationships between operating conditions (irradiation temperature, power density, burnup) and fuel performance outputs (volumetric swelling, fission gas release). Rather than presenting scattered phenomenology, this chapter establishes the physical basis for fission gas release and swelling as the primary performance-limiting mechanisms, documents the critical temperature and burnup regimes where swelling behavior transitions from gradual accumulation to rapid saturation, and identifies microstructure-dependent effects that will inform mechanistic model development. The unified database of historical measurements directly constrains model parameters and provides experimental benchmarks for validation of uncertainty quantification analysis.

3.2 Fundamental Mechanisms Governing Irradiation-Induced Fuel Performance

Before examining individual testing programs, the underlying physical mechanisms governing fission gas release and fuel swelling must be established, as these phenomena represent the primary limiting factors for fuel qualification and operational lifetime.

3.2.1 Fission Gas Generation and Transport

Each fission event in UN produces two fission fragments plus gaseous fission products with noble gas isotopes (principally ^{133}Xe and ^{85}Kr) comprising approximately 25–30% of the fission yield by atom count. Within microseconds after fission, thermalization of recoil energy dissipates the fission fragment kinetic energy, implanting fission gas atoms at random lattice positions within the fuel matrix. These implanted gas atoms follow several migration pathways where they may remain dissolved or precipitate into small bubbles within the grain interior at low temperatures, they may migrate toward grain boundaries through thermally-activated interstitial diffusion driven by concentration gradients, or they may coalesce at grain boundaries or dislocations into larger bubbles once local concentrations exceed the solubility limit, eventually releasing gas to the fuel surface. The kinetics of gas atom migration, bubble nucleation, and bubble coalescence depend critically on temperature (which governs diffusion rates via an Arrhenius dependence) and on the local defect concentration (which provides both sinks for gas atoms through recombination and enhanced diffusion pathways via radiation-accelerated diffusion).

The effective diffusivity of fission gas atoms in irradiated UN is substantially higher than uranium and nitrogen self-diffusion coefficients alone would suggest, due to enhancement from irradiation-produced point defects and defect complexes. At low temperatures (below approximately 1400 K), fission gas is largely retained within the fuel unless mechanisms for accelerated gas release such as grain boundary recoil knockout, stress-driven micro-cracking, or density-driven gas escape through interconnected porosity become active. At elevated temperatures and high fission rates (high power density), fission gas atoms migrate rapidly to grain boundaries where they coalesce into bubbles. If sufficient gas accumulates at grain boundaries and the resulting bubble pressures exceed the mechanical strength of the grain boundary, interconnected porosity networks form and gas begins to escape to the fuel surface and plenum space.

3.2.2 Fuel Swelling Mechanisms and Saturation Phenomena

Fuel swelling under irradiation arises from several coupled mechanisms. Fission gas atoms precipitated into bubbles occupy volume within the fuel matrix. The volume occupied by fission gas depends on bubble pressure, which is governed by the balance between surface tension, internal gas pressure, and matrix stress. Solid fission products from the fission

process remain at their point of formation and accumulate throughout irradiation; for UN, particulate swelling from solid fission products typically contributes approximately 0.5% volumetric increase per 1% %FIMA burnup [54]. Fission gas preferentially migrates to grain boundaries where lower surface energy reduces the thermodynamic driving force for bubble nucleation. Once grain-boundary bubbles form and grow through coalescence, they occupy progressively larger fractions of the grain boundary area, eventually linking to form interconnected porosity pathways that provide channels for fission gas escape to the fuel surface.

Beyond the volumetric contributions from fission gas and particulate products, neutron irradiation accelerates plastic deformation through irradiation-enhanced creep mechanisms involving mobile point defect fluxes driven by irradiation-induced defect production rates [55]. This irradiation-enhanced creep exhibits fundamentally different temperature and stress dependencies compared to thermal creep where the creep rate is linearly proportional to applied stress rather than following power-law stress dependence, and at temperatures below approximately $0.3 T_m$ (approximately 1000 K for UN), the mechanism is athermal, driven by stress-directed migration of self-interstitial clusters rather than thermally-activated vacancy diffusion. Irradiation-enhanced creep accommodates volume changes from fission gas bubble swelling and solid fission product accumulation without proportional increases in macroscopic dimensional change, provided stress levels remain moderate.

A critical phenomenon observed in multiple high-temperature testing programs is saturation swelling [51]. Early in irradiation, fission gases rapidly migrate to grain boundaries via interstitial diffusion, where they coalesce into bubbles. This rapid grain-boundary bubble formation causes fast volumetric increase during the initial saturation phase, typically resulting in 20–40% volumetric swelling at burnups of approximately 0.5–1.0 %FIMA. Once the fuel reaches approximately 20% volumetric swelling and grain-boundary bubble coalescence reaches interconnectedness (typically 75% or more of porosity becoming interconnected), fission gas production and release reach an equilibrium where newly generated fission gas is approximately balanced by gas release through interconnected pathways. Therefore, volumetric increase stabilizes and becomes approximately constant at higher burnups (plateau phase). In contrast, lower-temperature operation or fuel with smaller grain size exhibits more gradual (approximately linear) swelling accumulation, where fission gas migration to grain boundaries is slower and only a fraction of generated gas escapes, allowing continued volume accumulation with burnup.

3.3 Historical Irradiation Testing Programs and Key Findings

3.3.1 BR-10: Fast Reactor Testing

The Bystry Reactor-10 (BR-10) is a sodium-cooled fast test reactor that operated in Russia from 1959 to 2002 at the Institute for Physics and Power Engineering (IPPE) in

Obninsk. UN fuel testing occurred in the fourth and fifth reactor loadings between 1983 and 2002 [47, 48], making it the most contemporary large-scale irradiation program in the historical database. The program irradiated 1250 driver subassembly (SA) fuel pins at maximum fuel temperatures of 1175 K (driver pins) to 1910 K (experimental assemblies K-1, K-3), with fuel densities ranging from 85–94% theoretical density and maximum burnups reaching 8–8.8 %FIMA.

Three experimental fuel subassemblies (UN-2, K-1, K-3) were selected for detailed post-irradiation examination. A key distinguishing feature was the systematic investigation of fabrication process effects on fuel performance. UN-2 was intentionally divided into two variants, one segment with helium bonding and another with sodium bonding, enabling direct comparison of thermal conductance effects at controlled irradiation conditions. The bonding material directly affects the fuel-cladding gap conductance and thus the fuel centerline temperature; sodium provides approximately 5–10 times higher thermal conductance than helium, producing lower fuel centerline temperatures and modulating swelling behavior independent of neutron flux. K-1 and K-3 were both helium-bonded UNfuel pins but fabricated via distinctly different uranium metallurgical processes. K-1 employed the hydrogenation-nitriding process, while K-3 employed carbothermic reduction of uranium dioxide. The carbothermic process produced uranium oxide impurities that required post-synthesis purification; this design directly tested whether fabrication history and resulting oxygen impurity levels influence swelling and gas release behavior, controlling for temperature and neutron flux [48].

For volumetric swelling versus burnup, driver pin central pellets irradiated at 1473 K exhibited a swelling rate of 1.65% per 1 %FIMA burnup, while upper and lower pellets at lower temperature (1100 K) showed 1.35% per 1 %FIMA, demonstrating temperature acceleration of approximately $1.22\times$ across a 373 K temperature difference. The gas-bonded UN-2 at 1530 K showed a swelling rate approximately $2.0\times$ higher than the sodium-bonded UN-2 at 1100 K, confirming that thermal conductance directly modulates swelling accumulation rates.

Most remarkably, K-1 and K-3, irradiated at nominally identical temperature (1670 K) but fabricated by different metallurgical processes, yielded volumetric swelling rates of 2.0% and 2.6% per 1 %FIMA respectively [48]. This 30% difference in swelling rate at constant temperature is substantial and conclusively demonstrated that fabrication process and its attendant microstructural and impurity effects significantly influence swelling behavior independent of thermal conditions. Oxygen content analysis showed K-3 contained 0.94 wt% oxygen versus K-1 at 0.2 wt%, both exceeding specification. This higher oxygen content in K-3 correlates with the elevated swelling rate, suggesting that oxygen impurities either degrade thermal conductivity (increasing local fuel temperatures), modify creep behavior, or alter grain-boundary and bubble-nucleation characteristics.

A detailed temperature-dependent correlation for swelling rate was developed from BR-10 data across the entire temperature range 993 K to 1953 K

$$\text{Swelling Rate (per 1 \% \%FIMA)} = 6.95 - 1.1 \times 10^{-2}T + 5.25 \times 10^{-6}T^2 \quad (3.1)$$

where temperature T is in Kelvin [47]. This correlation reveals non-monotonic temperature dependence with a maximum swelling rate near 1000 K, contrary to simple Arrhenius expectations. At intermediate temperatures (1000–1300 K), fission gas migration is sufficiently thermally-activated to enable rapid grain-boundary coalescence, creating maximum volumetric growth rate; meanwhile, thermal creep strength remains high enough to provide mechanical resistance to further expansion. At higher temperatures (>1300 K), thermal creep strength decreases, allowing stress relaxation and elastic recovery that partially accommodates volumetric expansion, thus reducing net macroscopic swelling despite continued fission gas generation.

Fission gas release measurements revealed negligible release below 3 %FIMA burnup in all specimens, consistent with low-temperature gas migration limitations. Driver subassembly gas release ranged from approximately 1–3% across 3–8 %FIMA burnup. Experimental assemblies K-1 and K-3, despite identical irradiation temperature and neutron flux, showed markedly different fission gas release, with K-3 yielding 2–3 times higher release percentages than K-1 at equivalent burnups [48], again confirming fabrication-process sensitivity independent of temperature.

BR-10 experienced 26 fuel failures total (2 in Loading IV, 24 in Loading V), primarily at high burnup (6.3–8.0 %FIMA) and attributed to fuel-cladding mechanical interaction (FCMI). All tested fuel rods contained impurities exceeding standard limits, with oxygen 0.2–0.94 wt% (versus specification 0.15 wt%). The elevated impurity levels are noted as potential contributors to FCMI-driven failures through degradation of thermal conductivity and creep strength. Cladding analysis showed circumferential cracks and void formation at fuel-cladding interfaces, demonstrating that fuel qualification must account not only for swelling and gas release, but also for their mechanical interaction with cladding constraints [48].

3.3.2 SNAP-50: Space Reactor Program

The SNAP-50 program represents one of the first comprehensive UNqualification efforts for high-temperature space reactor applications. In the 1960s, a tri-agency agreement among the Atomic Energy Commission (DOE), NASA, and Air Force established objectives to develop a 300–1000 kW electric space power supply using a lithium-cooled fast reactor. Initial designs targeted 8 MW thermal, 10,000 operational hours, maintaining volumetric fuel swelling below 4% and fission gas release below 10%. Initial testing and subsequent post-irradiation testing revealed that these stringent conditions could not be satisfied at the proposed operating parameters. High power density conditions (2–3 kW/cm³) and high temperatures (1600–1900 K) produced swelling and gas release rates incompatible with design limits, necessitating a major redesign cycle. The final SNAP-50 reference design reduced thermal power to 2 MW, maximum fuel centerline temperature to 1533 K, maximum burnup to 1.5 %FIMA, and relaxed limits to permit volumetric swelling and fission gas release of 8% and 20% respectively.

Two separate irradiation testing campaigns were conducted: CANEL by Pratt and Whitney Aircraft [13], and ORNL (Oak Ridge National Laboratory) at the Materials Testing Reactor (MTR) [46]. CANEL testing (1960–1965) investigated high power density conditions ($1.0\text{--}2.0\text{ kW/cm}^3$) at nominal temperatures of 1375 K, 1575 K, and 1775 K. ORNL testing (1966–1975) focused on low to moderate power density ($0.3\text{--}1.4\text{ kW/cm}^3$) and temperatures of 1375–1575 K, with some capsules achieving higher burnups up to 4.6 %FIMA.

At CANEL, initial high power density testing (1.5 kW/cm^3 target) revealed severe swelling limitations that forced immediate redesign. Post-irradiation examination revealed volumetric swelling of approximately 8–10% at low burnup, confirming that high power density conditions produce accelerated swelling rates substantially above design projections. Analysis of the CANEL test series revealed approximately linear swelling accumulation with burnup, with no break-away swelling observed at low burnup. Comparing CANEL high power density ($1.5\text{--}2.0\text{ kW/cm}^3$) and low power density ($0.4\text{--}0.8\text{ kW/cm}^3$) data at nominally equivalent temperatures revealed that swelling rate increases with power density by factors of 20–30% [13]. DeCrescente et al. attributed this to enhanced temperature gradients, faster thermal migration of fission gases in regions of steep temperature gradient, and reduced creep strength at the higher average temperatures in high-power-density pins. Data scatter within each test series reached $\pm 30\%$ at constant nominal burnup and temperature, attributed to variations in neutron flux at different capsule positions, thermal transients, and fabrication variability.

ORNL testing validated the reduced design specifications under conditions closer to the final SNAP-50 reference design. Successful testing of 15 UNfuel capsules was completed; eight achieved the desired 10,000-hour irradiation time, and one extended to 12,000 hours. Burnups achieved ranged from approximately 1.5 to 4.6 %FIMA, with the highest-burnup capsules operated at reduced power to extend irradiation time. Post-irradiation examination included gamma scanning to determine burnup uniformity, macroscopic dimensional measurements, metallographic cross-sections and optical microscopy of grain structure, and in some cases SEM of fractured fuel surfaces. Diametral changes remained below the 2% design limit in all but three pins, despite achieving approximately 4 %FIMA burnup, nearly three times the initial 1.5 %FIMA design specification. This robust dimensional stability at extended burnup was a major program success and strengthened confidence in UNfuel for space applications.

Volumetric swelling remained below the 8% design limit in all but one capsule (which reached 8.7%), occurring in the specimen irradiated at the highest temperature and highest burnup. Fission gas release analysis showed all fuel pins released less than the 20% design limit, with maximum release approximately 12% in the highest-burnup, highest-temperature capsule. Capsule 665, at lower temperature but similar burnup, released approximately 2%, demonstrating strong temperature dependence of fission gas release. Mechanical fracturing of fuel during capsule puncture released substantial additional gas (up to 40% of total recovered), demonstrating that interconnected porosity facilitates gas escape beyond thermally-activated lattice diffusion alone [46].

3.3.3 NASA Advanced Power Reactor Program

The NASA Advanced Power Reactor (APR) program inherited the SNAP-50 design objectives in the 1970s–1980s [56]. Testing was conducted at NASA’s Plum Brook Reactor Facility near Cleveland, Ohio, representing a shift toward more conservative operational design. UNfuel pins were irradiated at substantially lower power densities (0.4–0.8 kW/cm³) and moderate temperatures (1185–1300 K cladding, approximately 1219–1350 K bulk fuel). This lower-temperature regime was motivated by the need for increased operational margins, reduced swelling to accommodate longer fuel lifetimes, and improved reliability for demanding space missions.

Five fuel pins were irradiated for approximately 3797 hours (158 days) [49], achieving burnups of 1.76–1.86 %FIMA. Volumetric swelling ranged from approximately 1.7–2.8%, corresponding to swelling rates of 0.95–1.5% per 1 %FIMA burnup. These rates are substantially lower than contemporary CANEL and ORNL results at higher temperatures, confirming exponential temperature dependence. Maximum diametral changes were 0.5–1.0%, well below 1–2% design limits and demonstrating excellent dimensional stability at moderate temperatures.

The most notable finding was that the low-temperature, low-power-density conditions yielded highly predictable, linear swelling accumulation with minimal data scatter. Standard deviation of swelling rates across the five pins was approximately 0.15% per %FIMA, compared to $\pm 30\%$ scatter in CANEL high-power-density testing. This dramatic reduction in scatter reflects the thermal and mechanical stability at lower power density. Reduced thermal gradients mean fewer localized hot spots, reduced stress means less mechanical feedback, and reduced temperature means slower gas migration and more gradual, equilibrium-controlled swelling. Fission gas release remained negligible at these moderate temperatures and burnups, with no detectable release reported in most pins. The thermal evolution of the fuel during 158 days of irradiation showed minimal microstructural degradation, with grain sizes remaining similar to pre-irradiation values and no secondary phases or dissociation products observed.

3.3.4 Battelle Research Reactor High-Temperature Program

The Battelle Memorial Institute conducted a systematic investigation of UNbehavior under sustained high-temperature irradiation (1950–2100 K bulk fuel) [51, 52], targeting development of dimensionally stable fuels for advanced reactor applications. Testing was performed at Battelle’s research reactor (BRR), equipped with a specialized high-temperature irradiation capsule capable of sustaining fuel temperatures up to 2250 K. The test series, designated HT-BRR, involved seven distinct irradiation experiments (HT-BRR-3, -4, -6, -7, -8, -10, -13) spanning 1968–1972, each investigating specific phenomena related to temperature, density, fabrication history, cladding mechanical properties, and microstructure effects.

3.3.4.1 Saturation Swelling Discovery

In HT-BRR-3, a single 97% TD UNfuel slug was irradiated at nominal cladding surface temperature 1950 K (with estimated fuel centerline temperature approximately 2100 K). Unlike the linear swelling observed in all prior lower-temperature programs, this specimen exhibited rapid initial volumetric swelling reaching approximately 20–22% at burnup 0.8 %FIMA (swelling rate approximately 25% per %FIMA), after which volumetric increase essentially ceased despite continued fission gas production and irradiation to higher burnups (eventually reaching 1.2 %FIMA). Final total volumetric swelling plateaued at approximately 22–24%. Diametral swelling plateaued at approximately 10%, corresponding to roughly 30% volumetric.

Post-irradiation metallography revealed that approximately 75% of the porosity at saturation was interconnected at the grain-boundary network, directly supporting the mechanistic interpretation that once grain-boundary bubble coalescence creates percolating porosity networks, newly generated fission gas preferentially escapes through these interconnected pathways to the fuel surface rather than accumulating further in new bubbles. This discovery of saturation swelling represented a fundamental departure from low-temperature linear accumulation behavior and established the existence of distinctly different swelling regimes depending on irradiation temperature and fuel microstructure.

3.3.4.2 Parametric Investigation of Saturation Mechanisms

Recognizing the significance of HT-BRR-3, subsequent tests investigated variables controlling saturation behavior. HT-BRR-4 tested the effect of a 0.1397 cm central hole in the fuel to enable mechanical gas release. Specimen without hole exhibited saturation swelling similar to HT-BRR-3, while the specimen with central hole showed enhanced saturation-phase swelling but experienced severe cladding bulging and circumferential cracking at high burnup (>3 %FIMA), with post-examination revealing fuel dissociation and uranium metal migration, demonstrating that mechanical gas release mechanisms introduce failure pathways in practice.

HT-BRR-6 investigated cladding mechanical properties by varying cladding thickness and composition while maintaining identical fuel and nominal irradiation temperature. Stronger and thicker cladding reduced volumetric swelling compared to nominal-thickness cladding, with up to 4% of diametral swelling being redirected inward due to cladding mechanical constraint. However, the absolute swelling reduction was modest (approximately 2–3% volumetric swelling decrease), and the mechanical properties benefit came at significant economic and neutron-economy penalty.

HT-BRR-7 studied porosity effects using two 70% TD UNspecimens (versus the standard 97% TD). Remarkably, these low-density specimens exhibited minimal volumetric increase and no saturation swelling phenomenon up to 0.3 %FIMA burnup, suggesting that pre-existing interconnected porosity eliminates the pressure-driven bubble formation and coalescence that drives saturation at high-density fuel. This observation indicated that

saturation swelling is fundamentally linked to grain-boundary bubble coalescence creating new porosity pathways (percolation transition), not merely to gas release through existing voids.

HT-BRR-8 tested fabrication-size effects by irradiating specimens of 85% and 92% TD fabricated with varying uranium powder mesh sizes at nominal 1950 K. Higher-density fuel exhibited higher volumetric swelling than lower-density fuel, contrary to the hypothesis that higher surface area in lower-density fuel would enable more rapid gas release. Post-hoc analysis suggests that density controls the initial porosity distribution, which modulates the percolation transition burnup.

HT-BRR-10 systematically investigated grain size effects by irradiating five UN specimens with intentionally varied grain sizes (ranging 5–350 μm) at 1900–1950 K, with real-time gas-pressure monitoring on one specimen to provide in-situ measurements of gas release kinetics [52]. This experiment was particularly significant because grain size directly controls grain-boundary surface area. All specimens except the large-grained specimen experienced cladding failure during irradiation. The large-grained specimen (approximately 350 μm) retained nearly all fission gases, showing negligible release until 0.8 %FIMA, then gradual increase to approximately 5–10% at 1.2 %FIMA, and exhibited little fission gas bubble growth at grain boundaries. Final swelling in this specimen was approximately 8–10%, substantially lower than the 20–25% saturation swelling observed in fine-grained fuel at comparable temperature.

The fine-grained specimens exhibited rapid volumetric swelling and cladding failure. Direct comparison showed that specimens irradiated approximately 100 K higher than others had substantially larger volumetric swelling (difference of 10–12% volumetric swelling at 0.6 %FIMA), confirming temperature-accelerated swelling. The vast majority of the swelling variation, however, was attributable to grain size. Fine-grained fuel swelled to near-saturation (20%+) while coarse-grained fuel remained at 8–10%, a 2–3 $\times$ difference at identical temperature and burnup. This finding has critical implications for mechanistic modeling and fuel design. Grain size appears to be a primary control variable for swelling behavior, not a secondary effect.

HT-BRR-13 investigated porosity effects on gas release using 93% TD (nominal cladding) and 73% TD (gas-monitored) specimens at 1975 K. Gas-pressure monitoring on the high-porosity 73% TD specimen revealed negligible fission gas release until 0.1 %FIMA, after which release accelerated rapidly. Cumulative gas release was estimated at 65–70% by 0.5 %FIMA. Comparison to HT-BRR-10 and HT-BRR-7 showed complex competing porosity effects. Increased porosity decreases the delay before gas release begins, increases gas release rate post-initiation, and increases total cumulative release. This highlights the importance of pore morphology (size, spatial distribution, interconnectedness) beyond mere density as a control variable.

3.3.4.3 Battelle Synthesis

The Battelle program yielded several critical conclusions. Saturation swelling occurs only in high-density fuel at elevated temperatures where rapid gas migration to grain boundaries creates interconnected porosity through a percolation transition. Grain size strongly influences both swelling rate and cladding failure probability, with coarse-grained fuel showing superior dimensional stability ($2\text{--}3\times$ lower swelling) and gas retention compared to fine-grained fuel at identical temperature and burnup. Cladding mechanical properties provide modest swelling reduction at significant economic and neutronic penalty, and are not a primary design lever. Fabrication process and resulting microstructure (grain size, pore distribution, impurity content) are as important as irradiation conditions in determining swelling behavior. Pre-existing porosity structure fundamentally alters swelling and gas release behavior, enabling careful fuel engineering to suppress saturation through controlled porosity design. These findings established grain-boundary engineering, microstructural control, and fabrication-process optimization as critical design levers for fuel performance, complementary to and possibly exceeding temperature and burnup operational constraints.

3.3.5 SP-100 Program

The SP-100 program represented the final major UNirradiation campaign in the historical database, conducted across multiple fast reactor facilities (EBR-II, FFTF) in the 1980s–1990s [53] and targeting demonstration of reliable UNperformance for space power applications. A total of 86 experimental fuel pins were irradiated with nominal cladding peak temperatures of either 1200 K or 1500 K, corresponding to fuel centerline temperatures of approximately 1700–1950 K depending on power density, thermal gap conductance, and axial position within the test reactor. Capsule design was refined based on decades of SNAP-50, NASA, and Battelle experience. However, detailed post-irradiation examination reports have not been comprehensively released in the open literature, limiting quantitative analysis of swelling and gas release data compared to earlier programs.

SP-1 consisted of four UNfuel pins at 87% TD clad with PWC-11 cladding and featuring free-standing chemical vapor-deposited (CVD) tungsten liners. Irradiation for 96 equivalent full-power days achieved burnups up to 0.83 %FIMA, with two pins continuing to reach 3.9 %FIMA. Measured volumetric swelling reached approximately 8–10% at 0.83 %FIMA (swelling rate approximately 10–12% per %FIMA), substantially higher than contemporary lower-temperature programs but somewhat comparable to high-temperature Battelle HT-BRR results. This extreme swelling was attributed to the high irradiation temperature (estimated 1900–1950 K fuel centerline), consistent with the exponential temperature dependence established by prior programs.

SP-2 tested four UNfuel pins at 87% TD with two variants of cladding materials, two pins with CVD-tungsten-lined PWC-11 and two with Mo–13Re cladding representing an alternative high-temperature material. Irradiation for 181 full-power days achieved 1.3

%FIMA burnup. Limited post-irradiation data were reported. Mo-13Re cladding showed no apparent mechanical failure or unusual corrosion, suggesting dimensional stability.

SP-3 and subsequent variants progressed to higher fuel density (96% TD), representing the most challenging fuel density among historical UN programs. Eight fuel pins were irradiated for 112 full-power days to 0.8 %FIMA burnup, with six pins plus two from SP-1 continuing irradiation to 3.3 %FIMA cumulative burnup.

FSP-1 at FFTF represented the largest SP-100 campaign, comprising 38 fuel pins of varied geometry and density. Irradiation for 269 full-power days achieved burnups up to 2.3 %FIMA at higher enrichment and 0.5 %FIMA at lower enrichment, with accumulated burnups reaching 5.6 %FIMA (original pins) and 3.1 %FIMA (fresh pins).

Detailed examination of two representative fuel pins (NBU-2, NBU-3 from SP-1) by Matthews et al. reported irradiation under approximately 1 kW/cc power density to 0.74 %FIMA burnup with peak cladding temperature 1500 K (estimated fuel centerline approximately 1950 K) [53]. Measured fission gas release was 5.35%, and maximum diametral swelling was 3.5%, yielding volumetric swelling approximately 8.6% at 0.74 %FIMA (approximately 11.6% per %FIMA), consistent with other high-temperature programs at 1900+ K. Higher-density fuel in later series showed reduced dissociation phenomena compared to lower-density fuel at equivalent temperature, suggesting improved structural integrity at extended burnup.

3.4 Dominant Trends Across the Historical Database

Examination of the complete historical database across BR-10, SNAP-50, NASA APR, Battelle, and SP-100 programs reveals several dominant trends that directly inform mechanistic model development and uncertainty quantification strategy.

3.4.1 Temperature and Power Density Dependence

A significant finding from all historical UN irradiation programs is the exponential (or approximately power-law) temperature dependence of both swelling and fission gas release. Volumetric swelling data, when plotted against burnup and color-coded by irradiation temperature and power density (Figure 3.1), show that at constant burnup, swelling rates increase by factors of 5–10 across the 1000–2100 K range [57]. Power density is a clear secondary driver, as higher power density at fixed temperature accelerates swelling by 20–30%.

At intermediate temperatures (1300–1700 K), typical of the SNAP-50 test conditions, swelling growth is near-linear with burnup, and the rate of increase tracks the Arrhenius-like behavior expected from thermally-activated fission gas atom diffusion and bubble nucleation. At higher temperatures (>1900 K), tested in Battelle and some fast reactor experiments, the swelling curve plateaus, indicating that rapid grain-boundary bubble coalescence and




Figure 3.1: Volumetric swelling versus burnup color-coded by irradiation temperature and power density across all historical UN testing programs.

interconnected porosity formation have enabled quasi-equilibrium between gas production and release [51].

Fission gas release data across the same temperature range (Figure 3.2) show even greater temperature sensitivity. Release is negligible ($<1\%$) below 1400 K but rises exponentially above this threshold, with over 50% observed at the highest experiment temperatures. This transition marks the regime where grain-boundary diffusion and interconnected pathways become rapid enough to outpace gas retention and drive bulk release, with far-reaching implications for operational limits and mechanistic modeling [58].

3.4.2 Program-Specific Trends and Microstructural Effects

Program-segregated data (Figures 3.3 and 3.4) reveal distinct clustering for each testing series, reflecting their primary differences. SNAP-50 and NASA cluster below 10% swelling at moderate burnup, consistent with lower temperatures and power densities; Battelle and SP-100 reach much higher swelling at lower burnup due to operation at significantly higher temperatures. Such clustering emphasizes that cross-program extrapolation must be performed cautiously and with mechanistic justification.

Within individual programs, scatter provides insight into controllable variables. The high scatter in CANEL data ($\pm 30\%$) reflects substantial variations in neutron flux at different




Figure 3.2: Fission gas release versus burnup color-coded by irradiation temperature and power density across all historical UN testing programs.

capsule positions, thermal transients during reactor operation, and fabrication variability. The dramatic reduction in scatter in NASA APR data ($\pm 0.15\%$ per %FIMA) reflects the thermal and mechanical stability at lower power density, demonstrating that operation at reduced temperatures and power densities achieves not only acceptable fuel performance within design margins, but also predictable, low-scatter performance desirable for reliability qualification.

A profound finding emerges from the Battelle and SP-100 data when segregated by grain size and fabrication method. Fine-grained fuel exhibits 2–3 times higher swelling rates than coarse-grained fuel at identical temperature and burnup. Oxygen content and fabrication process show comparable influence, as demonstrated by the 30% swelling difference between K-1 and K-3 in the BR-10 program at identical temperature. These observations establish microstructural variables as primary controls on fuel performance, not secondary corrections. Any empirical or mechanistic model omitting grain size, pore morphology, and impurity content cannot capture the observed variation across fuel specifications and fabrication methods.




Figure 3.3: Volumetric swelling versus burnup for all historical UN irradiation programs, color-coded by test program. Program-specific clustering reflects irradiation temperature and power density effects.

3.5 Empirical Correlations, Limitations, and Mechanistic Necessity

3.5.1 Ross Swelling Correlation

The comprehensive historical database from 1960–2002 provided the foundation for empirical approaches to predict volumetric swelling in UN fuel as a function of burnup, temperature, and fuel properties. Ross, El-Genk, and Matthews [57] developed one of the first systematic empirical correlations for UN fuel swelling, compiled from 75 data points from SNAP-50 (CANEL and ORNL testing). Their dataset spanned fuel densities of 90–98% theoretical density, volume-averaged temperatures from 1200 to 1600 K, and burnups between 0.05 and 4.58 at%. The resulting power-law correlation expresses volumetric swelling as

$$\frac{\Delta V}{V} = 4.71 \times 10^{-11} T_{av}^{3.12} B_u^{0.830} \rho^{0.5} \quad (3.2)$$




Figure 3.4: Fission gas release versus burnup for all historical UN irradiation programs, color-coded by test program. Different symbol shapes denote distinct testing programs.

where $\Delta V/V$ represents volumetric swelling (fractional), T_{av} is the volume-averaged fuel temperature in Kelvin, B_u is the fuel burnup in at%, and ρ is the as-fabricated fuel density as a fraction of theoretical density.

This power-law structure encodes three key controlling dependencies identified from the historical data. Volume-averaged fuel temperature emerges as the dominant control, capturing the transition where at low-to-intermediate temperatures, swelling increases rapidly, but at high temperatures (>1900 K), saturation effects plateau the growth. The burnup exponent reflects sublinear accumulation, where swelling increments diminish per unit incremental burnup as saturation mechanisms activate and interconnected porosity pathways enable gas escape. Fuel density enters through a square-root scaling, reflecting competing mechanisms where higher density increases fission gas production but simultaneously modulates initial porosity structure, grain boundary density, and mechanical accommodation pathways.

The Ross correlation achieved $\pm 25\%$ accuracy at moderate-to-high burnups (>1.12 at%), where saturation had stabilized data scatter. At very low burnups (<0.5 at%), scatter reached $\pm 60\%$, reflecting the erratic early swelling transient. The Ross formulation represented a major advance, providing the first quantitative framework for predicting UN

swelling trends from an expanded dataset. However, several limitations were inherent. The correlation does not explicitly account for microstructural variables such as grain size or pore morphology, parameters shown by Battelle investigations to vary swelling by $2\text{--}3\times$ at constant temperature and burnup. The fitted domain (1200–1600 K, 0.05–4.6 at%, 90–98% TD) constrained applicability, and extrapolation beyond these ranges lacks mechanistic justification. Thermal transients and time-varying histories, common in operational reactors, were not represented. The assumed separability of parametric effects breaks down when microstructural evolution couples temperature, burnup, and density dependencies.

3.5.2 Machine Learning Swelling Models

Recent work by Chen et al. [59] examined whether machine learning methods could improve upon empirical approaches by leveraging more flexible functional forms and explicit incorporation of microstructural variables. Using kernel ridge regression and symbolic regression, the analysis tested whether nonparametric algorithms could more effectively capture complex swelling behavior across the same historical UN database that informed the Ross correlation. The machine learning framework included a richer set of input features such as burnup, irradiation temperature, power density, and fuel geometry.

Across the datasets examined, machine learning methods achieved lower root mean square errors and demonstrated more generalizability compared to the original Ross correlation, particularly when trained on homogeneous datasets with consistent experimental conditions. Sensitivity analysis revealed dataset-dependent feature importance, with temperature and burnup emerging as dominant controls in low-to-moderate temperature regimes, though the relative importance varied across experimental test series. The machine learning models captured nonlinear interactions between temperature, burnup, and density that simple power laws cannot represent.

However, significant limitations persisted that prevent deployment of machine learning approaches for uncertainty quantification in mechanistic modeling. While machine learning models improved predictive accuracy within the training data envelope, they remained purely empirical tools limited in mechanistic insight into the underlying physics of swelling. Dataset heterogeneity and microstructural variability introduced substantial challenges, with single trained models showing poor performance when applied to data from different experimental programs with varying fuel geometries, cladding designs, and irradiation conditions. Critically, machine learning models cannot reliably extrapolate beyond the training domain or adapt to novel fuel compositions and microstructures without retraining on new experimental data. These limitations demonstrate that while machine learning offers incremental improvements in empirical fitting accuracy, neither machine learning nor traditional correlations provide a path toward accelerated fuel qualification grounded in mechanistic understanding and rigorous uncertainty quantification.

3.5.3 Impact of Mechanistic Modeling

Comparing Ross and ML approaches highlights both the potential and the ceiling of purely empirical methods. ML advances accuracy through algorithmic flexibility and explicit multiparameter inclusion, yet both frameworks share fundamental constraints. Neither can confidently predict swelling under conditions or fuel specifications substantially removed from the historical dataset. Operational reactors with transient thermal histories, varying power profiles, and stress-driven acceleration of swelling mechanisms introduce complexities absent in controlled laboratory irradiations that generated the training data.

The apparent advantages of empirical and ML methods are speed and low computational cost, however this comes at the cost of mechanistic versatility and extrapolation capacity. Bayesian uncertainty quantification frameworks require decomposition of uncertainties into contributions from measurement error, model uncertainty, and irreducible material variability.

The transition from empirical correlation or ML curve-fitting to mechanistic modeling is thus not incremental refinement but a step toward experimental extrapolation and future design. Mechanistic models, grounded in first-principles defect kinetics, bubble nucleation thermodynamics, grain-boundary mechanics, and porosity percolation theory, enable explicit incorporation of micro-structural variables (grain size, pore morphology), independence from most fitted parameter ranges, and rigorous uncertainty decomposition. Calibrated via Bayesian inverse analysis against historical data, mechanistic models preserve the empirical constraints while enabling extrapolation with quantified confidence bands. This capability underpins the research direction of the present dissertation and the subsequent development of physics-based fuel performance frameworks.

The historical UN irradiation database establishes four fundamental findings that drive the mechanistic modeling agenda developed in subsequent chapters. First, temperature is the dominant control on swelling and gas release, with exponential or power-law behavior reflecting thermally-activated diffusion. Temperature acceleration factors of 5–10 times across the 1000–2100 K range demand models capturing the underlying Arrhenius kinetics precisely. Second, micro-structural variables (grain size, density, pore morphology) produce swelling variations of $2\text{--}3\times$ at constant temperature and burnup, establishing these as primary design levers rather than second-order corrections. Empirical correlations cannot capture this without explicit microstructure modeling. Third, saturation swelling phenomena at high temperature represent a qualitatively distinct regime where rapid bubble coalescence and porosity percolation enable quasi-equilibrium between gas production and release. Fourth, fabrication process and irradiation history (time-varying temperature profiles, stress transients, damage-rate variations) produce substantial performance variation inadequately represented in static empirical correlations.

Chapter 4

Mechanistic Fission Gas Release and Swelling Model for Uranium Mononitride

4.1 Introduction and Mechanistic Model Motivation

Mechanistic fuel performance modeling represents an advancement over empirical correlations for predicting FGR and swelling in nuclear fuels. Chapter 3 synthesized six decades of experimental UN irradiation data (SNAP-50, BR-10, SP-100, and others) and developed empirical correlations relating operating conditions (temperature, power density, burnup, fuel density) to performance outcomes, volumetric swelling, gas release. While these correlations provide useful interpolations within their experimental domains, they suffer from two fundamental limitations for advanced fuel qualification.

First, the empirical database exhibits order-of-magnitude scatter at constant temperature and burnup. Second, and more critically for AFQ programs, empirical correlations cannot reliably extrapolate beyond the experimental domain. New fuel forms, operating regimes, or geometries differ systematically from those under which empirical fits were primarily developed.

Mechanistic models overcome these limitations by decomposing fission gas release and swelling into constituent physical processes: generation of fission gas atoms through fission events, migration of these atoms through the fuel lattice under thermal and irradiation-enhanced diffusion, nucleation and growth of bubbles at heterogeneous sites (grain boundaries, dislocations), resolution of gas atoms by energetic fission fragments, and interconnection of grain boundary bubbles leading to gas release to the fuel plenum. By capturing these mechanisms explicitly, mechanistic models become sensitive to variations in fuel microstructure (grain size, dislocation density, etc.), operating conditions, and fuel form—dependencies.

Critically, mechanistic models enable rigorous sensitivity analysis and uncertainty quan-

tification. By making model structure and parameter dependence explicit, mechanistic frameworks allow identification of which parameters most strongly influence predictions, quantification of how parameter uncertainties propagate to output uncertainties, and calibration of uncertain parameters using Bayesian inverse analysis with historical data. This probabilistic framework provides quantified confidence bands on predictions, a requirement for regulatory decision-making and confidence in extrapolation to novel conditions. Initial validation against experimental data typically reveals parameter sensitivities and may necessitate parameter refinement or temperature-dependent parameterization to achieve quantitative agreement across diverse operating conditions—a process that leverages mechanistic transparency to systematically identify and address model limitations.

This chapter presents a comprehensive mechanistic model for fission gas release and swelling in UN fuel, implemented within the BISON multiphysics fuel performance code through the Simple Integrated Fission Gas Release and Swelling (SIFGRS) model framework developed by Rizk et al. [60]. The model integrates lower-length-scale (LLS) calculations, primarily DFT, MD, and cluster dynamics (CD) simulations, to inform critical mechanistic parameters such as gas and self-defect diffusivities, bubble surface and grain boundary energies, and resolution rates by fission fragments. The model is validated against historical UN irradiation data, with emphasis on SNAP-50 test cases that provide the most comprehensive microscopic and integral fuel pin performance data available for UN.

4.2 Mechanistic Model Framework: SIFGRS Overview and Bubble Population Tracking

4.2.1 Fuel Performance Code

BISON [3] is a finite element based nuclear fuel performance code developed at Idaho National Laboratory (INL) as part of the Nuclear Energy Advanced Modeling and Simulation (NEAMS) [61] program. The code solves fully coupled nonlinear partial differential equations governing thermo-mechanics and species diffusion across one-, two-, and three-dimensional geometries, enabling multiphysics simulations of nuclear fuel behavior under irradiation. BISON is built on the Multiphysics Object-Oriented Simulation Environment (MOOSE) [62] computational framework and incorporates material and behavioral models for thermal properties, creep, swelling, porosity, and mechanical interaction for a variety of fuel forms. The code is designed with flexibility to accept LLS-informed models from atomistic calculations, operate across multiple length scales and dimensions. BISON was chosen as the primary fuel performance code due to increased regulatory use and aligned directly with the AFQ framework.

4.2.2 Conceptual Framework

The SIFGRS model explicitly tracks fission gas segregation into spatially distinct bubble populations reflecting their distinct nucleation sites and growth kinetics. Intragranular bubbles (in the lattice) experience preferential trapping at dislocations and point defects, while grain boundary bubbles benefit from enhanced diffusion and serve as the primary release pathway to the plenum. This population-based approach captures the selective migration pathways and phase transitions that govern fuel performance.

The SIFGRS model for UN distinguishes four primary bubble populations [60]:

1. **Bulk bubbles:** Intragranular bubbles nucleating homogeneously in undisturbed fuel lattice at zero-energy sites (perfect lattice). These bubbles typically form at low temperatures where thermally-enhanced diffusion is minimal and bubbles remain small (radii < 30 nm). Their contribution to volumetric swelling is secondary at elevated temperatures but can become significant at low-to-moderate temperature regimes where dislocation saturation limits dislocation bubble growth.
2. **Dislocation bubbles:** Intragranular bubbles nucleating and growing preferentially on the dislocation network. Dislocations provide energetically favorable nucleation sites (lower barrier for cluster formation) and enhanced diffusion pathways along dislocation cores. Dislocation bubbles dominate the swelling contribution at elevated temperatures ($T > 1300$ K) where they grow to large radii (50–200 nm) and constitute the majority of retained gas volume.
3. **Grain boundary bubbles:** Bubbles accumulating at grain boundaries as gas atoms migrate via temperature gradients and concentration-driven diffusion. These bubbles grow preferentially along grain boundaries due to lower surface energy at the grain boundary-bubble interface compared to bubble-lattice interfaces. Crucially, grain boundary bubbles represent the transition state between gas retention and release. As grain boundary bubbles coalesce and grow, they interconnect to form continuous porosity pathways that become permeable to gas flow to the fuel surface.
4. **Released gas:** Once grain boundary bubble networks interconnect, generated fission gas escapes to the fuel plenum and is no longer retained as bubbles within the fuel matrix.

For each bubble population, the model tracks three key state variables: bubble number density n_i (m^{-3}), mean bubble radius r_i (m), and total gas atoms N_i (primarily ^{133}Xe and ^{85}Kr , treated as a single effective species). Global conservation of gas atoms is enforced such that the sum of retained gas in all populations plus released gas equals the total generated fission gas.

4.2.3 Evolution Equations and Rate Processes

Each bubble population evolves according to coupled differential equations balancing multiple rate processes:

$$\frac{dn_i}{dt} = \dot{n}_{\text{nuc},i} - R_{\text{res},i} - \dot{n}_{\text{coal},i} \quad (4.1)$$

where $\dot{n}_{\text{nuc},i}$ is the nucleation rate (new bubbles forming per unit volume per unit time), $R_{\text{res},i}$ is the resolution rate (bubbles destroyed by fission fragment collisions), and $\dot{n}_{\text{coal},i}$ is the coalescence rate (bubble merging and sintering).

The mean radius of bubbles in population i evolves via vacancy absorption or gas atom release:

$$\frac{dr_i}{dt} = v_{\text{atom}} (J_{\text{gas},i} - J_{\text{vac},i} - J_{\text{res},i}) \quad (4.2)$$

where v_{atom} is the atomic volume of the fuel, $J_{\text{gas},i}$ is the flux of gas atoms arriving at bubbles (from diffusion), $J_{\text{vac},i}$ is the flux of vacancies being absorbed by bubbles (contributing lattice expansion to accommodate pressurized gas), and $J_{\text{res},i}$ is the flux of gas atoms being knocked out by resolution. The detailed rate processes (nucleation $\dot{n}_{\text{nuc},i}$, resolution $R_{\text{res},i}$, coalescence $\dot{n}_{\text{coal},i}$) are developed in Sections 4.4 and 4.5. The full BISON implementation (detailed in Appendix A of Rizk et al. [60]) couples these equations with spatial diffusion via an effective diffusivity concept. The balance between these competing fluxes determines whether bubbles grow, shrink, or remain at quasi-steady-state pressure.

4.3 Fission Gas Diffusion: Three Mechanistic Regimes

Fission gas atom migration through UN is governed by fundamentally different physical mechanisms depending on temperature and irradiation dose rate. Rather than assuming a single temperature-dependent diffusivity, the SIFGRS model explicitly decomposes the total effective diffusivity into three additive components, each dominating in different physical regimes. This decomposition is essential for capturing the transition in diffusion mechanism that underlies breakaway swelling phenomena observed in UN at high temperatures.

4.3.1 Thermal Equilibrium Diffusivity: D_1

At low fission rates and sufficiently high temperatures, fission gas atom migration is controlled by thermally-activated diffusion through the perfect lattice, limited by the activation energy barrier for forming and migrating point defects (typically vacancies). This thermal component is given by a standard Arrhenius form:

$$D_1 = D_{0,1} \exp\left(-\frac{E_{A,1}}{k_B T}\right) \quad (4.3)$$

where $D_{0,1}$ is the kinetic coefficient, $E_{A,1}$ is the activation energy barrier, k_B is Boltzmann's constant, and T is absolute temperature. For xenon diffusion in UN, lower-length-scale calculations using DFT combined with empirical potential simulations have provided estimates of these parameters. Cooper et al. [36] calculated Xe migration barriers in UN by mapping the minimum-energy pathway for Xe displacement through the lattice, accounting for both uranium and nitrogen sublattices and the crystal symmetry.

The physical interpretation of D_1 dominance is straightforward. At temperatures below approximately 1000 K, thermal energy is insufficient to populate the activated states required for diffusion, so lattice diffusion rates remain negligible and gas atoms remain spatially localized. As temperature increases, D_1 increases exponentially. However, at the same time, the irradiation environment introduces additional diffusion mechanisms that become progressively more important.

4.3.2 Irradiation-Enhanced Diffusivity: D_2

Even at moderate temperatures where thermal diffusivity is low, neutron irradiation continuously generates point defects (vacancies and interstitials) through displacement cascades. These excess point defects create additional diffusion pathways. Gas atoms can migrate by exchanging with vacancies that are in excess concentration due to irradiation, and interstitial clusters can propagate through the lattice. The enhanced diffusivity from this excess defect concentration is captured in the D_2 component.

Unlike D_1 , which is independent of irradiation conditions, D_2 depends explicitly on the fission rate density (or equivalently, dose rate) through its dependence on the steady-state excess defect concentration. At steady state, the rate of defect production by collision cascades is balanced by the rate of defect recombination and annihilation at sinks (dislocations, grain boundaries, free surfaces). This balance yields a point defect concentration that scales as $\sqrt{\dot{f}}$, where $\dot{f}$ is the fission rate density. Consequently, the irradiation-enhanced diffusivity takes the form:

$$D_2 = D_{0,2} \sqrt{\dot{f}} \exp\left(\sum_{i=0}^3 c_i \left(\frac{1000}{T}\right)^i\right) \quad (4.4)$$

where the coefficients c_i represent polynomial coefficients from fitting cluster dynamics calculations of excess point defect concentrations as a function of temperature. The choice of polynomial form in inverse temperature reflects the complex temperature dependence of defect production and recombination rates in cluster dynamics models. For UN, Cooper et al. [36] performed cluster dynamics simulations to generate the temperature-dependent coefficients. For xenon specifically, cluster dynamics calculations indicate that D_2 remains negligible compared to $D_1 + D_3$ across all relevant irradiation conditions [36]; the form

in Equation 4.4 is retained for completeness and represents defect-enhanced mobility for vacancy diffusion.

The key distinction between D_1 and D_2 is their dependence on irradiation conditions: D_1 is a material property independent of neutron flux, while D_2 is directly proportional to fission rate and therefore higher in high-power-density fuel or in fast neutron spectra with higher fission rates. This means that identical fuel at identical temperature but different power densities exhibits markedly different fission gas release rates due to diffusion enhancement. Empirical correlations that do not explicitly account for power density sensitivity miss this important mechanism.

4.3.3 Irradiation-Induced Athermal Diffusivity: D_3

The third diffusivity component arises from direct ballistic mixing within collision cascade volumes. Unlike D_2 , which assumes steady-state defect concentrations, D_3 captures direct ballistic displacement of gas atoms and lattice atoms during the collision cascade process itself. This mechanism occurs even at very low temperatures where thermal activation is negligible—hence the term “athermal.” The athermal diffusivity scales linearly with fission rate density:

$$D_3 = D_{0,3}\dot{f} \quad (4.5)$$

where $D_{0,3}$ is a coefficient characterizing the cross-section for ballistic displacement per fission event. This linear dependence reflects the fact that higher fission rates produce more cascades, and in each cascade there is a fixed probability of ballistic mixing.

The physical significance of D_3 is profound. It explains why fission gas release can occur at low temperatures below 1000 K) where thermal diffusion is negligible and irradiation-enhanced diffusion is still modest. At low temperatures, D_3 dominates the total diffusivity, enabling gas atom displacement and bubble nucleation even under conditions where a purely thermal model would predict near-zero gas mobility. This athermal mechanism is responsible for the low-temperature gas release plateau observed in historical UN data (Chapter 3).

4.3.4 Total Diffusivity and Transition Regimes

The total fission gas diffusivity is the sum of the three mechanistic components:

$$D_{\text{total}} = D_1 + D_2 + D_3 \quad (4.6)$$

The relative dominance of each component depends on temperature, fission rate, and burnup. At low temperatures ($T < 1000$ K) and moderate fission rates typical of earlier-generation reactors, D_3 dominates, yielding relatively constant (approximately linear) fission gas release independent of further temperature decreases. At intermediate temperatures (1000–1300 K), all three components contribute, and the behavior transitions. At high

temperatures ($T > 1400$ K), D_1 becomes the dominant mechanism, exhibiting strong Arrhenius temperature dependence.

This compositional transition in diffusivity mechanisms underlies the non-monotonic swelling behavior observed in UN. At temperatures just below 1000 K, swelling accelerates due to enhanced diffusivity from D_2 and D_3 , enabling rapid bubble nucleation and growth. As temperature increases toward 1500 K, D_1 begins to dominate, and the effective activation energy barrier increases (since D_1 has a higher E_A than the combined effect of D_2 and D_3). Simultaneously, thermal creep strength decreases, allowing plastic relaxation of stresses generated by bubble swelling. These competing effects create the maximum in swelling rate near 1000–1300 K observed in BR-10 data, which simple thermal-diffusion models fail to reproduce.

For grain boundary diffusion, the model assumes an enhancement factor $f_{D,GB}$ multiplying the bulk thermal diffusivity, reflecting pipe diffusion along grain boundary dislocation cores. Typical values from the literature suggest $f_{D,GB} \sim 10^6$ (grain boundary diffusion 1 million times faster than bulk), though significant uncertainty remains. This parameter is treated as uncertain and included in the sensitivity analysis (Chapter 6).

4.4 Bubble Nucleation and Growth Kinetics

4.4.1 Homogeneous Nucleation: Bulk Bubbles

At the earliest stages of irradiation, before significant nucleation of dislocation or grain boundary bubbles occurs, fission gas atoms must overcome an energy barrier to nucleate a stable bubble cluster. Homogeneous nucleation in the perfect lattice occurs when individual gas atoms, diffusing randomly through the fuel, encounter each other and form small clusters. The probability of cluster formation depends on the gas atom concentration and the critical cluster size needed to achieve thermodynamic stability.

The nucleation barrier arises from competition between two energetic contributions: surface energy (which increases as cluster size increases, penalizing small clusters) and bulk gas pressure (which decreases with cluster size, favoring larger clusters). For a spherical bubble of radius r_b , the work required to form a stable nucleus is given by the classical nucleation theory formula:

$$W_{\text{nuc}} = \frac{4}{3}\pi r_b^3 (p_{\text{eq}} - p_{\text{ext}}) + 4\pi r_b^2 \gamma \quad (4.7)$$

where $p_{\text{eq}} = \frac{2\gamma}{r_b}$ is the equilibrium pressure inside the bubble (Young-Laplace equation), p_{ext} is the external hydrostatic stress, and γ is the bubble-matrix surface energy. The first term represents the bulk energy gain (pressure difference times volume) and the second represents the surface energy cost. The critical radius r_* at which the nucleation barrier reaches its maximum is found by setting $\frac{dW_{\text{nuc}}}{dr_b} = 0$, yielding:

$$r_* = \frac{2\gamma}{p_{\text{eq}} - p_{\text{ext}}} \quad (4.8)$$

Only clusters exceeding this critical size grow spontaneously while smaller clusters redissolve. The nucleation rate (number of stable clusters forming per unit volume per unit time) depends on the gas atom concentration and the probability that multiple atoms aggregate to exceed the critical size. Following standard nucleation theory, the homogeneous nucleation rate for bulk bubbles is approximated as:

$$\dot{n}_{\text{nuc},1} = \nu_1 C_1^2 \quad (4.9)$$

where ν_1 is an effective nucleation rate constant parameterizing the probability of dimer formation (first stable cluster), and C_1 is the single gas atom concentration. The nucleation rate constant ν_1 implicitly captures the density of nucleation sites and the collision frequency of gas atoms in the lattice. The quadratic dependence on C_1 reflects the probability that two gas atoms encounter each other.

4.4.2 Heterogeneous Nucleation: Dislocation Bubbles

Dislocations and other extended defects provide energetically favorable nucleation sites by reducing the surface energy barrier. Bubbles at dislocation cores experience partial interface with both the lattice and the dislocation, reducing the effective surface energy and thus the nucleation barrier. As a result, dislocation bubbles nucleate at lower gas atom concentrations than bulk bubbles and begin growing earlier in irradiation.

The number density of dislocation bubbles is assumed to be proportional to the dislocation line density:

$$n_2 = \alpha_\rho \rho_{\text{disl}} \quad (4.10)$$

where α_ρ is the bubble-to-line density ratio (bubbles per meter of dislocation line), and ρ_{disl} is the total dislocation density (in m^{-2}). The coefficient α_ρ captures the saturation of nucleation sites along dislocations—dislocations can accommodate only so many bubbles per unit length before geometric packing limits are reached. Literature values and LLS predictions suggest $\alpha_\rho \sim 10^5 - 10^6$ bubbles/meter, representing the saturation of nucleation sites along dislocation lines in Chapter 6).

The dislocation density itself evolves under irradiation according to balance between irradiation-induced hardening (creating new dislocation loops) and dynamic recovery (dislocation annihilation). The evolution model for dislocation density is:

$$\frac{d\rho_{\text{disl}}}{dt} = \dot{\rho}_{\text{prod}} - \dot{\rho}_{\text{recomb}} \quad (4.11)$$

where the production rate $\dot{\rho}_{\text{prod}}$ typically follows $\dot{\rho}_{\text{prod}} \propto \sqrt{\dot{f}}$ (from collision cascade density scaling), and the recombination rate $\dot{\rho}_{\text{recomb}} \propto \rho_{\text{disl}}$ (representing mutual annihilation of

complementary dislocation types). At early burnup, $\dot{\rho}_{\text{prod}}$ dominates and dislocation density increases. At high burnup, a quasi-steady-state is reached where production and recombination balance, yielding a saturation dislocation density. In SIFGRS implementations for UN, dislocation density is typically held constant during transient simulations based on post-irradiation transmission electron microscopy (TEM) measurements, rather than dynamically evolved. This simplification is discussed in Section 4.8.

4.4.3 Bubble Coalescence and Mean Size Evolution

As bubbles grow, their number density may decrease through coalescence—the physical merging of adjacent bubbles. In the mean-field approximation used by SIFGRS, coalescence is modeled by coupling the evolution of number density to the evolution of mean radius. As the mean radius increases, the surface area density available for gas atoms decreases (fewer, larger bubbles occupy the same volume), and the number density correspondingly decreases. This is captured through a coalescence rate:

$$\left. \frac{dn_i}{dt} \right|_{\text{coal}} = -\alpha_{\text{coal}} n_i \frac{dr_i}{dt} \quad (4.12)$$

where α_{coal} is a dimensionless coalescence coefficient of order 1 representing the rate at which number density decreases as mean radius increases. This phenomenological form assumes coalescence rate is proportional to both bubble number density and growth rate, capturing both Ostwald ripening of small bubbles and direct mechanical coalescence as bubbles grow closer together.

For dislocation bubbles specifically, hard-sphere packing geometry [63] suggests a coalescence coefficient of approximately $\alpha_{\text{coal}} \approx 4.2\sqrt{2}\pi r_{eq}^2$, derived from geometric packing arguments. For bulk bubbles, coalescence is often assumed to be less efficient (smaller α_{coal}), reflecting the more dispersed spatial arrangement and lower collision rates compared to bubbles confined to dislocation networks.

4.5 Resolution of Fission Gas from Bubbles

Resolution is the process by which energetic fission fragments knock gas atoms out of bubbles back into the surrounding lattice, thereby reducing the number of retained gas atoms. This process is particularly important at low temperatures where thermal remigration of knocked-out atoms to other bubbles is unlikely, and at high burnups where accumulated bubbles experience frequent fission fragment collisions.

For non-oxide fuels such as UN with high thermal conductivity, fission fragment energy is dissipated primarily through electronic stopping (ionization) rather than nuclear stopping (atomic collisions). The relatively thin region in which atomic collision cascades occur means that only bubbles in close proximity to the fission track experience significant resolution. This limit is captured by the homogeneous resolution model. Resolution occurs when fission

fragments pass through or near a bubble and transfer momentum directly to trapped gas atoms.

The total resolution rate from bubble population i is given by:

$$R_{\text{res},i} = \alpha_0(r_b) N_i n_b \dot{f} \quad (4.13)$$

where $\alpha_0(r_b)$ is the resolution parameter in units of $[10^{-25} \text{m}^6/\text{fission}]$, N_i is the number of gas atoms per bubble, n_b is the bubble number density, and $\dot{f}$ is the fission rate density. The resolution parameter is a decreasing function of bubble radius, reflecting the physical fact that small bubbles are more easily emptied (fewer atoms need to be displaced) while large bubbles are relatively resistant to resolution (more atoms distributed over a larger cross-section).

Matthews et al. [64] calculated the resolution parameter for UN using binary collision approximation (BCA) simulations with the 3DOT code, accounting for fission fragment energy distributions from ENDF libraries and the uranium and nitrogen nuclear cross-sections. The radius dependence was fit to the functional form:

$$\alpha_0(r_b) = 10^{-25} (2.64 - 2.02 \exp(-2.61 \times 10^{-9} r_b)) \quad [\text{m}^6/\text{fission}] \quad (4.14)$$

where r_b is the bubble radius in meters. At small radii, the exponent is small, so $\exp(-\text{small})$ remains close to 1, yielding $\alpha_0 \approx 0.62 \times 10^{-25} \text{ m}^6/\text{fission}$. At large radii, the exponent becomes more negative, so $\exp(-\text{large})$ approaches 0, yielding $\alpha_0 \approx 2.64 \times 10^{-25} \text{ m}^6/\text{fission}$, approximately a 4-fold increase. This radius dependence is critical for capturing the quasi-steady-state pressure of intragranular bubbles at low temperatures, where resolution maintains bubbles near equilibrium rather than allowing unlimited pressurization.

The resolution coefficient α_{res} and efficiency factors represent significant uncertainties in the model. Experimental measurements of resolution rates are sparse and indirect, typically inferred from post-irradiation examination of bubble populations rather than measured directly. Furthermore, resolution efficiency likely varies with bubble size, pressure, and local radiation damage state—dependencies not captured in the current constant-efficiency parameterization. Subsequent validation and parametric investigation will assess the sensitivity of fission gas release predictions to these resolution parameters and identify whether refinement is warranted for specific temperature or burnup regimes.

4.6 Grain Boundary Bubble Evolution and Interconnection-Driven Release

4.6.1 Migration Toward and Accumulation at Grain Boundaries

Fission gas atoms diffusing through the fuel eventually reach grain boundaries, particularly under temperature gradients that create concentration-driven diffusion (grain-directed

diffusion). The diffusion flux toward grain boundaries is driven by gradients in gas concentration and stress fields around bubbles, as well as by temperature gradients. Once gas atoms reach the grain boundary, the energy landscape changes. The grain boundary-bubble interface has lower energy than the bulk-bubble interface, making grain boundaries thermodynamically favorable sites for bubble nucleation and accumulation.

The flux of gas atoms arriving at grain boundaries depends on the bulk gas atom concentration and the diffusivity. In SIFGRS, this flux is approximated as:

$$J_{\text{GB}} = D_{\text{eff}} \nabla C_1 \quad (4.15)$$

where D_{eff} is an effective diffusivity incorporating both bulk diffusion toward the grain boundary and pipe diffusion along the boundary, and ∇C_1 is the concentration gradient of single gas atoms where D_{eff} incorporates pipe diffusion enhancement ($f_{D,\text{GB}} \sim 10^6$, discussed in Section 4.3.4). The grain boundary bubble number density increases as gas atoms arrive and coalesce:

$$\frac{dn_3}{dt} = J_{\text{GB}} A_{\text{GB}} n_{\text{sites}} \quad (4.16)$$

where A_{GB} is the grain boundary area per unit fuel volume and n_{sites} is the available nucleation site density on grain boundaries.

4.6.2 Saturation, Interconnection, and Gas Release

As grain boundary bubbles grow and coalesce through continued gas atom arrival, they eventually occupy a significant fraction of the grain boundary area. At low coverages (fraction of grain boundary area covered by bubbles less than saturation threshold f_{sat}), bubbles remain isolated and gas is retained. However, once the fractional coverage exceeds f_{sat} , neighboring bubbles begin to coalesce directly into a continuous interconnected pathway along the grain boundary.

The percolation transition from isolated to interconnected geometry is a geometric phase transition. Below the saturation threshold, bubbles are separated by fuel bridges. Above it, the pathway becomes continuous. Experimentally, saturation coverage varies depending on bubble morphology, grain boundary misorientation, and boundary impurities, but typical values from oxide fuel literature suggest $f_{\text{sat}} \approx 0.5$ – 0.6 [65]. The SIFGRS model uses f_{sat} as an adjustable parameter, currently set nominally to 0.5 but retained as an uncertain parameter in sensitivity analysis (Chapter 6) to account for material system variability.

The saturation coverage threshold represents a geometric percolation parameter with inherent uncertainty. The value $\theta_{\text{sat}} = 0.5$ derives from percolation theory for randomly distributed circles on a plane, but actual grain boundary bubble distributions exhibit spatial correlation, size distributions, and grain boundary groove geometries that may lower the effective saturation threshold. Values ranging from 0.25 to 0.5 are physically plausible depending on bubble morphology and grain boundary character. Sensitivity to this parameter will be assessed in subsequent parametric studies.

Once interconnection occurs, the continuous grain boundary porosity network becomes permeable to gas flow. Any fission gas generated in the fuel near an interconnected boundary is essentially released immediately, while gas generated far from an interconnected boundary must still migrate to reach the pathway. This creates a transition from the regime where gas is retained (pre-interconnection) to the regime where a substantial fraction of generated gas is released (post-interconnection). In BISON simulations, once interconnection is triggered, the grain boundary bubble population is assumed to instantly release all subsequent generated gas at the rate of fission, leading to rapid increase in cumulative FGR.

The physical mechanism driving this transition is mechanical. As grain boundary bubble coverage increases, the fuel bridges between bubbles become narrower and weaker. Eventually, the internal pressure in the bubbles exceeds the tensile strength of the remaining fuel bridges, leading to mechanical failure of the bridge and bubble coalescence. This is a stress-driven process that depends on both the bubble pressure and the mechanical properties of the fuel.

4.7 Model Parameters and Uncertainty Sources

The mechanistic model depends on 22 uncertain parameters spanning multiple physical processes in Chapter 6, of which global sensitivity analysis identifies 16 as most influential). Rather than relying entirely on fitting these parameters to historical data, the SIFGRS framework prioritizes physics-based estimation from lower-length-scale (LLS) calculations where possible, reserving empirical fitting for phenomenological quantities that cannot be calculated from first principles.

4.7.1 Lower-Length-Scale Calculations

LLS calculations—DFT, MD, and CD—provide quantitative parameter estimates with systematic uncertainties from methodological choices.

Cooper et al. [36] calculated Xe diffusion coefficients in UN using DFT nudged elastic band (NEB) techniques to map migration barriers, followed by empirical potential (EP) calculations for wider temperature sampling. Cluster dynamics simulations integrated these inputs to model steady-state point defect concentrations under irradiation, yielding the polynomial coefficients for D_2 in Equation 4.4. Systematic uncertainties of 10–100% arise from exchange-correlation functional choices, empirical potential limitations, and finite-size effects.

Kocevski et al. [66] performed MD calculations of UN surface and grain boundary energies by creating supercells with free surfaces or grain boundaries, performing MD at 1000 K for relaxation, and computing interface energies as excess energy per unit area. Five representative grain boundaries yielded an average grain boundary energy of 1.139 J/m² (bulk-vacuum surface energy: 1.11 J/m²), with uncertainties of ± 10 –15% from empirical potential variations.

Matthews et al. [64] and Schneider et al. [67] calculated resolution parameters using binary collision approximation (BCA) and MD damage cascade simulations, tracking fission fragment trajectories and gas atom displacements. The radius-dependent resolution parameter in Equation 4.14 carries uncertainties of $\pm 30\text{--}50\%$ from fission fragment energy distributions (ENDF nuclear data) and cascade modeling approximations.

4.7.2 Empirical Parameters

Parameters not computable from first principles are measured experimentally or inferred from historical data:

- **Dislocation density** (ρ_{disl}): TEM measurements from BR-10, SNAP-50, and SP-100 (documented in Chapter 3) show 2–3 orders of magnitude scatter ($10^{12}\text{--}10^{14} \text{ m}^{-2}$) depending on temperature and dose. SIFGRS employs constant dislocation density as input rather than tracking time evolution.
- **Grain size** (r_g): Fabrication-controlled grain radii range from 5–35 μm in SNAP-50 samples. BISON simulations hold grain radius constant (neglecting grain growth) as a fixed input parameter.
- **Saturation coverage** (f_{sat}): The percolation threshold for grain boundary bubble interconnection depends on bubble morphology, packing geometry, and fuel bridge failure. Oxide fuel literature suggests 0.5–0.6; UN validation is absent. Sensitivity analysis spans 0.2–0.7 to reflect epistemic uncertainty.
- **Nucleation factors** (ν_1 , etc.): Homogeneous nucleation probability factors are fitted phenomenologically from classical nucleation theory or experimental matching, lacking direct measurement. Treated as high-uncertainty scaling factors.
- **Bubble-to-dislocation-line ratio** (α_ρ): Controls bubble nucleation density per unit dislocation length. Literature values: $10^5\text{--}10^6$ bubbles/meter; direct measurement challenging.
- **Solid fission product swelling**: Treated phenomenologically at fixed 0.5% per %FIMA, not mechanistically derived.

4.7.3 Primary Uncertainty Sources

Model uncertainties arise from four distinct sources:

(1) **LLS Calculation Uncertainties**: DFT functional selection, empirical potential transferability to high-temperature defect-rich systems, finite-size effects in supercells, and temperature extrapolation beyond calculated ranges introduce 10–100% variations in diffusion coefficients and 5–20% variations in surface energies [36].

(2) Experimental Data Scatter: Measurement technique variations (gas puncture, radiometric counting, mass spectrometry for FGR; density measurements or dilatometry for swelling) introduce $\pm 1\text{--}3\%$ measurement error. Fabrication variability—oxygen/carbon impurities, grain size, density—drives 30% swelling differences (BR-10 K-1 vs. K-3, Chapter 3) not mechanistically parameterized. Limited temperature (1300–1700 K) and burnup ($< 3\%$ FIMA) coverage to pressure-bonded fuel geometries and specific cladding alloys restricts extrapolation confidence for micro-reactor designs.

(3) Phenomenological Parameterizations: Four-population discretization (bulk, dislocation, grain boundary, released) uses mean-field approximations that miss rare large bubbles or percolating networks. Coalescence rate proportionality α_{coal} from geometric packing may not capture merging kinetics at high bubble densities. Fixed saturation coverage f_{sat} simplifies stress-dependent interconnection physics. Prescribed dislocation density neglects time/burnup evolution of hardening and recovery.

(4) Model Scope Limitations: Chemistry effects—nitrogen partial pressure dependence and stoichiometry and oxygen impurity effects (30% swelling sensitivity) are not mechanistically represented. Quasi-static evolution assumptions neglect power transients and thermal cycling. Anisotropic swelling and grain boundary strength degradation are simplified. These four uncertainty source categories are the primary drivers of parameter ranges in the sensitivity analysis (Chapter 6).

4.8 Model Limitations and Validity Domain

The SIFGRS mechanistic model, while comprehensive, incorporates several important limitations and simplifications that define its range of applicability:

1. **Stoichiometry and oxygen impurities:** The model assumes constant N/U stoichiometry and does not account for oxygen impurities or N-rich or U-rich off-stoichiometry. Chapter 3 documented that oxygen content affects swelling by up to 30% (K-1 vs. K-3 comparison), yet the current mechanistic model does not parameterize this effect. Nitrogen partial pressure dependence is similarly not mechanistically represented. Future refinements should include stoichiometry dependence of diffusivity and surface energy, requiring additional LLS calculations.
2. **Anisotropic swelling:** The model predicts spherically-averaged volumetric swelling. In practice, swelling often exhibits anisotropy (preferential expansion in certain crystallographic directions), particularly when grain boundaries are not randomly oriented. The polycrystal model in BISON averages over all grain orientations, but local anisotropy is not captured. This simplification has minimal impact on integral fuel pin performance but affects prediction of fuel relocation and gap dynamics in multi-pellet fuel column.

3. **Grain boundary strength and mechanical failure:** The grain boundary saturation model treats the percolation threshold as fixed. In reality, mechanical failure of fuel bridges depends on local stress state (which depends on thermal gradients, fission gas pressure, and cladding constraint) and on grain boundary strength (which degrades under irradiation). A more refined model would explicitly compute stress state and failure criterion, currently neglected.
4. **Fission fragment spectrum:** The resolution parameter (Equation 4.14) is calculated assuming uranium-235 fission spectrum. For fuel undergoing fast fission (fuel in fast reactor neutron spectrum) or breeding scenarios, the fission fragment spectrum differs, potentially affecting resolution rates. This effect is not captured in the current model.
5. **Gas solubility and trapping at point defects:** The model assumes negligible gas solubility in the fuel lattice at the temperatures of interest. In reality, xenon and krypton have small but finite solubility (on the order of 0.01 at. % or less at 1500 K based on thermodynamic estimates), and they can occupy substitutional or interstitial sites temporarily. Additionally, gas atoms can be transiently trapped at point defect complexes (divacancies, trivacancies) before escaping. These transient trapping effects are not mechanistically represented and would require multi-state cluster dynamics for accurate treatment. Trapping rate parameters represent substantial uncertainties, as they depend on bubble nucleation kinetics, point defect concentrations, and irradiation-enhanced diffusion—all of which exhibit order-of-magnitude variations depending on stoichiometry, grain size, and radiation damage state. The baseline trapping rates employed here are adapted from UO_2 literature values and may require temperature-dependent or burnup-dependent adjustment to match UN experimental observations. Parametric investigation of trapping rate sensitivity will be presented in Chapter 5.
6. **Low-burnup kinetics:** At very low burnup (<0.1 %FIMA), the mechanistic model may overpredict gas retention due to assumption of quasi-steady-state dislocation density. Early in irradiation, dislocation density is still evolving (building up from initially annealed state), and bubble nucleation kinetics depend sensitively on this evolving density. The model assumes equilibrium dislocation density is reached quickly, but in reality there may be an induction period where gas release is suppressed due to incomplete dislocation network formation.
7. **Fabrication porosity and pore migration:** Initial porosity from fuel fabrication (sintering pores) can serve as pre-existing gas collection sites and potentially migrate under thermal gradients or stress. The current model does not account for fabrication porosity evolution or its interaction with fission gas. This effect becomes important at high temperatures where thermal gradients are steep.
8. **Dislocation loop punching:** At high bubble pressures, grain boundary bubbles under certain stress states can nucleate dislocation loops, a mechanism not explicitly captured in the current model. This could affect local stress relief and migration rates.

9. **Fission-induced recoil and stoichiometric mixing:** During fission, recoil atoms can cause local disruption of stoichiometry, creating transient N/U imbalances. These recoil-induced point defects contribute to point defect concentrations (feeding into D_2) but are treated implicitly through cluster dynamics averaging rather than explicitly. This is acceptable as a mean-field approximation but misses potential non-linear effects if recoil-induced defect spikes reach high local densities.

These limitations do not invalidate the mechanistic model for its intended purpose—predicting FGR and swelling under representative micro-reactor conditions—but they do define boundaries beyond which model predictions should be interpreted with caution. Section 4.9 and subsequent validation chapters will quantify the magnitude of agreement between model predictions and experimental data, thereby empirically constraining the impact of these limitations.

4.9 Transition to Model Validation Against Experimental Data

The mechanistic framework developed thus far provides a rigorous physics-based structure for understanding fission gas release and swelling. However, mechanistic rigor alone does not guarantee predictive accuracy. Validation against experimental data is essential to demonstrate that:

1. The mechanistic processes identified (diffusion mechanisms, bubble nucleation, grain boundary saturation) are indeed the dominant physics governing observed fuel behavior;
2. The lower-length-scale calculation inputs (diffusion coefficients, surface energies, resolution rates) provide appropriate nominal values when integrated into the fuel-scale model;
3. The phenomenological parameters (nucleation factors, saturation coverage, coalescence rates) selected are physically reasonable and consistent with experimental observations;
4. Discrepancies between initial predictions and observations—which are expected given the substantial parameter uncertainties can be systematically diagnosed through gas partitioning analysis and parametric investigation to identify which mechanisms limit predictive capability and which parameters require refinement or temperature-dependent formulation.

Initial validation typically reveals parameter sensitivities that necessitate systematic parametric investigation. Chapter 5 presents baseline validation against SNAP-50 data, identifying where and why discrepancies arise. Chapter 5 then documents a comprehensive

parametric investigation exploring parameter adjustments designed to overcome identified limitations while maintaining physical grounding. This progression—from baseline validation through diagnostic analysis to parameter optimization—represents standard practice in mechanistic model development and leverages the mechanistic framework’s transparency to systematically improve predictive capability.

Specifically, the model validation will establish:

- **Integral validation:** Comparison of predicted volumetric swelling and fission gas release (as functions of temperature, burnup, and power density) to measured fuel pin values from SNAP-50 capsules;
- **Residual analysis:** Systematic examination of where and why the model over- or under-predicts, revealing whether discrepancies correlate with temperature, burnup, power density, or other variables that might indicate missing physics or incorrect parameter values.

The detailed validation analysis will be presented in a dedicated validation and implementation chapter that complements the mechanistic framework presented here. That chapter will include the complete BISON model geometry, mesh, boundary condition selections, power history reproductions, and detailed comparison plots and statistical metrics.

4.10 Chapter Summary

The SIFGRS mechanistic model for fission gas release and swelling in UN successfully integrates LLS calculations (DFT, MD, CD), fuel-scale mechanistic processes (diffusion, nucleation, growth, coalescence, resolution), and phenomenological parameterizations to create a physically-grounded predictive framework. The three-regime diffusion decomposition ($D_1 + D_2 + D_3$) explains the complex temperature dependence and power density sensitivity observed empirically. Grain boundary saturation and interconnection provide a mechanistic basis for the rapid transition from gas retention to gas release observed at high temperatures.

The model’s parameter structure explicitly distinguishes high-confidence parameters (derived from LLS calculations and widely-supported literature) from high-uncertainty parameters (phenomenologically parameterized or known to be variable). This transparent structure enables systematic sensitivity analysis (Chapter 6) to identify which parameters most strongly influence predictions and therefore warrant prioritized experimental or computational refinement. It further enables Bayesian calibration (Chapter 7) to constrain parameter distributions using historical data while maintaining physical grounding.

With this mechanistic foundation established, Chapter 5 validates baseline model predictions against SNAP-50 experimental data, identifying which mechanisms dominate different temperature and burnup regimes and where parameter adjustments or mechanistic refinements may be necessary to achieve quantitative agreement. Chapter 5 documents a systematic parametric investigation exploring parameter space to identify optimal parameter

combinations and diagnose physical bottlenecks limiting gas transport. Chapter 6 then leverages this mechanistic transparency to quantify parameter uncertainties and propagate them through fuel performance predictions, providing probabilistic confidence bands essential for regulatory decision-making.

Validation against SNAP-50 data will establish the mechanistic model's fidelity for the temperature, burnup, and power density ranges relevant to advanced micro-reactor applications, while residual analysis will identify directions for model refinement and prioritize future experimental or LLS work. The mechanistic framework thus serves as the foundation for both parametric investigation that reveals physical bottlenecks, sensitivity analysis that reveals parameter importance, and uncertainty quantification that provides probabilistic confidence bands on fuel performance predictions.

Chapter 5

UN Fission Gas Release and Swelling Model Implementation and Experimental Validation

5.1 Introduction and Validation Strategy

The mechanistic fission gas release and swelling model developed in Chapter 4 provides a rigorous physics-based framework for predicting UN fuel performance under irradiation. However, mechanistic models are only valuable insofar as they accurately capture dominant physical phenomena and provide reliable predictions across the experimental domain. This chapter validates the SIFGRS model against historical UN irradiation data, with particular emphasis on the SNAP-50 irradiation program, which provides the most comprehensive characterization of UN fuel performance available.

The validation strategy employs a hierarchical approach. First, model predictions of burnup, cladding temperature, and fuel centerline temperature are compared to experimental measurements to verify that boundary conditions faithfully represent actual irradiation conditions. Once model inputs are verified, integral fuel pin performance metrics, volumetric swelling and fission gas release, are compared to experimentally measured values across a range of irradiation conditions (temperature, burnup, power density). These comparisons establish whether the model captures overall fuel behavior. Second, examination of residual patterns, where and why the model over- or under-predicts, reveals whether discrepancies correlate with specific operating regimes, suggesting missing physics or parameter misprediction. Third, the consistency of model predictions with observed phenomena (e.g., breakaway swelling transitions at high temperature, plateau release at low temperature) validates not just the output accuracy, but also that the correct mechanisms are being captured.

Validation against SNAP-50 data is particularly valuable because the SNAP-50 program provides comprehensive characterization spanning a wide range of conditions relevant to

advanced reactor designs. The SNAP-50 experiments (reported by DeCrescente et al. [13] and Weaver et al. [46]) encompass fuel pin irradiations across fuel temperatures from 1213 K to 1839 K, burnups from 0.18% FIMA to 4.58% FIMA, and power densities from 0.32 kW/cm³ to 2.12 kW/cm³. Both integral measurements (fuel pin diameter change, density change, fission gas release) and microscopic characterization (transmission electron microscopy observations of bubble populations and grain growth) are available for many cases, enabling validation at multiple scales.

The goal of this chapter is threefold. First, establish that the SIFGRS mechanistic model provides credible predictions of UN fuel performance across a well-defined operating domain, validating the fundamental mechanistic assumptions while identifying where those assumptions potentially break down. Second, identify model gaps or structural limitations that cannot be resolved through parameter calibration alone, thereby defining the validated domain for subsequent calibration in Chapter 6 and 7. Third, demonstrate model maturity sufficient for predictive applications within the validated regime while clearly documenting boundaries beyond which model development is required.

5.2 SNAP-50 Experimental Context and Data

5.2.1 Historical Overview and Motivation

The SNAP-50 irradiation program was conducted in the 1960s by Pratt & Whitney Aircraft and advanced by researchers at Oak Ridge National Laboratory, primarily to qualify UN as a fuel for space reactor applications. The SNAP (Systems for Nuclear Auxiliary Power) program required fuel materials with high thermal conductivity, high uranium density, and acceptable swelling behavior under high-temperature, space-reactor neutron spectra. Uranium mononitride was identified as a promising candidate due to its thermal conductivity (approximately 0.25 W/cm·K at 1500 K, compared to 0.08 W/cm·K for UO₂) and uranium density (13.6 g/cm³ compared to 10.96 g/cm³ for UO₂). The SNAP-50 experiments were designed to characterize fuel performance across a range of reactor-relevant conditions and to understand the fundamental mechanisms governing fission gas release and swelling.

The DeCrescente et al. [13] report details initial UN fuel development and high-power irradiation experiments, while the Weaver et al. [46] Oak Ridge report provides comprehensive post-irradiation examination data spanning a broader range of irradiation conditions and includes detailed metallographic analysis and bubble characterization. Chapter 3 contains a complete discussion of the SNAP-50 program.

5.2.2 BISON Modeled Test Cases and Irradiation Conditions

BISON fuel performance simulations were developed for 39 individual SNAP-50 test cases spanning comprehensive operating conditions from published experimental data [13, 46]. The test matrix includes 13 specimens with multiple axial locations (top, middle, bottom) per

fuel pin, providing spatial resolution of temperature gradients and their effects on gas release and swelling.

5.2.2.1 Operating Envelope

Table 5.1 summarizes the operating conditions spanned by the validation dataset. The exceptional breadth of conditions, temperatures spanning 626 K, burnups varying by a factor of 25, and power densities varying by a factor of 7, enables rigorous testing of mechanistic model predictions across regimes where different physical phenomena dominate.

Table 5.1: SNAP-50 Operating Envelope and Measurement Ranges

Parameter	Range
Fuel centerline temperature	1213–1839 K
Burnup	0.18–4.58 %FIMA
Power density	0.32–2.12 kW/cm ³
Experimental FGR	0.01–36.9 %
Experimental swelling	1.0–10.0 % volumetric

The temperature range encompasses the full transition from athermal to diffusion-dominated behavior ($< 1300\text{K}$) through the thermal transition regime (1300–1600 K) to thermal diffusion-dominated behavior ($> 1600\text{ K}$), directly enabling validation of the three-regime diffusivity decomposition developed in Chapter 4. The burnup range spans both short-term exploratory tests and extended-duration mission profiles, while the power density variation enables assessment of fission-rate-dependent phenomena independent of temperature effects.

Experimental measurements of fission gas release span five orders of magnitude, from suppressed release ($< 0.01\%$) in low-temperature, low-burnup cases to substantial release ($> 30\%$) in high-temperature cases. Volumetric swelling ranges from 1% to 10%, capturing both stable swelling regimes and potential breakaway swelling behavior. This wide range of responses provides stringent testing of whether the mechanistic model correctly predicts both the magnitude and the temperature-burnup dependence of fuel behavior.

5.2.2.2 Fuel Fabrication and Physical Properties

UN fuel pellets were fabricated using powder-metallurgical synthesis at the Connecticut Advanced Nuclear Engineering Laboratory [13]. Uranium metal was hydrided, vacuum-decomposed to fine uranium powder, nitrated to sesquinitride (U_2N_3), then decomposed to stoichiometric UN. Pellets were cold-pressed and sintered at 2530 K under nitrogen, producing 4.78 mm diameter pellets at 95–98% theoretical density (14.3 g/cm³). Pre-irradiation characterization confirmed stoichiometric nitrogen content (5.56 wt%), low oxygen ($< 1000\text{ppm}$) and carbon ($< 1000\text{ppm}$), and grain sizes ranging 5–50 μm [13, 46].

All fuel pins were clad with PWC-11 (Nb-1% Zr-0.1% C) tubing with a tungsten diffusion barrier on the inner surface. Fuel-cladding gaps were filled with helium, and tungsten end plugs provided fission gas plenum volume. Cladding dimensions varied across specimens to accommodate different power densities and thermal conditions [46].

5.2.2.3 Experimental Measurements and Uncertainties

Post-irradiation examination provided the integral measurements compared to BISON predictions. Cladding temperature was monitored during irradiation via Type K thermocouples positioned at the fuel pin midplane with accuracy $\pm 10^\circ\text{C}$ and total uncertainty $\pm 15^\circ\text{C}$. Four thermocouples per capsule verified axial temperature uniformity [46]. Volumetric swelling was determined from fuel pellet diameter change (optical micrometry, ± 0.0002 in. resolution) and density change (immersion or geometric measurements). Combined uncertainties are $\pm 1\text{--}3\%$ depending on measurement technique [46]. Fission gas release was measured by puncturing pressurized pins, collecting gas in calibrated volumes, and analyzing Xe-131 and Kr-85 content via mass spectrometry. FGR is reported as percentage of total gas generated based on measured burnup. Measurement precision is $\pm 0.03\text{ cm}^3$ at STP, corresponding to $\pm 2\text{--}5\%$ relative error [46]. Burnup was determined from isotopic analysis (^{137}Cs , ^{144}Ce , or U-235 depletion) with $\pm 5\%$ uncertainty [46]. Transmission electron microscopy on select specimens provided direct observation of bubble populations and grain boundary features supporting the mechanistic model.

5.3 BISON Model Implementation

BISON fuel performance simulations of SNAP-50 test cases were developed to predict fuel swelling and fission gas release under experimentally-measured irradiation conditions. This section details the computational model architecture, novel multi-pellet mesh implementation, thermal boundary condition derivation from ABAQUS [68], and integration of mechanistic material property correlations developed in Chapter 2. ABAQUS is a commercial finite element analysis software used for simulating mechanical behavior, stress-strain analysis, and complex structural deformation across a wide range of materials and engineering applications.

5.3.1 Fuel Pin Geometry and Mesh Design

Two computational geometries were employed depending on the experimental configuration. Single-pin capsules utilized the standard BISON fuel pin mesh generator, producing a 2D-axisymmetric (RZ) geometry with independent fuel, cladding, and gap discretizations as shown in Figure 5.1. Three-pin stack capsules required a novel meshing architecture to enable simultaneous simulation of three fuel pellets with independent material response while maintaining shared cladding behavior.




Figure 5.1: Single fuel pin mesh utilizing BISON fuel pin mesh generator.

The three-pin mesh was constructed by stitching three instances of the fuel pin mesh generator such that the cladding formed a single continuous component encompassing all three pellets (top, middle, bottom), while the tungsten diffusion liner and fuel pellet for each location remained individual, independent mesh components. This decomposition enabled critical simulation features: independent power density assignment to each pellet based on experimentally-reported burnup values, independent calculation of postprocessor quantities (fission gas release and swelling) for each axial location, and shared cladding mechanical behavior ensuring proper fuel-cladding load transfer along the entire pin length (Figure 5.2). The result was efficient simultaneous simulation of three separate fuel pellet experiments within a single BISON execution, reducing computational overhead compared to three separate single-pin simulations while maintaining full physical coupling through the shared cladding component.



Figure 5.2: Three fuel pin configuration showing the variation in fuel temperature due to power influence and cladding boundary condition.

Mesh resolution was carefully selected to balance fidelity with computational cost. Typical employment included 18 radial elements in fuel (refined toward centerline for accurate capture of peak temperature and burnup), 50 axial elements per fuel pellet (enabling resolution of axial temperature gradients from the fitted polynomial boundary condition), 4 radial elements in cladding, and 2 radial elements in tungsten liner. Bilinear quadrilateral (Quad4) elements were used throughout for robust convergence and to ensure compatibility between stitched component interfaces.

5.3.2 Thermal Boundary Conditions

Establishing accurate thermal boundary conditions was critical for fidelity yet challenging due to the complexity of the experimental capsule geometry and multiple irradiation conditions. A novel approach was adopted: finite-element thermal analysis in ABAQUS, calibrated against experimental thermocouple data, was used to derive the complete cladding outer surface temperature distribution for each test case.

For each of the test cases, a dedicated ABAQUS thermal model replicated the fuel pin geometry, experimentally-reported power density, and fuel stack configuration (single or three-pellet). The model solved steady-state heat conduction with volumetric fission heat source distribution using material thermal conductivities appropriate to each case.

The critical calibration step involved iterative adjustment of the outer boundary condition (convective heat transfer coefficient to the lithium bath environment) until the calculated cladding centerline temperature matched the experimentally-reported nominal cladding temperature measured via thermocouples at the fuel pin midplane. This ensured that the ABAQUS model accurately captured the thermal resistance of the capsule and coolant environment.

Once calibrated to experimental thermocouple data, the complete axial temperature profile was extracted from the ABAQUS solution and fit to a high-order polynomial function of axial position:

$$T_{\text{clad,outer}}(z) = \sum_{i=0}^n a_i \cdot z^i \quad (5.1)$$

Polynomial orders between 11 and 13 captured the axial temperature variation, with higher orders necessary for three-pellet cases exhibiting multiple thermal peaks. Coefficient scaling factors near 1.0–1.2 improved numerical conditioning without altering the profile shape.

The fitted polynomial boundary condition was applied as a prescribed temperature on the cladding outer surface in BISON. This approach provided several advantages. The temperature distribution varied smoothly without discontinuities or artificial step-changes that could introduce numerical artifacts, the boundary condition matched experimentally-measured thermocouple values at midplane locations, spatial temperature variation reflected the actual neutron flux distribution (via differential power density input) and heat removal efficiency of the capsule design, and for multi-pellet cases, the polynomial captured the complex multi-peak structure, enabling each pellet to experience appropriately different thermal environments.

5.3.3 Thermal Model and Material Properties

In this work, the integration of mechanistic material property correlations, developed in Chapter 2, were directly implemented into the BISON material codes to replace existing or default constant property values. This ensured thermal and mechanical predictions were grounded in updated physically-consistent, temperature-dependent behavior.

Fuel thermal conductivity was calculated using UN correlations from Chapter 2 that account for both temperature dependence and porosity effects. The BISON material code was updated with these correlations as user-defined functions. This enabled the model to capture the significant variation in UN thermal conductivity across the 1200–1900 K temperature range spanned by SNAP-50 cases, directly impacting predicted fuel centerline temperatures and consequently the predicted gas release rates (which are highly temperature-sensitive).

Cladding thermal conductivity for PWC-11 (Nb-1% Zr-0.1% C) was similarly implemented with temperature dependence derived from literature data across the relevant range




Figure 5.3: Representative ABAQUS-fitted temperature profile showing characteristic peaks and valleys across the fuel pin axial length, with polynomial fit curve overlaid on ABAQUS data points to demonstrate fit quality. Include the three peaks for a three-pellet stack case to illustrate the methodology.

(300–1500 K). Historically, BISON implementations of niobium-alloy claddings employed constant thermal conductivity values, a simplification that becomes problematic at SNAP-50 cladding temperatures exceeding 1400 K. New equations were developed and coded into the material code to provide temperature-dependent conductivity, improving thermal accuracy at elevated temperatures relevant to advanced space reactor designs.

The fuel-cladding gap was modeled using BISON’s thermal contact mortar capability, implementing both thermal contact conductance and mechanical contact through mortar constraints. The thermal conductance was inherently gap-geometry dependent, accounting dynamically for gap closure due to fuel swelling and cladding thermal expansion over the irradiation duration. Mechanical contact enforcement prevented numerical instabilities from fuel-cladding separation and enabled accurate stress transfer between fuel and cladding, important for coupled mechanical-thermal predictions of fuel-cladding interaction.

5.3.4 Power History and Fission Rate Definition

Power history for each test case was determined by inverse calculation from experimentally measured quantities (burnup, irradiation time, and heavy metal density) rather than relying on potentially uncertain direct power measurements or design specifications. This inverse approach leveraged the experimentally determined burnup as the most reliable con-

straint on the cumulative fission history.

The average power density required to achieve the measured burnup was calculated by integrating the burnup equation:

$$\text{Burnup} = \frac{1}{\rho_u} \int_0^{t_{\text{irr}}} P(t) dt \quad (5.2)$$

where $P(t)$ is power density (W/cm^3), t_{irr} is irradiation time, and ρ_u is heavy metal (uranium) atomic density. However, fission rates are not uniform radially within the fuel pin due to thermal neutron self-shielding. The normalized radial power distribution was extracted from published neutron transport analysis and experimental data available in Weaver et al. [46], which provided a dimensionless radial flux depression profile showing the ratio of local-to-surface fission rate as a function of normalized fuel radius.

To ensure that the volume-averaged fission rate correctly reproduced the experimentally-measured average burnup, a volume integral of the radial power profile was performed:

$$P_{\text{avg}} = \frac{\int_0^1 P_{\text{nom}}(r) \cdot r dr}{\int_0^1 r dr} \quad (5.3)$$

where $P_{\text{nom}}(r)$ is the normalized radial power distribution and P_{avg} is the corrected average power density. The volume-averaged power depression factor (approximately 0.78 for the SNAP-50 capsule geometry) was applied to the nominal power density to yield the true volume-averaged fission rate. This correction ensured that despite the radial flux variations inherent to the SNAP-50 geometry, the spatial integral of fission rate over the fuel cross-section produced the experimentally-measured average burnup when integrated over irradiation time.

Incorporating the radial power profile was essential for mechanistic validation. The predicted spatial distribution of void swelling and fission gas bubble populations follows steep fission-rate dependencies, so accurate representation of the radial fission rate variation directly impacts whether the mechanistic model correctly predicts the observed spatial heterogeneity in fuel behavior. Conversely, constraining the volume-averaged fission rate to match experimental burnup ensured that model inputs were consistent with experimental observations.

For three-pellet capsules where individual pellet burnups were documented as varying (per the Weaver et al. [46] and DeCrescente et al. [13] reports), each pellet in the BISON simulation received an independent, tailored power profile such that its calculated average burnup matched the reported value. This capability was essential for accurate representation of spatially-varying irradiation damage and fission gas generation along the pin length.

Irradiation was modeled with a piecewise-linear power history consisting of three phases: ramp-up over 12 hours from zero to nominal power, steady-state operation at nominal power for the bulk of the irradiation duration, and ramp-down over 12 hours to zero at end-of-life. This profile reflects typical reactor operational cycles and avoids numerical singularities or artificial stress transients that can arise from instantaneous power step-changes.

5.3.5 Summary: Model Validation Strategy

The BISON model implementation prioritizes fidelity to experimental observations and consistency with mechanistic understanding. Thermal boundary conditions derive from dedicated ABAQUS calibration against each individual test case’s experimental thermocouple measurements, ensuring boundary conditions represent the actual thermal environment without generic assumptions. Power density and power history are back-calculated from experimentally-measured burnup rather than relying on potentially uncertain direct measurements, ensuring that predicted fuel burnup matches experimental data by construction. Radial neutron flux depression and radial temperature gradients are explicitly modeled, enabling validation of predicted spatial heterogeneity in fuel swelling and fission gas release. Material properties throughout the model employ correlations developed from fundamental mechanistic understanding in Chapter 2, ensuring that predictions are grounded in physically-consistent behavior.

This implementation approach ensures that BISON model inputs are tightly coupled to experimental data and mechanistic understanding, establishing a rigorous foundation for validation of mechanistic model predictions against integral fuel pin measurements of swelling and fission gas release.

5.4 Baseline Model Performance and Gas Partitioning Analysis

5.4.1 Collective Model-Experiment Comparison

Initial validation of the SIFGRS mechanistic model against the complete SNAP-50 experimental database reveals systematic prediction patterns that motivate targeted investigation of specific model parameters and physical mechanisms. Figure 5.4 presents model-experiment comparisons for both volumetric swelling (left panel) and fission gas release (right panel) across all 39 SNAP-50 test cases, with data points colored by fuel centerline temperature to reveal temperature-dependent trends. The swelling comparison (Figure 5.4, left panel) shows that the model captures temperature and burnup dependencies correctly but systematically overpredicts by 1–3 volume percent. This uniform bias suggests the mechanistic treatment of bubble growth may overestimate bubble populations or that gas partitioning requires adjustment.

In contrast, the fission gas release comparison (Figure 5.4, right panel) reveals more severe and temperature-dependent underprediction. At all temperatures and burnups, the model significantly underpredicts release. At moderate to high temperatures (1400 - 1800 K) experimental values reach 5–7% while predictions remain below 2%. At low temperatures ($T < 1400$ K, blue points), experimental fission gas release values are small (0.01–0.5%) but the model predicts essentially zero release, missing the onset of venting entirely.




Figure 5.4: Baseline SIFGRS model predictions (circle) compared to experimental measurements (square) for all SNAP-50 test cases. Left panel: Volumetric swelling versus burnup, Right panel: Fission gas release versus burnup. Data points are colored by fuel centerline temperature to reveal temperature-dependent prediction patterns. Green highlights under prediction for both swelling and fission gas release.

This temperature-dependent underprediction pattern suggests that the limiting mechanisms for fission gas release differ across temperature regimes and that the baseline model parameters may not adequately capture these transitions. At low temperatures, the model has a difficulty achieving grain boundary saturation, preventing onset of release, and limiting release all together. At intermediate temperatures, even when venting onset occurs, the model fails to sustain release rates commensurate with the observed high-burnup behavior, suggesting insufficient gas supply to grain boundaries or inadequate venting efficiency once interconnection is achieved.

5.4.2 Strategic Specimen Selection for Parametric Investigation

To systematically investigate the physical mechanisms responsible for the observed model-experiment discrepancies, five representative SNAP-50 specimens spanning distinct temperature and burnup regimes were selected for detailed parametric sensitivity analysis. These specimens were chosen to isolate specific physical phenomena (onset behavior, sustained release, high-temperature and high-burnup success) while minimizing confounding variables such as grain size variation or fabrication differences. Table 5.2 summarizes the

selected specimens and their experimental regimes.

Table 5.2: Specimen selection for parametric investigation including operating conditions.

Specimen	Temp (K)	Burnup (%FIMA)	Power (kW/cc)	FGR Exp (%)	FGR BISON (%)
220-M	1714	0.95	1.50	12.0	0.953
667-M	1498	0.78	0.40	0.02	0.029
602-M	1422	0.78	0.40	0.20	0.00
665-M	1426	4.58	1.00	5.96	0.061
669-M	1348	2.72	0.60	0.13	0.00
602-T	1269	0.97	0.47	0.02	0.00

Specimen 220-M represents the high-temperature regime where the baseline model achieves FGR, serving as a control case to ensure that parameter adjustments intended to improve intermediate- and low-temperature predictions do not degrade high-temperature, high-fgr accuracy. Specimens 602-M and 667-M probe onset behavior at intermediate temperatures with low power densities and modest burnups, isolating the grain boundary coverage accumulation rate from complications arising from high fission gas production rates. Specimen 665-M is the critical case for sustained release investigation, as it combines intermediate temperature with the highest burnup in the SNAP-50 database (4.58% FIMA), providing the most severe test of whether the model can sustain fission gas release over extended irradiation. Specimen 669-M and 602-T represents the low-temperature onset challenge, where even moderate burnup produces minimal experimental release but the model predicts essentially zero venting.

By focusing parametric investigations on this strategically selected subset, subsequent sections isolate specific physical mechanisms limiting model performance without the confounding effects of analyzing all 39 cases simultaneously. The selected specimens span temperatures from 1348 K to 1714 K, burnups from 0.78% to 4.58% FIMA, and power densities from 0.4 to 1.5 kW/cc, collectively representing the breadth of physical behavior observed in the SNAP-50 database.

5.4.3 Gas Partitioning Diagnosis: Baseline Model Limitations

To diagnose the physical origin of fission gas release underprediction, the spatial distribution of fission gas atoms throughout the fuel microstructure was examined at the completion of each irradiation cycle for the baseline model. Figure 5.5 presents stacked area plots for six representative specimens spanning the temperature range from 1348 K to 1714 K, showing the fraction of total produced gas residing in five distinct microstructural locations as a function of burnup: solid solution in the fuel matrix (blue), intragranular bubbles (orange), dislocation bubbles (purple), grain boundary bubbles (green), and released to the plenum (maroon).




Figure 5.5: Baseline model gas partitioning for six representative SNAP-50 specimens showing fraction of total produced fission gas in each microstructural location versus burnup. Colors indicate microstructural location for matrix (blue), intragranular bubbles (orange), dislocation bubbles (purple), grain boundary bubbles (green), and released gas (maroon).

The gas partitioning analysis reveals a universal pattern across all specimens. Early in the irradiation, fission gas atoms are rapidly trapped at dislocation networks (purple region), but as burnup progresses, this gas is released through dislocation bubble resolution and subsequently re-trapped in growing intragranular bubbles (orange region) rather than diffusing to grain boundaries. This dislocation to intragranular bubble re-trapping cycle dominates the gas transport pathway across the entire temperature range investigated.

For Specimen 665-M, which represents a difficult sustained release challenge, based on burnup and moderate temperature, the gas partitioning evolution is particularly instructive. At the start of irradiation, dislocation bubbles (purple) contain approximately 80% of total

gas, with the matrix (blue) holding approximately 20%. As burnup increases beyond 0.5% FIMA, the matrix bubbles lower as they diffuse through the matrix and trap at either intragranular bubbles or dislocations. Dislocation bubble fraction decreases as resolution releases gas back to the matrix, while the intragranular bubble fraction increases correspondingly, reaching approximately 85% by final burnup. The grain boundary gas inventory (green) remains below 3% throughout the irradiation, and the released fraction (maroon, barely visible) reaches only 0.06% despite the high burnup.

The high-temperature case, Specimen 220-M, exhibits qualitatively different behavior. Dislocation bubble trapping (purple) dominates early burnup ($\approx 100\%$ at 0.1% FIMA) but decreases as resolution begins to release bubbles from the dislocations and they trap in intragranular bubbles. Intragranular bubble trapping (orange) reaches a maximum of only 40% at intermediate burnup. The grain boundary gas fraction (green, visible but small) remains below 3%, and the released fraction (maroon) reaches approximately 1% by 0.95% FIMA. This demonstrates that at high temperatures, even modest grain boundary gas inventories are not sufficient for sustained release because gas diffusion to grain boundaries is still not rapid enough to maintain venting due to ongoing re-trapping and limited resolution.

The low-temperature specimens (669-M and 602-T) exhibit the most severe trapping behavior. For Specimen 669-M, the matrix fraction (blue) remains elevated (approximately 20%) throughout the irradiation, indicating that gas diffusion is slow enough that substantial fractions persist in solid solution rather than being captured by bubbles. However, the intragranular bubble fraction (orange) still dominates, reaching 75% by final burnup. The grain boundary gas inventory (green) barely registers (less than 1%), and correspondingly, release (maroon) is effectively zero. The persistence of substantial matrix gas at low temperatures might initially suggest that intragranular diffusion is the limitation, but comparison with the moderate-temperature cases reveals that this matrix inventory does not translate to grain boundary accumulation, indicating that re-trapping during diffusion prevents gas from completing the transport pathway.

The intermediate-temperature specimens (602-M and 667-M) exhibit behavior transitional between the high- and low-temperature extremes. Intragranular bubble trapping (orange) reaches 50–65% by final burnup, dislocation trapping (purple) contributes 20–35%, and grain boundary gas inventories (green) range from 1–3%. These specimens demonstrate that even at temperatures where gas diffusion is moderately fast, the re-trapping competition limits grain boundary accumulation and consequently suppresses fission gas release.

The grain boundary gas inventory remains below 3 percent across all specimens regardless of conditions, immediately identifying the limitation as inadequate gas supply rather than venting efficiency. Grain boundary coverage in these simulations reaches the saturation threshold (0.5 fractional coverage) required for bubble interconnection and venting for most cases, indicating that the venting model itself is functioning. However, because only 1–3% of total gas is available at grain boundaries, the absolute inventory of gas available for venting is insufficient to produce experimentally observed release rates, even when venting efficiency is near 100%.

The dominant physical process preventing gas from reaching grain boundaries is re-

trapping in intragranular bubbles following resolution of dislocation bubbles. Gas atoms execute a cyclic dislocation to resolution to intragranular bubble re-trapping pathway rather than completing the intragranular to grain boundary diffusion pathway required for sustained venting. This cyclic behavior persists across all temperatures and burnups, with only the relative importance of diffusion versus trapping kinetics varying with temperature. At high temperatures, rapid diffusion partially overcomes re-trapping, allowing modest release. At low temperatures, slow diffusion exacerbates re-trapping, preventing any measurable release. This diagnosis motivates systematic investigation of trapping rates, resolution rates, and matrix diffusion. However, adjustments must balance competing requirements across temperature regimes. Sufficient bubble population for grain boundary saturation, but weak enough trapping to maintain gas supply for sustained venting.

5.5 Systematic Parametric Investigation

The gas partitioning diagnosis established that the baseline model's fission gas release underprediction originates from excessive trapping of gas in intragranular and dislocation bubbles with restricted diffusion and limited resolution, preventing adequate gas supply to grain boundaries for venting. This section documents a systematic investigation of model parameters controlling gas transport and partitioning, progressing chronologically through the parameter adjustments attempted and the physical insights gained from each investigation. The parametric study focuses on the six strategic specimens identified in Table 5.2, with particular emphasis on Specimen 220-M, 665-M, and 669-M as the primary test cases for sustained release behavior.

5.5.1 Saturation Coverage Reduction

The first parameter investigated was the grain boundary saturation coverage threshold, which defines the fractional coverage at which grain boundary bubbles achieve geometric interconnection and enable venting to the fuel-cladding gap. The baseline model employs a saturation coverage of 0.5 (50% of grain boundary area covered by bubbles), derived from geometric percolation theory for randomly distributed circles on a plane.

Reducing the saturation coverage threshold to 0.25 lowers the grain boundary gas inventory required to initiate venting, potentially enabling earlier onset of fission gas release at low temperatures where coverage accumulation is slow. Figure 5.6 compares fission gas release predictions for the baseline saturation coverage (0.5) and reduced saturation coverage (0.25) across the five strategic specimens.

For Specimen 669-M (low temperature, 1348 K), reducing saturation coverage from 0.5 to 0.25 advances the onset of fission gas release from zero (baseline never achieves saturation) to approximately 1.5% FIMA. This demonstrates that at low temperatures, where intragranular diffusion rates are slow and grain boundary gas accumulation is correspondingly sluggish, the baseline saturation threshold is overly restrictive. However, the total release achieved




Figure 5.6: Effect of grain boundary saturation coverage reduction on fission gas release predictions for three specimen: ORNL-220 (red), ORNL-665 (purple), ORNL-669 (blue). Solid lines indicate the baseline saturation coverage (0.5) while dashed lines show reduced saturation coverage (0.25).

even with reduced saturation coverage is approximately 0.5%, overshooting the experimental value of 0.13%.

For Specimen 665-M (intermediate temperature, high burnup), the reduced saturation coverage produces earlier venting onset, increasing release from 1% to approximately 2% by final burnup. However, this improvement is modest relative to the experimental value of 5.96%, and the characteristic plateau behavior observed in the baseline simulation persists. After initial venting onset, the release rate diminishes despite continued burnup accumulation, indicating that the fundamental limitation is not the coverage threshold but rather the supply of gas to grain boundaries.

For Specimen 220-M (high temperature, 1714 K), saturation coverage has negligible effect on predictions, as the baseline model already achieves saturation early in the irradiation and sustains release throughout burnup accumulation. At high temperatures, gas diffusion to grain boundaries is sufficiently rapid that the saturation threshold does not constitute a bottleneck.

While onset behavior at low temperatures is sensitive to the coverage threshold, sustained release at high burnup is limited by gas supply rather than interconnection requirements.

Subsequent parametric investigations must therefore address the physical mechanisms governing gas transport from intragranular bubbles to grain boundaries.

5.5.2 Grain Boundary Vacancy Diffusion Enhancement

Having established that gas supply to grain boundaries limits sustained fission gas release, attention turned to the mechanisms governing gas atom diffusion and bubble growth at grain boundaries. One hypothesis for the observed release plateau was that vacancy diffusion at grain boundaries might be insufficient to accommodate the influx of gas atoms, leading to pressure buildup within grain boundary bubbles that suppresses further gas incorporation and arrests venting.

The grain boundary vacancy diffusion coefficient was increased by a factor of 10^3 relative to the baseline value to test whether enhanced vacancy mobility could sustain bubble growth and venting at higher gas supply rates. This extreme multiplier was chosen to ensure that vacancy diffusion would not constitute a bottleneck, allowing isolation of whether accommodation physics limits venting.




Figure 5.7: Effect of grain boundary vacancy diffusion enhancement on fission gas release predictions for three specimen: ORNL-220 (red), ORNL-665 (purple), ORNL-669 (blue). Solid lines indicate the baseline grain boundary vacancy diffusion (10^6) while dashed lines show reduced grain boundary vacancy diffusion (10^9).

Similar to saturation coverage, the vacancy diffusion enhancement shifts the onset of FGR earlier in burnup allowing some cases to achieve FGR under the same conditions. For Specimen 669-M and 665-M, release reaches approximately the same FGR as saturation coverage condition indicating that vacancy grain boundary diffusion no longer increase total FGR. This is seen when observing 220-M (high temperature), where the increase in vacancy grain boundary diffusion did not change the onset of FGR nor the total FGR release.

The vacancy diffusion investigation established definitively that accommodation physics is not the limiting mechanism in the baseline model though it does allow for enhanced grain boundary coverage, especially in lower temperature conditions. Indeed, the baseline model operates in a near-optimal accommodation regime where vacancy supply is adequate to support bubble growth but not so excessive as to suppress venting. The fission gas release limitation must therefore originate earlier in the transport pathway, specifically in the intragranular region where gas partitioning between bubbles, dislocations, and the matrix determines the supply rate to grain boundaries.

5.5.3 Dislocation Density Reduction

The preceding parametric investigations established that excessive intragranular bubble trapping limits fission gas release by sequestering gas atoms in stable bubble populations before they can diffuse to grain boundaries. Since dislocations provide the primary nucleation sites for the larger intragranular bubbles that dominate trapping, the dislocation density directly controls the total sink strength available for gas capture. The UN SIFGRS model acknowledges that dislocation density in irradiated UN exhibits order of magnitude scatter based on transmission electron microscopy measurements, ranging from 10^{12} to 10^{14} m^{-2} across different irradiation conditions and material fabrication routes [60]. The baseline simulations employ a constant dislocation density as a simplifying assumption, which may overestimate the sink strength relative to actual fuel microstructures.

Reducing the dislocation density decreases the total length of preferential nucleation sites available, which should reduce the rate at which gas atoms are captured into dislocation bubbles and increase the concentration of dissolved gas available for matrix diffusion and grain boundary transport. To test this hypothesis, the dislocation density was systematically reduced from the baseline value of $4 \times 10^{13} \text{ m}^{-2}$ to 10^{12} m^{-2} , spanning the experimental range observed in irradiated UN fuel.

Figure 5.8 presents the fission gas release predictions as a function of burnup for three dislocation density values across the temperature range. The results demonstrate that dislocation density reduction produces modest FGR improvements while still preventing low temperature cases (669-M) to initiate release. The only thing to note is that there is a slight increase in final FGR, suggesting that less gas is trapped in the dislocations and can migrate to the grain boundaries. However, the characteristic plateau behavior persists, suggesting that even with reduced dislocation trapping, gas transport remains insufficient to sustain continuous release.




Figure 5.8: Effect of dislocation density reduction on fission gas release predictions for three specimen: ORNL-220 (red), ORNL-665 (purple), ORNL-669 (blue-doesn't appear on plot). Solid lines indicate the baseline dislocation density ($4 \times 10^{13} \text{ m}^{-2}$) while dashed lines show reduced dislocation density ($1 \times 10^{12} \text{ m}^{-2}$).

The gas partitioning analysis shown in Figure 5.9 reveals the mechanistic explanation for the limited FGR improvement. As dislocation density decreases, the fractional gas inventory in dislocation bubbles (purple region) declines dramatically, confirming that fewer nucleation sites are available for bubble formation. However, this gas does not transition to the mobile matrix or grain boundary populations. Instead, the intragranular bubble fraction (orange region) increases equally, indicating that gas atoms not captured at dislocations are trapped in bulk bubbles nucleated homogeneously in the lattice.

This redistribution demonstrates that the model contains multiple competing sink populations with comparable trapping strengths. When dislocation sink strength is reduced, intragranular bubbles become the dominant trapping mechanism, capturing gas that would otherwise have been sequestered at dislocations. The matrix gas concentration (blue region) remains low across all dislocation density values, indicating that gas atoms are captured into bubbles, either dislocation or bulk, faster than they can diffuse to grain boundaries. The grain boundary gas fraction (green region) shows only modest increases at intermediate temperatures and remains negligible at low temperatures, confirming that reducing dislocation density alone cannot overcome the fundamental transport limitation.

The dislocation density investigation establishes that the baseline model's high disloca-




Figure 5.9: Gas partitioning evolution for dislocation density reduction across three temperature regimes. Baseline dislocation density (grey background) compared to reduced dislocation density (color overlay). Colors indicate microstructural location for matrix (blue), intragranular bubbles (orange), dislocation bubbles (purple), grain boundary bubbles (green), and released gas (maroon).

tion sink strength is not the sole cause of fission gas release underprediction. While dislocation trapping represents the dominant sink in the baseline configuration, the model's physics contains sufficient alternative trapping mechanisms (intragranular bubbles) that reduction in dislocation density merely shifts the locus of gas capture rather than enabling grain boundary transport.

5.5.4 Intragranular Diffusion Enhancement

The gas partitioning analysis thus far has established that the model traps approximately 85% of fission gas in intragranular or dislocation bubbles, with only 1–3% reaching grain boundaries. While this diagnosis identified gas supply as the limiting mechanism, an alternative hypothesis is that gas atoms successfully escape intragranular trapping but diffuse too slowly through the fuel matrix to reach grain boundaries before being re-trapped. If matrix diffusion constitutes the rate-limiting step, then enhancing the intragranular gas diffusion coefficient should enable gas atoms to complete the transport pathway from intragranular bubbles to grain boundaries more efficiently.

To test this hypothesis, the intragranular gas diffusion coefficient was uniformly multiplied by a factor of 5. This enhancement maintains the temperature-dependent functional form of the baseline diffusion model while scaling the magnitude for both Xe and vacancy diffusion, effectively accelerating the rate at which dissolved gas atoms and vacancies traverse the fuel matrix. The results displayed in Figure 5.10 demonstrate that enhancing intragranular diffusion by $5\times$ produces systematically modest but consistent improvements in gas transport to the grain boundary across the temperature range examined. The effect manifests through

several observable mechanisms that reveal the underlying gas transport dynamics within the fuel.




Figure 5.10: Effect of intragranular diffusion enhancement on fission gas release. Baseline intragranular diffusion coefficient (grey background) compared to enhanced diffusion with $5\times$ multiplier (color overlay). Colors indicate microstructural location for matrix (blue), intragranular bubbles (orange), dislocation bubbles (purple), grain boundary bubbles (green), and released gas (maroon).

In all three specimens shown, enhanced diffusion causes the dissolved matrix gas concentration (blue region) to decline, indicating that gas atoms diffuse away from the bulk matrix more rapidly before being trapped. Simultaneously, the dislocation bubble population (purple region) increases its fractional share of total gas, suggesting that faster diffusion allows gas to preferentially accumulate at the highest-affinity trapping sites—dislocations and intragranular bubbles—rather than migrating to the grain boundaries. This redistribution toward dislocation bubbles reflects a fundamental competition between diffusion pathways. When gas atoms can move more quickly, they preferentially seek out the strongest traps rather than being immobilized at random nucleation sites.

The magnitude of improvement in gas transport varies substantially across the temperature range examined. At the lowest temperature (669-M, 1348 K), the enhancement produces minimal fission gas release, and the majority of gas remains trapped in intragranular bubbles despite the accelerated diffusion.

The fivefold diffusion enhancement confirms that matrix diffusion is a secondary contributor, not the primary limitation. Enhanced diffusion enables slightly more efficient gas routing toward dislocations and modestly increases grain boundary supply, yet the persistent low-to-moderate grain boundary gas accumulation demonstrates that the primary bottleneck remains the excessive partitioning of gas into intra-granular bubbles during the early stages of irradiation. This finding redirects the investigation away from transport coefficients and toward more mechanisms which govern bubble trapping and resolution rates. The model's

failure to release significant fission gas is not due entirely due to slow diffusion to the grain boundary, but rather due to gas atoms being permanently captured in traps before they have the opportunity to reach the boundary at all. The subsequent investigation therefore focuses on directly manipulating the trapping and resolution rate parameters to test whether reducing gas capture in bubbles and enhancing bubble destruction can overcome the supply limitation.

5.5.5 Trapping and Resolution Rate Adjustment

The gas partitioning diagnosis and diffusion enhancement investigations established that excessive bubble trapping constitutes the primary limitation on fission gas release. The baseline model traps approximately 85% of generated fission gas in intragranular bubbles, with resolution (bubble destruction by fission fragments) insufficient to return this gas to the mobile matrix population. To test whether directly reducing trapping rates and enhancing resolution rates could overcome this bottleneck, a moderate parameter adjustment was implemented. Intragranular trapping rates (intragranular bubble trapping, dislocation bubble trapping, and dislocation line trapping) were reduced by a factor of 0.5, and resolution rates (intragranular bubble and dislocation bubble resolution) were increased by a factor of 5.

The rationale for this combined trapping-resolution adjustment is mechanistically straightforward. Reducing trapping rates decreases the probability that a dissolved gas atom will be captured by a bubble, maintaining a larger population of mobile gas atoms available for grain boundary transport. Enhancing resolution rates increases the frequency with which bubbles are destroyed by fission fragments, releasing trapped gas back to the fuel matrix before it can accumulate in large, stable bubbles. The moderate multipliers (0.5 for trapping, 5 for resolution) were selected to test whether physically plausible parameter adjustments, within the uncertainty bounds of the underlying mechanistic models, could produce meaningful improvements in gas transport.

Figure 5.11 reveals an unexpected and mechanistically revealing outcome. Across all three temperature regimes examined, the moderate trapping and resolution adjustment successfully reduces intragranular bubble gas fraction, confirming that the parameter modifications function as intended, but produces a counterintuitive redistribution of gas toward the fuel matrix rather than toward grain boundaries. For specimen 665-M (center panel, 1426 K), the matrix gas fraction (blue region) increases from approximately 10% (baseline) to 60–80% at BOL, while grain boundary gas (green region) remains below 10% throughout the irradiation. The sharp change in gas concentration as noted earlier for specimen 220-M is now more pronounced.

This gas redistribution pattern reveals a critical physical insight such that the model exhibits a two-stage bottleneck structure. The first bottleneck, intragranular bubble trapping, successfully captures most fission gas in the baseline model, preventing transport to grain boundaries. When this first bottleneck is partially relieved through reduced trapping and enhanced resolution, gas atoms successfully escape permanent capture in bubbles and




Figure 5.11: Gas partitioning for moderate trapping and resolution rate adjustments (trapping = $0.5\times$, resolution = $5\times$) across three temperature regimes. Baseline intragranular diffusion coefficient (grey background) compared to enhanced diffusion with $5\times$ multiplier (color overlay). Colors indicate microstructural location for matrix (blue), intragranular bubbles (orange), dislocation bubbles (purple), grain boundary bubbles (green), and released gas (maroon).

enter the mobile matrix population. However, a second bottleneck then becomes rate-limiting such that the matrix diffusion is insufficiently rapid to clear dissolved gas atoms to grain boundaries before continued fission gas production overwhelms the available transport capacity. Gas atoms are generated faster than they can diffuse through the fuel matrix to reach the grain boundary sink, resulting in accumulation of dissolved gas in solid solution.

The fission gas release consequences of this matrix accumulation are shown in Figure 5.12. With the dramatic shift in gas partitioning from bubbles to matrix, fission gas release improvements are sufficient but still well below experimental. For Specimen 220-M (high temperature), release increases from 1.5% (baseline) to approximately 6.5%. For Specimen 665-M (intermediate temperature), release increases from 1% to approximately 3%, still far below the experimental values. For Specimen 669-M (low temperature), release remains effectively zero despite 80% of gas residing in the matrix.

The moderate trapping and resolution investigation establishes that the baseline model's fission gas release limitation arises from a coupled bottleneck structure rather than a single rate-limiting mechanism. Intragranular bubble trapping constitutes the first bottleneck, preventing gas from entering the transport pathway. Matrix diffusion constitutes the second bottleneck, preventing gas that does escape trapping from reaching grain boundaries. Addressing only the trapping bottleneck, even with combined reduction in trapping rates and enhancement in resolution rates, is insufficient to achieve sustained fission gas release because gas accumulates in the matrix faster than diffusion can clear it to grain boundaries.

This finding motivates the aggressive combined parameter optimization presented in




Figure 5.12: Effect of trapping and resolution on fission gas release predictions for three specimen: ORNL-220 (red), ORNL-665 (purple), ORNL-669 (blue-doesn't appear on plot). Solid lines indicate the baseline trapping and resolution (Trap=1.0 Resolution=1.0) while dashed lines show reduced trapping and increase resolution (Trap=0.5 Resolution=5.0).

next, which simultaneously addresses both bottlenecks through coordinated enhancement of resolution rates, reduction of trapping rates, and acceleration of matrix diffusion.

5.5.6 Combined Parameter Investigation

The parametric investigations presented in Sections 5.5.1–5.5.5 established that the baseline model's fission gas release underprediction arises from a coupled two-stage bottleneck. Excessive intragranular bubble trapping (first bottleneck) followed by insufficient matrix diffusion to clear dissolved gas to grain boundaries (second bottleneck). The moderate trapping and resolution adjustment (Section 5.5.5) successfully demonstrated that reducing bubble trapping and enhancing bubble destruction shifts gas from intragranular bubbles to the fuel matrix, but revealed that matrix diffusion alone is insufficient to transport this dissolved gas to grain boundaries before continued production overwhelms the available diffusive capacity.

These findings motivate an aggressive three-parameter optimization strategy: simultaneously maximizing bubble resolution rates to continuously destroy traps and return gas to solution, minimizing trapping rates to prevent re-capture of dissolved gas, and maximizing

matrix diffusion rates to rapidly clear dissolved gas to grain boundaries before it can be re-trapped. The rationale for this combined approach is as follows. Minimized trapping (scaling factor = 0.1, a factor of 10 reduction from baseline) dramatically slows gas capture by bubbles, maintaining the majority of fission gas in mobile solid solution. Maximized resolution (scaling factor = 10, a factor of 10 increase from baseline) ensures that bubbles that do nucleate are destroyed rapidly before they can grow and stabilize, maintaining a dynamic steady state where bubble populations remain transient rather than permanent trapping sites. Maximized intragranular diffusion (multiplier = 10, enhancing both vacancy and xenon diffusion coefficients) accelerates the rate at which dissolved gas traverses the fuel matrix to reach grain boundaries, overcoming the matrix diffusion bottleneck identified in Section 5.5.5.




Figure 5.13: Gas partitioning comparison for aggressive combined parameter optimization (trapping = 0.1, resolution = 10, diffusivity = 10) across three temperature regimes. Left: Specimen 220-M (1714 K). Center: Specimen 665-M (1426 K). Right: Specimen 669-M (1348 K).

Figure 5.13 demonstrates the dramatic effect of the combined parameter optimization on gas partitioning. Across all three temperature regimes, the aggressive parameter combination produces qualitatively similar gas redistribution patterns, though with temperature-dependent magnitudes. For Specimen 665-M (center panel, 1426 K), the primary test case, intragranular bubble gas decreases from 85% (baseline) to <10%, matrix gas remains controlled at 10–20% (compared to 60–80% in the moderate case, Figure 5.11), grain boundary gas increases from 1% to 15%, and released gas reaches 6%.

The success of the combined optimization arises from synergistic effects that address both identified bottlenecks simultaneously. Minimized trapping (scaling factor = 0.1) prevents the majority of dissolved gas from being captured in bubbles, but unlike the moderate case (Section 5.5.5), does not cause matrix accumulation because enhanced diffusion provides sufficient transport capacity. Enhanced resolution (scaling factor = 10) increases the rate at

which bubbles that do form are destroyed by fission fragments, limiting the time available for gas accumulation within bubbles and maintaining a dynamic equilibrium where bubble gas content remains low. Enhanced intragranular diffusivity (multiplier = 10) overcomes the matrix diffusion bottleneck observed in the moderate case, clearing dissolved gas to grain boundaries faster than production can overwhelm transport capacity. The critical insight is that all three parameters must be adjusted aggressively and simultaneously with minimal trapping to prevent capture, maximal resolution to destroy bubbles that do form, and maximal diffusion to clear matrix gas before re-trapping.




Figure 5.14: Effect of trapping, resolution and diffusivity on fission gas release predictions for three specimen: ORNL-220 (red), ORNL-665 (purple), ORNL-669 (blue-doesn't appear on plot). Solid lines indicate the baseline trapping, resolution and diffusion (Trap=1.0 Resolution=1.0 Diff=1.0) while dashed lines show reduced trapping, increased resolution and diffusion (Trap=0.1 Resolution=10.0 Diff=10.0).

The cross-specimen comparison (Figure 5.14) reveals both the success and the limitations of the combined parameter optimization. The combined optimization produces moderate overprediction, as the baseline model already performed adequately at high temperatures where gas diffusion is rapid. The combined parameter set enhances gas transport beyond what is necessary, resulting in excessive release. For Specimen 665-M, the sustained release behavior qualitatively matches experimental trends, demonstrating that the physical mechanisms governing gas transport have been correctly identified even though the quantitative parameter values require refinement.

For Specimen 669-M (low temperature), the combined optimization produces sufficient release. The baseline model was fundamentally unable to release gas due to the low temperature migration of gas. The combined optimization demonstrates that the physical pathway exists for significant fission gas release when trapping is minimized, resolution is maximized, and diffusion is enhanced, but optimal parameter values must be temperature-dependent to match experimental observations across the full operating range.

The combined parameter optimization investigation establishes that the fission gas release limitation at all temperatures can be overcome through synergistic adjustment of trapping, resolution, and diffusivity parameters, confirming that the baseline model contains the necessary physical mechanisms but employs parameter values that excessively limit gas transport. However, the temperature-dependent response to these parameter adjustments, with systematic overprediction increasing in severity from high to low temperatures, indicates that optimal parameter values are temperature-dependent, and that a single parameter set cannot simultaneously match experimental fission gas release across the full temperature range. Before proceeding to regime-based analysis and mechanistic interpretation, it is essential to evaluate the complete suite of parameter variations investigated in Sections 5.5.1–5.5.6 against both primary validation metrics, fission gas release and fuel swelling, to assess whether multi-objective optimization is achievable and to identify the fundamental limitations constraining model predictive capability.

5.5.7 Multi-Objective Validation

Throughout Sections 5.5.1–5.5.6, parameter adjustments were evaluated primarily through fission gas release predictions. However, mechanistic fuel performance models must simultaneously satisfy multiple validation objectives. Fuel swelling constitutes a critical second metric that constrains fuel-cladding mechanical interaction and licensing margins. Because parameter adjustments alter gas partitioning between microstructural locations, they directly affect both fission gas release and swelling. This section provides a comprehensive multi-objective assessment.

Table 5.3: Fission gas release predictions across parameter variations.

Case	Exp FGR (%)	Sat Cov. FGR (%)	Vac. GB FGR (%)	IG Diff. FGR (%)	Trap/Res FGR (%)	Combined FGR (%)
220-M	12.0	1.7	1.04	1.14	3.85	9.98
665-M	5.96	0.92	1.49	0.11	0.77	5.97
669-M	0.13	0.15	1.20	0.00	0.00	1.46
602-B	0.08	0.09	0.20	0.00	0.00	1.88
667-M	0.03	1.01	0.19	0.11	1.67	6.17

Tables 5.3 and 5.4 reveal asymmetric model performance across the two validation metrics. Volumetric swelling remains robust across the parameter space explored, showing

Table 5.4: Volumetric swelling predictions across parameter variations.

Case	Exp Swell (%)	Sat Cov. Swell (%)	Vac. GB Swell (%)	IG Diff. Swell (%)	Trap/Res Swell (%)	Combined Swell (%)
220-M	6.0	6.1	6.9	9.2	4.1	5.1
665-M	7.9	12.6	15.1	16.5	8.8	8.1
669-M	6.9	7.7	9.2	10.1	4.8	4.7
602-B	2.4	3.1	3.8	4.2	1.7	2.1
667-M	1.8	3.3	3.9	5.1	2.4	2.7

that as swelling scales directly with dislocation and intragranular bubble inventory as seen with an increase in the first three tests. As the resolution and trapping was reduced, the swelling was reduced towards experimental while FGR was increased.




Figure 5.15: Predicted versus experimental scatter plots for volumetric swelling (left) and fission gas release (right) across all parameter variations. Shaded regions indicate factor-of-2 uncertainty bands. Swelling clustering near 1:1 agreement line demonstrates robustness to parameter adjustments.

These results establish that fuel swelling prediction capability is fundamentally coupled to fission gas release performance. More interestingly, fission gas release and swelling behavior with respect to parameter manipulation appears to be experimental regime-dependent. Optimal trapping, resolution, and diffusivity values vary systematically with burnup and temperature, and no single temperature-independent parameter set can satisfy validation objectives across all operating regimes simultaneously. For this reason, an investigation into how the tests are grouped based on temperature, burnup and power must be conducted.

5.6 Regime Identification and Model Limitation Analysis

The SNAP-50 experiments separate naturally into distinct operating regimes when mapped in temperature–burnup space. Figure 5.16 summarizes the experimental data and baseline model predictions across these regimes. Panel (a) identifies three characteristic clusters in temperature–burnup space and colors each specimen by experimental FGR. The bulk of the data falls within a compact region between 1200 and 1585 K and below roughly 3 %FIMA. In this standard regime experimental FGR remains very low, typically between 0.01 and 0.15 %. Two specimens lie outside this cluster. Specimen 220 occupies the high-temperature corner above 1600 K, where experimental FGR reaches 12 %. Specimen 665 occupies the high-burnup corner near 4.6 %FIMA at about 1425 K, where experimental FGR is almost 6 %. These outliers demonstrate that either very high temperature or very high burnup can activate additional release mechanisms that are not strongly expressed in the bulk of the dataset.

Panel (b) shows model error, which is small for most specimens but grows near the upper edge of the temperature range. Panel (c) compares predicted and experimental FGR on a parity plot. Standard-regime points cluster near the one-to-one line, while the high-temperature and high-burnup specimens deviate strongly. Panel (d) presents the average microstructural gas distribution for each regime and reveals that the outliers already exhibit elevated grain boundary inventories relative to the standard regime. This redistribution hints that the model begins to push gas to grain boundaries in conditions where experiments still show limited release.

To understand the origin of this behavior, the analysis focuses on the subset of specimens that remain within the main temperature–burnup cluster. Removing the two extreme outliers allows closer examination of how model error evolves as temperature and burnup increase. Figure 5.17 shows this in more depth. Panel (a) shows fission gas release error as a function of temperature with marker size proportional to burnup. Within the standard regime, predictions remain close to experimental values until a sharp transition near 1485 K and 1 %FIMA. Specimens below either threshold show errors near zero. Four specimens that exceed both thresholds exhibit a sudden increase in error, with model predictions between 0.6 and 1.0 % FGR while experiments remain near 0.1 %. This discontinuity does not appear in the experimental data, which vary smoothly across the same conditions, so it must arise from model structure rather than from a physical transition in the fuel.

The experimental FGR map in panel (b) reinforces this conclusion. Points within the shaded high temperature–high burnup region continue to show low experimental FGR. In contrast, the model predicts a step increase in release once both thresholds are exceeded as shown in panel a. The diagnostic plot of grain boundary coverage in panel (c) identifies the specific mechanism. For most specimens, grain boundary coverage increases gradually with temperature and burnup and remains between about 20 and 40 %. In the four regime shifts, coverage reaches 50 %. At this point the FGR is triggered.




Figure 5.16: Regime analysis of SNAP-50 experimental data and baseline model performance. (a) Temperature–burnup map colored by experimental FGR showing standard, high-temperature, and high-burnup regimes. (b) Model error map showing where predictions deviate from measurements. (c) Parity plot comparing predicted and experimental FGR. (d) Average gas distribution by regime showing the shift in microstructural partitioning.

Panel (d) shows the corresponding change in gas partitioning. In the standard regime, an average of roughly 10 % of the gas remains in the matrix, about half resides in intragranular bubbles, and the remainder is divided between dislocation loops and very small grain boundary and released fractions. In the high temperature–high burnup specimens, the model shifts gas from dislocation bubbles to grain boundaries and the released category once the percolation threshold activates. Experimental FGR, however, does not reflect such a transition. The data therefore indicate that real grain boundaries do not suddenly become fully connected at 50 % coverage and that additional resistance to venting remains, possibly due to microstructural changes due to temperature or burnup.

Swelling diagnostics lead to the same conclusion. Figure 5.18 shows swelling error versus temperature, again with marker size proportional to burnup. Errors remain modest through most of the domain, then increase sharply in the same shaded region defined by 1485 K and




Figure 5.17: Fission gas release diagnostics for the standard regime. (a) Prediction error versus temperature with marker size proportional to burnup (b) Experimental FGR map showing smooth variation across the same conditions (c) Grain boundary coverage showing the binary percolation switch at 50 % (d) Gas partitioning shift between standard and high temperature-high burnup conditions.

1 %FIMA. The experimental swelling distribution covers a wide range, from roughly 2 to 8 % volumetric strain, but does not show a sudden step at the identified thresholds. Instead, the enhanced swelling predicted in the high temperature-high burnup region is an artifact of the same binary percolation switch that forces gas to the grain boundaries and out of the fuel. Once this switch activates, the model both releases too much gas and retains an unrealistically altered internal gas inventory, which in turn perturbs swelling predictions.

These diagnostics motivate a regime-based definition of the calibration domain. The




Figure 5.18: Swelling diagnostics for the standard regime. (a) Prediction error versus temperature with marker size proportional to burnup. (b) Experimental swelling distribution across the full temperature to burnup domain.

standard regime includes all specimens that remain below 1485 K and below 1 %FIMA. Within this subset the model response is continuous. Both FGR and swelling errors are small and do not show abrupt changes with operating conditions. This regime spans a sufficient combinations of temperature, burnup, and power density to constrain the sensitive parameters identified in the preceding parametric studies. Specimens that exceed both thresholds are excluded from parameter calibration because their discrepancies originate from the structural grain boundary percolation assumption rather than from uncertain rate parameters. Attempting to fit them would drive trapping, resolution, and diffusion coefficients away from physically plausible values.

The external regimes retain important roles in the validation process. Specimen 220 defines a high temperature thermal release regime where additional mechanisms such as enhanced grain boundary sweeping, restructuring, or high-temperature coarsening likely contribute to the observed 12 % release. Specimen 665 defines a high burnup regime where cumulative damage and high fission gas production enable significant release even at moderate temperature. Both conditions remain as dedicated validation cases that bracket the standard regime and highlight where new physics are required before extrapolation can be trusted.

5.7 Summary and Transition to Sensitivity Analysis

This chapter established the initial validation of the SIFGRS mechanistic model for UN fuel by comparing BISON predictions to the SNAP-50 experimental database. The work demonstrated that the model can reproduce key features of fuel behavior within a defined operating domain and also exposed structural limitations that appear when both temperature and burnup become large.

The early sections described the experimental program, model implementation, and baseline performance. The BISON simulations incorporated detailed thermal boundary conditions derived from ABAQUS analyses, temperature-dependent material properties, and power histories that reproduce measured burnup. With this framework in place, baseline comparisons showed that volumetric swelling is predicted with good fidelity throughout the dataset. Swelling errors fall within experimental uncertainty for most specimens, which indicates that the overall magnitude and burnup dependence of retained gas and vacancy populations are represented reasonably well in the current formulation.

Fission gas release proved more challenging. Even after model inputs were tied closely to experimental conditions, baseline simulations significantly underpredicted FGR across much of the temperature range. Gas partitioning diagnostics revealed that the model traps roughly 85 % of produced gas in intragranular and dislocation bubbles, while only a small fraction reaches grain boundaries. Systematic parametric studies then investigated the roles of saturation coverage, grain boundary vacancy diffusion, dislocation density, intragranular diffusion, and coupled trapping-resolution rates. These studies showed that adjustments to any single mechanism produce only modest improvements and often trade accuracy between temperature regimes. The investigations also revealed a coupled two-stage bottleneck. Strong trapping at dislocations and intragranular bubbles forms the first bottleneck. Limited matrix diffusion into grain boundaries forms the second. When trapping and resolution are adjusted without enhancing diffusion, gas simply accumulates in the matrix. When diffusion is increased without reducing trapping, gas is routed preferentially to strong sinks instead of to grain boundaries.

A combined parameter study demonstrated that the model physics are capable in principle of matching high FGR values when trapping is strongly reduced, resolution is enhanced, and diffusion is increased together. Under these conditions the gas inventory shifts from intragranular bubbles to grain boundaries and the plenum, and the model reproduces both the magnitude and qualitative shape of release for the most demanding specimens. This success comes at the cost of overprediction in other parts of the dataset. The response therefore cannot be captured by a single temperature-independent parameter set. Instead, the model behaves as a family of regimes, each with different effective balances among trapping, resolution, and diffusion.

The regime analysis formalized this observation. Using temperature-burnup maps and focused diagnostics, the chapter identified a standard regime where the model response is smooth and both swelling and FGR errors remain small, along with high-temperature and high-burnup regimes where predictions deviate systematically. The binary grain

boundary percolation rule in the BISON venting model was shown to introduce an artificial discontinuity at 1485 K and 1 %FIMA. Above this threshold, the model forces grain boundary coverage to act as a hard switch and produces unrealistically high release and swelling. Experimental data do not exhibit such a sharp transition, which confirms that the discrepancy is structural rather than parametric. The calibration domain was therefore restricted to the 22 specimens that remain below the 1485 K and 1 %FIMA thresholds. Within this domain the model behavior is continuous and amenable to parameter refinement.

This regime-guided validation provides the foundation for the next two chapters. Chapter 6 applies a comprehensive sensitivity analysis to the standard regime dataset. The analysis ranks model parameters according to their influence on swelling and fission gas release and quantifies how each mechanism contributes within the validated operating space. Particular attention is given to parameters that govern intragranular diffusion, bubble trapping and resolution, and grain boundary transport, since these processes emerged in this chapter as the controlling features of the gas migration pathway. The sensitivity results define a reduced set of influential parameters that will be carried forward into the calibration stage.

Chapter 7 then performs Bayesian calibration and uncertainty quantification using the standard-regime specimens as the likelihood-defining data. The regime definition developed here ensures that posterior distributions reflect genuine parameter uncertainty rather than compensating for missing physics or structural discontinuities in the model. Calibration is carried out while maintaining consistency with the mechanistic interpretation of the traps-diffusion-venting pathway. The high-temperature and high-burnup specimens, including 220 and 665, are reserved as out-of-sample validation cases that test whether the calibrated model extrapolates in a physically reasonable manner and highlight where additional mechanistic development is still required.

In summary, this chapter moved the SIFGRS model from implementation to a regime-aware validated state. The work identified the domain where the current physics are credible, isolated the structural limitations associated with the grain boundary percolation assumption, and established a clear path for subsequent sensitivity analysis and Bayesian uncertainty quantification. Together, these steps position the model for calibrated predictions in advanced reactor conditions while transparently documenting the boundaries beyond which further model development is necessary.

Chapter 6

Global Sensitivity Analysis

The parametric investigation in Chapter 5 established that SIFGRS predictions of fission gas release are highly sensitive to trapping, resolution, and diffusion processes, while swelling predictions remain comparatively robust. That work relied on targeted parameter sweeps in a reduced space. The goal of this chapter is to extend those insights to the full 22-parameter model using a formal global sensitivity analysis. The resulting ranking of influential mechanisms provides the bridge between the regime-based validation in Chapter 5 and the Bayesian calibration and uncertainty quantification developed in Chapter 7.

The parametric investigation demonstrated that fission gas release predictions are sensitive to trapping rates, resolution rates, intragranular diffusion coefficients, and grain boundary saturation coverage, while swelling predictions remain robust across parameter variations. However, systematic quantification of parameter importance across the full 22-dimensional parameter space is required to prioritize parameters for Bayesian calibration and identify potential compensatory interactions that were not previously investigated.

This chapter presents a global sensitivity analysis using the Morris Elementary Effects method to quantify parameter importance for both fission gas release and volumetric swelling predictions across the SNAP-50 experimental database. The analysis identifies the subset of influential parameters requiring calibration, confirms regime-dependent parameter sensitivity suggested by Chapter 5 findings, and quantifies interaction effects that inform the Bayesian calibration framework (Chapter 7).

6.1 Morris Elementary Effects Methodology

The Morris method [69, 70] quantifies parameter importance through two metrics, the μ^* (mean absolute elementary effect), which measures overall parameter influence on model outputs, and σ (standard deviation of elementary effects), which measures nonlinearity and parameter interactions. The method uses randomized one-at-a-time parameter perturbations, providing computational efficiency suitable for expensive simulation models while maintaining qualitative agreement with variance-based methods [71].

Each of the 22 uncertain SIFGRS parameters was perturbed over predetermined ranges, representing plausible parametric uncertainty based on literature values and physical constraints. The Morris analysis employed 168 randomized trajectories through 12 levels, providing robust estimates of μ^* and σ for each parameter. Sensitivity metrics were computed separately for fission gas release and volumetric swelling at end-of-life conditions for representative specimen-locations, enabling assessment of whether parameter importance varies with operating conditions (temperature, burnup, power density).

6.2 Fission Gas Release Sensitivity

Figure 6.1 presents Morris sensitivity metrics for fission gas release predictions, with results colored by temperature (left), power density (center), and burnup (right) to assess regime-dependent sensitivity.

Eight parameters exhibit substantial influence on fission gas release ($\mu^* > 2$), ranked by mean μ^* across all specimens (Table 6.1):

1. **Grain Radius** ($\mu^* = 11.0$, $\sigma = 17.5$): Dominant influence consistent with grain size controlling diffusion length scales and grain boundary area. High σ indicates strong nonlinear effects and interactions with transport parameters.
2. **Saturation Coverage** ($\mu^* = 7.3$, $\sigma = 13.4$): Controls grain boundary venting onset, consistent with Chapter 5 finding that coverage threshold affects venting activation. Moderate σ reflects interaction with gas supply mechanisms.
3. **Dislocation Trap Z** ($\mu^* = 6.6$, $\sigma = 12.4$): Controls dislocation line trapping strength, confirming Chapter 5 identification of intragranular trapping as first bottleneck in two-stage gas transport limitation.
4. **Effective Diffusion Coefficient Scale Factor** ($\mu^* = 9.8$, $\sigma = 9.1$): Scales intragranular gas diffusion, confirming matrix diffusion as second bottleneck (Section 5.5.5).
5. **Dislocation Density** ($\mu^* = 6.1$, $\sigma = 9.7$): Controls trapping site density and bubble nucleation, with sensitivity confirming dislocation trapping mechanism importance.
6. **Resolution Scale Factor** ($\mu^* = 4.0$, $\sigma = 6.6$): Scales bubble resolution rates. Chapter 5 demonstrated that $10\times$ resolution enhancement is required to overcome trapping bottleneck, validating strong sensitivity.
7. **Trapping Scale Factor** ($\mu^* = 2.7$, $\sigma = 6.3$): Scales intragranular bubble trapping rates. Strong sensitivity validates Chapter 5 finding that $0.1\times$ trapping reduction is required for sustained release.

8. **Intra-granular Diffusion Coefficient** ($\mu^* = 2.5$, $\sigma = 6.2$): Scales Xe and vacancy diffusion rates. Strong sensitivity validates Chapter 5 finding that $10\times$ diffusion increase is required for sustained release.
9. **Vacancy GB Diffusivity Multiplier** ($\mu^* = 2.31$, $\sigma = 5.6$): Scales grain boundary vacancy diffusion. Moderate sensitivity consistent with Section 5.5.2 finding that enhanced GB diffusion improves gas coverage accumulation.

Figure 6.1 presents parameter importance ranked by μ^* . The top eight parameters maintain consistent ranking across all operating conditions, confirming that parameter importance is regime-independent despite regime-dependent model response identified in Chapter 5. This consistency shows that the same physical mechanisms, trapping, diffusion, and grain boundary saturation, control fission gas release in each regime. What differs between regimes is not which parameters matter but which values of those parameters best describe the data. That distinction is central to the Bayesian calibration strategy in Chapter 7.




Figure 6.1: Fission gas release parameter importance ranked by μ^* . Results colored by temperature (left), power density (center), and burnup (right).

Comparing sensitivity results with Chapter 5 parametric findings reveals excellent consistency:

- **Trapping-Diffusion Coupling:** Dislocation Trap Z and Effective Diffusion Coefficient ($\mu^* = 9.8$) both exhibit strong influence, validating the two-stage bottleneck mechanism where both parameters must be adjusted simultaneously (Section 5.5.6).

- **Resolution Importance:** Resolution Scale Factor ranks among top parameters, confirming Chapter 5 finding that resolution enhancement is critical for sustained release.
- **Saturation Coverage:** Saturation Coverage exhibits high sensitivity despite Chapter 5 finding that coverage adjustment alone produced minimal improvement. This apparent contradiction reflects that coverage threshold strongly influences onset timing (high sensitivity) but cannot overcome gas supply limitation (low parametric improvement).
- **Grain Radius Dominance:** Grain radius exhibits highest sensitivity, validating its role in controlling diffusion length scales. However, grain radius is microstructurally constrained (measured $7.5 \mu m$ in SNAP-50 fuel [13]), not a calibration parameter.

Table 6.1: Summary of top parameters for fission gas release based on mean and maximum values of μ^* and σ

Parameter Display	Mean μ^*	Max μ^*	Mean σ	Max σ
Grain Radius	11.0	24.2	17.5	30.9
Sat. Coverage	7.3	10.6	13.4	16.9
Dis. Trap Z-factor	6.6	12.2	12.4	18.7
Dis. Density	4.0	8.0	9.7	16.5
Resolution Scalar	3.6	5.1	6.6	8.1
Trapping Scalar	2.7	4.0	6.3	8.5
Intra. Diff. Coeff. Scalar	2.5	3.9	6.2	8.6
Vac. GB Diff. Multiplier	2.3	4.0	5.6	9.6
Surface Energy	1.8	3.6	4.3	8.7
Fractional Yield	1.5	2.6	3.7	5.9
Nucleation Rate Scalar	1.2	1.8	3.3	4.4
Trapping SScalar (Dis. Line)	0.8	1.9	3.6	7.1
Resolution Scalar (Dis.)	0.8	2.1	3.6	6.7
Trapping Scalar (Dis.)	0.2	0.6	0.9	2.0

6.3 Volumetric Swelling Sensitivity

Figure 6.2 presents Morris sensitivity metrics for volumetric swelling predictions. The sensitivity validates the multi-objective validation finding that swelling predictions remain accurate across all parameter variations tested, including aggressive combined optimization that produced $40\times$ fission gas release overcorrection. In other words, the model can distort the release pathway dramatically before volumetric swelling moves outside experimental uncertainty. Lower parameter sensitivity for swelling indicates that local fission product

accommodation (bubble size distributions, intragranular porosity) is less sensitive to transport parameter uncertainties than long-range gas transport to grain boundaries.

Parameter ranking as seen in Table 6.2 for swelling closely parallels fission gas release ranking, with trapping, resolution and diffusion occupying the top positions for both outputs. Figure 6.2 confirms consistent parameter importance across operating conditions, with no systematic variation in sensitivity with temperature, power density, or burnup.

Table 6.2: Summary of top parameters for swelling based on mean and maximum values of μ^* and σ

Parameter Display	Mean μ^*	Max μ^*	Mean σ	Max σ
Grain Radius	6.3	15.1	10.7	25.8
Sat. Coverage	4.2	11.7	6.9	18.0
Dis. Trap Z-factor	3.3	5.5	7.2	12.1
Dis. Density	3.0	5.1	7.9	14.9
Fractional Yield	2.8	3.9	3.1	4.6
Intra. Diff. Coeff. Scalar	1.7	2.3	2.9	4.4
Vac. GB Diff. Multiplier	1.4	2.6	2.9	5.4
Surface Energy	1.4	2.3	2.5	3.9
Resolution Scalar	1.1	1.8	1.9	3.1
Trapping Scalar	1.0	1.4	1.6	2.5
Resolution Scalar (Dis.)	0.8	1.5	2.6	5.6
Trapping SScalar (Dis. Line)	0.6	1.0	2.0	4.0
Nucleation Rate Scalar	0.4	0.7	0.9	1.8
Trapping Scalar (Dis.)	0.1	0.2	0.4	0.7

6.4 Parameter Selection for Bayesian Calibration

Based on sensitivity analysis results and physical considerations from Chapter 5, 14 parameters were selected for a full Bayesian analysis (seen in Tables 6.1 and 6.2).

These 14 parameters correspond directly to the mechanistic bottlenecks identified in Chapter 5. Trapping strength and site density (Dislocation Trap Z, Trapping Scale Factor, Dislocation Density) govern the first stage of the transport pathway, where gas leaves the matrix and is captured in intragranular features. Intragranular diffusion coefficients control migration through the matrix. Grain boundary vacancy diffusivity sets the rate at which gas accumulates along boundaries once it reaches the grain edge. Resolution Scale Factor allows trapped gas to re-enter the mobile population, while Saturation Coverage regulates the onset of interconnected grain boundary networks and venting. This parameter subset enables regime-specific calibration (Chapter 7) where optimal values vary between the three identified regimes (low-burnup, high-temperature, and high-burnup) while maintaining physical mechanism consistency.




Figure 6.2: Volumetric swelling parameter importance ranked by μ^* . Results colored by temperature (left), power density (center), and burnup (right).

6.5 Summary

Global sensitivity analysis using the Morris Elementary Effects method quantified parameter importance for fission gas release and volumetric swelling predictions across the SNAP-50 experimental database. In the context of the overall dissertation, this chapter serves as the quantitative link between the diagnostic model evaluation in Chapter 5 and the Bayesian calibration and uncertainty quantification in Chapter 7.

Key findings from the global sensitivity analysis can be summarized as follows. The same 14 parameters dominate fission gas release sensitivity across the entire operating space, which spans temperatures from 1200 to 1750 K, burnups from 0.3 to 4.6% FIMA, and power densities from 0.3 to 1.5 kW/cc, confirming that the importance of the underlying mechanisms is regime-independent despite regime-dependent optimal parameter values identified in Chapter 5. The high sensitivity of trapping parameters, diffusion coefficients, and the resolution validates the two-stage bottleneck mechanism and confirms that simultaneous adjustment of multiple parameters is required for sustained release. Volumetric swelling exhibits a compressed sensitivity structure relative to fission gas release, with $\mu_{\max}^* \approx 15$ compared to $\mu_{\max}^* \approx 24$ for FGR, demonstrating that local accommodation is less sensitive than long-range transport. Several parameters exhibit $\sigma \approx \mu^*$, indicating strong nonlinear effects and parameter interactions, which motivates the Bayesian calibration framework in Chapter 7 that quantifies parameter correlations through posterior covariance structure.

The sensitivity results therefore do more than rank parameters. They define a reduced but physically meaningful calibration space and show where correlations and trade-offs must be expected in the posterior distributions. The 14 parameters identified here proceed to regime-stratified Bayesian calibration in Chapter 7, where optimal values are estimated separately for the three burnup regimes identified in Chapter 5. This approach enables quantification of regime-dependent parameter evolution while maintaining computational tractability through dimensionality reduction from 22 to 14 calibration parameters.

Chapter 7

Uncertainty Quantification and Bayesian Calibration

7.1 Introduction and Scientific Motivation

Computational models of physical systems inherently contain uncertainty across multiple dimensions. Parameters governing material behavior cannot be measured directly under operating conditions. Numerical approximations introduce discretization error. Boundary conditions and initial states reflect limited observational resolution. The accumulated effect of these uncertainties produces predictions that carry substantial confidence intervals, restricting the utility of mechanistic models for engineering decision-making and regulatory assessment.

Uncertainty quantification provides a formal framework for identifying, characterizing, and propagating these uncertainties through computational models. Rather than treating model outputs as deterministic point estimates, UQ methods recognize that both model inputs and structural limitations generate distributions of plausible outcomes. This probabilistic perspective transforms mechanistic models from prediction tools into confidence-bounding tools, enabling designers to understand not only what a model predicts but also the range of confidence warranted by available information.

For mechanistic fuel performance models such as BISON, this distinction proves critical. Sensitivity analysis identified that the UN swelling and fission gas release model contains approximately 14 uncertain parameters spanning a wide range of mechanistic influence. These parameters control predicted fission gas release and swelling through coupled nonlinear physics across scales ranging from atomic diffusion to centimeter-scale pellet behavior. Direct experimental measurements of many of these parameters are limited or remain infeasible. Instead, observable quantities like fission gas release and swelling measured from irradiation experiments may constrain parameter values inversely through the model structure.

This inverse relationship creates two distinct but complementary UQ problems. Forward uncertainty quantification propagates parameter uncertainties through the mechanistic

model to generate predicted distributions of observable quantities. By identifying which parameters most strongly influence fission gas release and swelling predictions, forward UQ reveals where experimental efforts should focus to reduce predictive uncertainty most efficiently.

Inverse uncertainty quantification reverses this direction. Given experimental measurements of fission gas release and swelling, inverse UQ estimates probability distributions for the model parameters themselves. This calibration process updates prior beliefs about parameter values using the information content of experimental data. The posterior distributions resulting from inverse UQ quantify remaining parameter uncertainty after accounting for what experiments reveal. These posterior distributions then propagate forward through the model to generate realistic prediction bounds that reflect both parameter uncertainty and model structural limitations.

The mechanistic UN swelling and fission gas release model integrated into BISON requires inverse calibration against historical UN irradiation data spanning multiple burnup levels and temperatures. This calibration challenge reduces to a high-dimensional parameter inference problem constrained by expensive forward model evaluations. Each BISON simulation requires 10 to 30 minutes of computational time. Rigorous parameter calibration through direct sampling methods demands thousands to tens of thousands of model evaluations to achieve adequate posterior convergence in 14-dimensional parameter space. Standard Bayesian methods applied directly to the full BISON model would therefore require 100–500 days of dedicated computing resources per calibration attempt, rendering the approach impractical for nearly all applications.

This computational bottleneck necessitates acceleration through surrogate modeling. Gaussian Process surrogate models trained on a Latin Hypercube sampling design of BISON simulations provide fast approximations to full model outputs while quantifying approximation uncertainty. By replacing expensive BISON evaluations with trained surrogates during Bayesian parameter inference, computational demands reduce from weeks to hours while maintaining statistical rigor through explicit treatment of surrogate prediction uncertainty in the likelihood function.

The present chapter documents a complete inverse uncertainty quantification workflow for the UN swelling and fission gas release model. Initial investigations into direct MCMC inference using the INL MOOSE Stochastic Tools framework were undertaken to understand computational requirements for the 14-dimensional parameter estimation problem. Development of multi-observable extensions to the framework, including AIE sampler architecture modifications and multi-variable likelihood formulations, revealed that computational costs remained prohibitive despite algorithmic optimization and HPC parallelization. As documented in Appendix A, the direct MCMC approach required approximately 95 core-days of computing resources per iteration, necessitating 400 to 800 days of continuous computing for adequate posterior convergence. This fundamental computational barrier motivated the transition to surrogate-based acceleration.

Development of the direct INL framework provided essential insights into multi-variable likelihood structures, variance parameter handling, and ensemble sampling that directly

informed surrogate-accelerated calibration design. Gaussian Process emulators trained on a Latin Hypercube sampling design across the 14-dimensional parameter space provide millisecond-scale predictions with quantified approximation uncertainty. By replacing expensive full-model evaluations with surrogate predictions during MCMC sampling, computational demands reduce dramatically while maintaining Bayesian statistical rigor through explicit treatment of surrogate approximation error in the likelihood.

Surrogate-model-accelerated MCMC enables systematic exploration of parameter space: Latin Hypercube sampling with 2,100 BISON evaluations provides far denser coverage of the 14-dimensional parameter space than sequential MCMC chains could achieve. This comprehensive sampling reveals the full structure of the parameter-to-output mapping, including nonlinearities and regime-dependent interactions that sequential MCMC chains might never discover within practical computational budgets. Multi-chain No-U-Turn Sampler (NUTS) sampling with eight independent chains generates posterior parameter distributions accounting for regime-dependent mechanistic complexity. Regime analysis partitions the parameter space into coherent subsets where model behavior remains stable. Finally, forward uncertainty propagation using posterior samples produces realistic prediction distributions incorporating both calibrated parameter uncertainty and persistent epistemic model limitations inherent to mechanistic fuel performance models.

This framework transforms inverse UQ from a computationally intractable problem into one that is achievable within practical timescales while maintaining statistical defensibility. The resulting posterior parameter distributions quantify mechanistic model uncertainty, supporting improved predictive accuracy for accelerated fuel qualification assessments and guiding targeted experimental efforts to constrain parameters that remain poorly identified by existing data.

7.2 Bayesian Inference and MCMC Theory

Bayesian inference provides a rigorous probabilistic framework for updating beliefs about model parameters in light of observational data. Rather than seeking a single point estimate of parameters, Bayesian methods produce full probability distributions that quantify parameter uncertainty while accounting for both data constraints and prior knowledge. This distribution-based approach proves particularly valuable for expensive computational models where each evaluation demands significant computing resources.

The foundation of Bayesian inference is the Bayes' theorem, which relates the posterior distribution (updated parameter beliefs after observing data) to three foundational quantities [72]. The prior distribution $p(\boldsymbol{\theta})$ encodes initial beliefs about parameter values before experiments are conducted. The likelihood function $p(\mathbf{d} \mid \boldsymbol{\theta})$ quantifies how well a given parameter set reproduces observed experimental data. Bayes' theorem combines these through

$$p(\boldsymbol{\theta} \mid \mathbf{d}) = \frac{p(\mathbf{d} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{d})}, \quad (7.1)$$

where the posterior $p(\boldsymbol{\theta} \mid \mathbf{d})$ represents the updated parameter distribution after assimilating experimental observations. The denominator $p(\mathbf{d})$, called the model evidence, represents a normalization constant that typically cannot be computed analytically for high-dimensional problems.

For fuel performance modeling, the posterior distribution enables answering questions such as the following. Given measured swelling of 4.5% and fission gas release of 0.08% at 1.16% burnup, what values of dislocation density, diffusivity, and creep parameters are most consistent with these observations? The posterior $p(\boldsymbol{\theta} \mid \mathbf{d})$ provides not a single answer but a probability distribution over the 14-dimensional parameter space, with high probability density in regions where parameter combinations reproduce observations and low probability density where predictions deviate substantially from data.

Direct evaluation of the posterior distribution proves intractable when parameters number more than a handful because the integration required to compute the evidence demands evaluating the likelihood across the entire parameter space. MCMC methods circumvent this analytical barrier by generating samples that asymptotically follow the posterior distribution without explicitly computing the evidence. By constructing a Markov chain where each state represents a parameter proposal and transition probabilities depend only on the current state and the likelihood function, MCMC algorithms produce a sequence of samples whose empirical distribution converges to the true posterior.

The essential requirement for MCMC validity is that transitions satisfy detailed balance, ensuring that the chain's stationary distribution equals the posterior. Different MCMC samplers achieve this through different proposal mechanisms. Metropolis-Hastings samplers [73] propose candidate parameter values and accept or reject them based on a likelihood ratio test. Modern samplers such as the NUTS [74] employ Hamiltonian dynamics to generate proposals that explore parameter space more efficiently, particularly in high-dimensional settings.

For parameter inference in fuel performance models, MCMC methods prove essential because they enable rigorous treatment of high-dimensional uncertainty. Evaluating all 14 uncertain parameters and their correlations through classical experimental design would require thousands of dedicated irradiation experiments. MCMC instead leverages the information content in existing historical data by computing the likelihood at each proposed parameter configuration and accepting or rejecting proposals based on agreement with observations. Over many iterations, the chain converges to regions where model predictions align with data, producing posterior samples that quantify remaining parameter uncertainty after accounting for all experimental constraints.

7.3 Surrogate-Accelerated Bayesian Inference Framework

7.3.1 Motivation and Computational Challenge

Initial investigations into direct Bayesian calibration employed the INL MOOSE Stochastic Tools framework to establish a baseline understanding of computational requirements for the 14-dimensional parameter estimation problem. These exploratory efforts, detailed fully in Appendix A, implemented an Affine Invariant Ensemble sampler with multi-observable likelihood extensions. While the framework demonstrated algorithmic soundness and proper treatment of multi-variable observations, computational performance rapidly became prohibitive. The extended framework was deployed on a high-performance computing cluster with 28 walkers operating in parallel across 336 cores total. Each BISON simulation required 10 to 30 minutes depending on burnup and temperature conditions, with significant variability arising from nonlinear solver convergence behavior at extreme burnup values. Over a 7-day maximum-allocation continuous run, the framework completed approximately 400 iterations of the AIE sampler, corresponding to approximately 95 core-days of computing resources per MCMC iteration. This computational investment extrapolates to 400 to 800 days of dedicated computing resources to achieve adequate posterior convergence in 14-dimensional parameter space, with convergence requiring 5,000 to 10,000 iterations per chain to reach stationarity and explore the full posterior support with reliable quantile estimates.

This fundamental computational barrier stems not from algorithmic deficiency (the extended Bayesian theory is rigorous, the ensemble MCMC implementation is sound, and the multi-variable architecture is properly validated) but rather from the inherent expense of mechanistic modeling. Each model evaluation involves solving coupled nonlinear PDEs across 2D fuel pin geometry with temperature-dependent properties and state-dependent source terms. While advanced implicit solvers and HPC parallelization have reduced individual evaluation time to 10 to 30 minutes, this fundamental limitation renders direct inference impractical.

Parameter dimensionality reduction through fixing insensitive parameters would discard potentially informative physics, biasing posteriors away from experimental constraints. Crude approximations to BISON through reduced geometry or simplified physics would compromise mechanistic rigor that motivates expensive forward modeling in the first place.

Instead, computational acceleration was achieved through surrogate modeling, a mathematically principled approach that maintains Bayesian validity while achieving three to four orders of magnitude speedup. Development of the direct INL framework, while ultimately superseded by surrogate acceleration, provided essential insights into multi-variable likelihood structures, variance parameter handling, and ensemble sampling that directly informed surrogate-accelerated calibration design. Rather than continuing with expensive BISON simulations, Gaussian Process emulators trained on a Latin Hypercube sampling design of BISON simulations provide millisecond-scale predictions with quantified uncertainty. By

replacing expensive full-model evaluations with surrogate predictions during MCMC sampling, computational demands reduce dramatically while maintaining Bayesian statistical rigor through explicit treatment of surrogate approximation error in the likelihood.

The transition from direct BISON-coupled MCMC to surrogate-accelerated inference represents more than computational expediency. Surrogate modeling enables systematic exploration of parameter space that would be impossible via direct coupling: Latin Hypercube sampling with 2,100 evaluations provides far denser coverage of the 14-dimensional parameter space than MCMC's 400 sequential evaluations. This comprehensive sampling not only trains accurate emulators but also reveals the full structure of the parameter-to-output mapping, including nonlinearities and interactions that sequential MCMC chains might never discover within practical computational budgets.

7.4 Surrogate Model Development

7.4.1 Design of Experiments and Training Data

The surrogate development pipeline began with systematic exploration of the mechanistic model across the 14-dimensional parameter space via Latin Hypercube sampling (LHS). LHS provides space-filling properties superior to naive random sampling, ensuring that parameter combinations explore the full feasible region rather than clustering near prior means. An initial LHS design generated 2,100 samples uniformly distributed in normalized parameter space $\xi \in [0, 1]^{14}$, transformed to physical parameter ranges reflecting three distinct scaling strategies corresponding to different sources of parameter uncertainty.

The LHS design was evaluated across nine experimental test conditions drawn from the SNAP-50 UN irradiation database: ORNL specimens 602, 603, 658, 662, 664, 667, 669, and 665, and CANEL specimens 602, 603, and 220. Each test encompasses three axial fuel pin positions (bottom, middle, top) where swelling and fission gas release were measured. The choice of multiple test cases reflects the reality that fuel model parameters must produce predictions consistent across diverse burnup levels (0.6 to 4.6 %FIMA), temperatures (1200 to 1800 K), and power densities (0.2 to 1.5 kW/cc), each representing distinct physical regimes within irradiation experiments.

The raw LHS-based simulation data contained numerical artifacts arising from occasional BISON convergence difficulties or parameter combinations far outside the physically relevant regime. Outlier detection and removal proceeded via an interquartile range (IQR) method applied to the maximum value observed across axial positions for each observable. Statistical outliers (defined as values exceeding $Q_3 + 2.5 \times \text{IQR}$) were excluded from surrogate training. This outlier removal strategy successfully eliminated approximately 5 to 15% of samples per test case while preserving the core distribution of physically meaningful outputs and maintaining 99% of the dynamic range in parameter values.

7.4.2 Gaussian Process Surrogate Specification

Gaussian Process (GP) surrogates model the relationship between normalized parameters $\xi \in \mathbb{R}^{14}$ and BISON predictions $y(\xi) \in \mathbb{R}$ as a realization from a Gaussian process. The covariance function determines smoothness and length-scale structures captured by the surrogate

$$\text{Cov}[y(\xi), y(\xi')] = k(\xi, \xi'; \phi) + \sigma_n^2 \delta(\xi - \xi'), \quad (7.2)$$

where k is the kernel function, σ_n^2 represents observational noise and ϕ denotes kernel hyperparameters. Three kernel families were systematically compared: the Matérn 3/2 kernel, the Matérn 5/2 kernel, and the Exponential Quadratic (RBF) kernel. Matérn kernels are particularly well-suited to fuel performance modeling because their finite smoothness (Matérn $\nu = 3/2, 5/2$) appropriately reflects the nonlinear, sometimes non-smooth responses of complex mechanistic models. RBF kernels assume infinite smoothness inappropriate for fuel performance predictions, and comparison studies confirmed that Matérn kernels outperformed RBF approaches across all test conditions.

The choice between Matérn and RBF kernels carries physical implications for fuel performance modeling. RBF kernels assume infinitely differentiable functions, implying that small parameter changes always produce smooth, gradual output changes. This assumption fails for fuel performance models where threshold phenomena exist. For example, the transition from closed to interconnected porosity occurs abruptly as swelling exceeds critical values. Matérn kernels with finite smoothness ($\nu = 3/2, 5/2$) permit the surrogate to capture these transitions through piecewise-smooth behavior, making them more appropriate for mechanistic models with distinct physical regimes.

Kernel hyperparameters (length scales $\ell \in \mathbb{R}^{14}$, signal variance η , and noise variance σ_n) were inferred via Maximum A Posteriori (MAP) estimation of the marginal likelihood. Two gradient-based optimization algorithms were compared: L-BFGS-B (limited-memory BFGS with box constraints) for computational efficiency and Powell's derivative-free method for robustness when L-BFGS-B converged to poor local optima. For each surrogate, both optimizers and all three kernels were executed independently, and the result yielding the higher test-set R^2 (coefficient of determination) was retained. R^2 defines the overall accuracy of the predicted surrogate outputs compared to the known experimental data, providing an overall prediction quality. Therefore this dual-optimizer strategy balanced efficiency with reliability.

The output space transformation proved essential for numerical stability and improved surrogate accuracy. Both swelling and fission gas release exhibit heteroscedastic behavior with output uncertainty and derivatives varying by orders of magnitude across the parameter space. Adaptive log-transformation with shift parameters handled zero and near-zero values appropriately

$$y_{\log} = \log(y_{\text{phys}} + \Delta), \quad (7.3)$$

where the shift Δ was chosen adaptively per quantity to prevent ill-conditioning near zero. This transformation reduced heteroscedasticity, improved Gaussian assumptions, automatically handled the zero-inflated fission gas release distribution, and enabled symmetric uncertainty bands in physical space when inverted.

7.4.3 Surrogate Training and Accuracy Validation

Training of 54 independent surrogates proceeded with a 75% / 25% train/test split (approximately 1500 training and 500 test samples per surrogate), ensuring reliable unbiased accuracy estimates on held-out data. Performance metrics across the surrogate ensemble demonstrate excellent accuracy suitable for MCMC acceleration. Swelling surrogates consistently achieved mean R^2 of 0.88, reflecting the physically deterministic nature of irradiation-induced void swelling. Fission gas release surrogates showed greater variability (mean R^2 of 0.84, range 0.74–0.91) reflecting the inherent stochasticity of bubble nucleation, coalescence, and release. Mean absolute error in log-space were 0.13% for swelling and 0.35% for fission gas release, small compared to experimental observations typically spanning 1 to 7% for swelling and 0.01 to 12% for fission gas release.

Prediction accuracy is quantified by R^2 and mean absolute error (MAE), where MAE is described by the following Equation 7.4:

$$\text{MAE} = \frac{1}{n} \sum_{k=1}^n |\hat{y}_k - y_k| \quad (7.4)$$

where $\hat{y}_k$ is the posterior predictive mean and y_k is the observation for location k . Surrogate uncertainty calibration proves critical for Bayesian inference. An accurate surrogate that underestimates its own uncertainty will induce artificially narrow posteriors, producing overconfident parameter estimates and misleading credible intervals. Conversely, excessive conservatism in uncertainty estimates dilutes the information content of experimental data, preventing the posterior from concentrating near parameter values that reproduce observations.

7.4.3.1 Empirical Uncertainty Calibration

This section directly addresses the challenge through ensuring that surrogate-predicted uncertainty intervals achieve nominal 95% coverage on validation data, guaranteeing that the surrogate’s self-assessment of approximation error is neither optimistic nor excessively pessimistic.

Adaptive observational noise variance estimation, where the noise term σ_n^2 was inferred rather than fixed, proved essential for accurate predictions. Nonadaptive noise assumptions led to systematic underfitting when local model variance differed from global assumptions. The learned noise parameters provide realistic estimates of surrogate prediction uncertainty suitable for Bayesian likelihood formulations. Beyond point predictions, GP surrogates

provide posterior predictive uncertainties reflecting both aleatory uncertainty (intrinsic variability in model predictions) and epistemic uncertainty (knowledge limitations from finite training data).

Surrogate predictions must account for systematic biases and miscalibration of predictive uncertainties. Calibration inflation factors were computed from test-set residuals to ensure empirically correct coverage. The standardized residuals on test data were computed as

$$z_i = \frac{y_i^{\text{true}} - \mu(\boldsymbol{\xi}_i)}{\sigma(\boldsymbol{\xi}_i)}, \quad (7.5)$$

where $\mu(\boldsymbol{\xi}_i)$ and $\sigma(\boldsymbol{\xi}_i)$ are the GP posterior mean and standard deviation. For a well-calibrated surrogate, these residuals should follow $\mathcal{N}(0, 1)$. The standard deviation of standardized residuals,

$$\sigma_z = \sqrt{\frac{1}{N_{\text{test}}} \sum_i z_i^2}, \quad (7.6)$$

quantifies miscalibration. Recommended inflation factors are computed as $\max(1.0, \sigma_z)$. Across the 54 surrogates (9 tests $\times$ 3 positions $\times$ 2 quantities), median calibration inflation was 1.1 to 1.3, indicating systematic modest underestimation of GP uncertainty. This underestimation is expected when GPs are trained on BISON output affected by numerical approximation errors, and the empirically-determined inflation factors ensure valid posterior credible regions.

7.5 MCMC Calibration: Methodology, Convergence, and Parameter Uncertainty

7.5.1 Surrogate-Based Likelihood and MCMC Architecture

The Gaussian Process surrogates developed in Section 7.4 provide fast approximations to BISON outputs across the 14-dimensional normalized parameter space. For each experimental test position and each observable quantity (swelling and fission gas release), an independent GP surrogate maps parameter vectors $\boldsymbol{\xi} \in [0, 1]^{14}$ to log-transformed predictions through a posterior predictive distribution accounting for both epistemic uncertainty from finite training data and approximation error from kernel specification.

The surrogate-based likelihood function explicitly accounts for both surrogate approximation uncertainty (varying spatially across parameter space based on distance to training data) and experimental measurement uncertainty (constant across parameter space but potentially observable-specific). This formulation maintains Bayesian validity by propagating surrogate uncertainty through the likelihood rather than treating surrogate means as deterministic BISON replacements. Calibration inflation factors computed from validation data ensure that surrogate uncertainty estimates achieve nominal 95% coverage on held-out

test sets, preventing artificially narrow posteriors that would result from underestimating approximation error.

The critical insight enabling surrogate-accelerated MCMC is that the GP predictive variance $\sigma_{\text{GP}}^2(\boldsymbol{\xi}^*)$ quantifies approximation uncertainty at any parameter configuration. This variance increases in regions far from training data and decreases near densely sampled areas, providing a spatially-varying confidence estimate for surrogate predictions. Rather than treating the GP mean μ_{GP} as a deterministic BISON replacement, the full predictive distribution must be propagated through the likelihood to avoid artificially narrow posteriors and overconfident parameter estimates.

7.5.2 No-U-Turn Sampler Selection and High-Dimensional Efficiency

The choice of MCMC algorithm profoundly affects computational efficiency, posterior convergence, and practical feasibility for high-dimensional calibration problems. The No-U-Turn Sampler (NUTS), a gradient-based algorithm implementing Hamiltonian Monte Carlo with automatic trajectory length adaptation, was selected for its exceptional performance in 14-dimensional parameter spaces. NUTS leverages gradient information from the log-posterior to propose distant parameter configurations that maintain acceptance probabilities of 60–90%, compared to $< 1\%$ for random-walk methods. Automatic tuning during warm-up eliminates the need for manual step size selection, and the U-turn criterion automatically terminates trajectories when they begin retracing explored parameter space, preventing wasted computation from excessively long proposals.

For surrogate-based likelihoods where gradient evaluations are fast (milliseconds per evaluation), NUTS explores the posterior 100 to 1,000 times more efficiently than gradient-free methods. Multi-chain NUTS sampling with eight independent chains ensures robust convergence assessment through Gelman-Rubin diagnostics, enabling identification of any chains that fail to mix properly. This parallelizable structure permits efficient use of HPC resources to achieve convergence in practical timescales (hours rather than weeks).

7.5.3 Prior Distribution Specification and Experimental Uncertainty

Prior distributions encode initial beliefs about parameter values before assimilating experimental data. For the UN fuel model, prior specification reflects the principle that sufficient experimental data exists to allow the likelihood to dominate posterior formation without strong prior constraints. All model parameters received uniform priors on the normalized interval $p(\xi_i) = \text{Uniform}(0, 1)$ for all $i = 1, \dots, 14$, where normalization ensures that all parameters are treated symmetrically in the NUTS algorithm. This maximum-entropy specification imposes no structure beyond the parameter bounds themselves, reflecting weak prior information about parameter values.

The experimental noise parameters $\sigma_{\text{exp,swell}}$ and $\sigma_{\text{exp,fgf}}$ received half-normal priors $p(\sigma_{\text{exp}}) = \text{HalfNormal}(\sigma = 0.05)$, reflecting the expectation that measurement uncertainties are positive and concentrated near small values. This prior remains weakly informative because the data likelihood strongly constrains noise parameters when residuals between predictions and observations exhibit systematic structure. Posterior samples show that $\sigma_{\text{exp,swell}}$ settled near the prior mode (mean posterior approximately 0.06), while $\sigma_{\text{exp,fgf}}$ adjusted upward (mean posterior approximately 0.12–0.18 depending on regime), reflecting the greater stochasticity in fission gas release measurements.

7.5.4 Convergence Diagnostics Validating MCMC Stationarity

Ensuring posterior convergence is critical for valid statistical inference. MCMC samplers generate correlated samples that only asymptotically follow the true posterior; premature termination produces biased parameter estimates and underestimated uncertainties. Rigorous convergence assessment employs multiple complementary diagnostics that collectively verify stationarity, adequate mixing, and exploration of the full posterior support.

The Gelman-Rubin $\hat{R}$ diagnostic compares between-chain variance to within-chain variance such that values $\hat{R} < 1.01$ indicate convergence where all chains have reached the same distribution. For all successful calibrations, $\hat{R} < 1.05$ for every parameter and experimental noise parameter, with $\hat{R} < 1.01$ for the 11-parameter configuration across all mechanistic tests.

Effective sample size (ESS) quantifies the reduction in independence due to autocorrelation, with $\text{ESS} > 1,000$ generally considered adequate for reliable posterior quantiles. All converged calibrations achieved $\text{ESS} > 1,500$ for dominant parameters (those with Morris $\mu^* > 1.0$) and $\text{ESS} > 500$ for less influential parameters, reflecting NUTS’s efficient exploration via gradient-guided proposals.

Visual inspection of trace plots confirms that chains mix well and explore the full posterior support without getting trapped in local modes. Well-mixed chains exhibit a “fuzzy caterpillar” appearance with no apparent trends, drifts, or abrupt transitions. Divergence statistics (NUTS-specific diagnostics identifying numerical issues during Hamiltonian trajectory integration) remained below 0.5% across all successful calibrations, confirming that NUTS effectively navigated the 14-dimensional posterior without numerical instabilities.

7.6 Prior Refinement Through Mechanistic Guidance and Regime-Stratified Calibration

7.6.1 Prior Specification Strategy: From Uninformed to Mechanistically-Guided

Bayesian inference requires specification of prior distributions encoding initial beliefs about parameter values before observing experimental data. For high-dimensional problems

where data provide strong constraints relative to prior uncertainty, prior selection proves less critical to posterior formation than in weakly-informed regimes. However, the strategic incorporation of mechanistic knowledge through prior specification can substantially improve MCMC convergence and sampler efficiency, particularly when exploring high-dimensional parameter spaces where naive uniform priors may leave the sampler to wander through physically implausible regions.

The present calibration employs a two-tier prior specification strategy. The baseline approach assigns uniform priors to all 14 mechanistic parameters, reflecting maximum-entropy assumptions where knowledge remains genuinely uncertain. A refined approach incorporates mechanistic guidance from Chapter 5 parametric studies and physical literature to specify more informative prior distributions for parameters where physical understanding permits reasonable constraints. Chapter 5 investigations revealed that fission gas release favored scaled down trapping mechanisms, with scaled up resolution and diffusion, for example. Literature bounds on fractional yield strength and surface energy, drawn from independent uranium characterization studies, provide additional prior guidance. Parameters without strong mechanistic constraints or literature support retain uniform priors, reflecting epistemic uncertainty appropriate to the current state of knowledge.

Where strong literature values provide sufficient information, a normal distribution was applied. When setting a prior based on Chapter 5, a weakly informed beta function shifted towards the domain of success was implemented. All other parameters were set with a uniform distribution. This hybrid prior specification maintains Bayesian validity (the likelihood dominates posterior formation in well-constrained regimes) while leveraging available mechanistic information to regularize parameter exploration in poorly-constrained directions. The refinement strategy improves MCMC efficiency without imposing artificial constraints that would bias posteriors away from experimental evidence.

7.6.1.1 Regime Definitions and Experimental Coverage

Table 7.1 summarizes the five calibration regimes from Chapter 5, listing specimens included in each partition, their operational ranges, and the number of experimental observations available per regime.

Table 7.1: MCMC calibration regimes with experimental coverage and operating ranges.

Regime	Specimens	Temperature (K)	Burnup (%FIMA)
Test a	220	1561–1714	0.95
Test b	602, 603, 658, 662, 664, 669	1213–1503	0.54–2.72
Test c	658, 662, 667	1388–1523	0.78–1.16
Test d	602, 603, 658, 662, 664	1269–1523	0.54–1.08
Test e	665	1323–1426	4.58

The primary calibration regime (Test d: Low-Temperature, Low-Burnup) encompasses 17 experimental observations across six specimens spanning temperatures from 1213 K to 1503

K and burnup from 0.54%FIMA to 1.08%FIMA. This regime aggregates data representing conditions where the mechanistic model was developed and where physical understanding of void swelling and fission gas migration mechanisms is most mature.

Three secondary regimes partition the experimental space to isolate regime-specific effects. Test a (High-Temperature) isolates specimen 220 (temperatures 1561–1714 K, 0.95%FIMA burnup), representing anomalously high temperatures where fuel microstructure transitions and diffusion mechanisms may differ from lower-temperature regimes. Test e (High-Burnup) isolates specimen 665 (4.58%FIMA, temperatures 1323–1426 K), representing high fission gas inventory conditions where bubble coalescence and interlinkage phenomena may dominate swelling behavior. Test b is similar to low-temperature, low-burnup but includes specimen 669 from the primary regime to test sensitivity of parameter posteriors to inclusion of the higher-burnup specimen within the nominally low-burnup category. Test 669 was on the edge of the regime identification performed earlier and therefore warrants investigation. Test c represents a combined high-temperature, high-burnup partition including specimens 658, 662, and 667 to assess whether moderate-temperature, moderate-to-high burnup conditions show distinct mechanistic behavior. This was the outlier regime that showed increased grain boundary coverage and fission gas release when temperatures increased above 1485K and burnup of $\approx 1\%$ FIMA.

7.6.2 Impact of Prior Refinement on MCMC Convergence

To quantify the benefits of mechanistically-informed priors, MCMC calibration was performed with both uninformed (uniform) and informed (Chapter 5 and literature-guided) prior specifications across multiple experimental regimes. Convergence diagnostics reveal the systematic impact of prior refinement on sampler behavior and posterior reliability.

Figure 7.1 illustrates convergence metrics across the calibration regimes. With uniform priors, several parameters exhibit $\hat{R}$ values exceeding 1.01, indicating incomplete convergence, particularly in temperature-extreme regimes where data scarcity prevents the likelihood from fully constraining parameter space. ESS values remain moderate (5,000 to 15,000), reflecting less efficient exploration. By contrast, informed priors systematically improve convergence: $\hat{R} < 1.01$ universally across all mechanistic parameters in the primary calibration domains, and ESS increases to 15,000 to 25,000, indicating that informed priors guide the sampler to high-probability regions more directly. This convergence improvement stems from priors concentrating prior mass near regions where data evidence concentrates, enabling the sampler to explore relevant parameter space more efficiently rather than dissipating effort on physically implausible configurations.

The convergence advantage proves most dramatic in the low-temperature, low-burnup regime (Figure 7.1, Test b), where informed priors yield universal $\hat{R} < 1.00$ and maximum ESS across parameters. Anomalous regimes (high-temperature specimen 220 [Test a] and high-burnup specimen 665 [Test e]) demonstrate persistent convergence challenges even with informed priors, a phenomenon explored in depth in Section 7.7.




Figure 7.1: Convergence diagnostics across five calibration regimes (Tests a–e) comparing uninformed (uniform) priors (top row) and informed (Chapter 5 and literature-guided) priors (bottom row). Left panels: Gelman-Rubin $\hat{R}$ diagnostic showing convergence status per parameter. Green ($\hat{R} < 1.01$) indicates convergence; yellow/orange ($\hat{R} > 1.02$) indicates incomplete mixing. Right panels: Effective sample size (ESS, bulk estimate) quantifying independent posterior samples. Dark blue (25,000+ ESS) indicates efficient sampling; light yellow (< 5,000 ESS) indicates inefficient exploration.

Comparison between Test b (Low-Temperature, Low-Burnup, including specimen 669) and Test d (excluding specimen 669) reveals regime-dependent shifts in convergence behavior. Although both tests remain within the low-temperature, low-burnup domain, inclusion of the higher-burnup specimen 669 produces moderately elevated $\hat{R}$ and reduced ESS relative to Test d, suggesting that this specimen occupies a regime boundary where mechanistic parameter relationships differ from the strictly low-burnup domain.

A critical observation emerges when comparing uniform and informed prior configurations: while most experimental regimes demonstrate improved convergence and elevated ESS with informed priors, Test a (high-temperature 220 alone) exhibits anomalous behavior—ESS decreases upon prior refinement despite $\hat{R}$ improvement. This counterintuitive finding

(improved convergence metrics yet reduced independent sampling) indicates fundamental regime-dependent differences in mechanistic model structure. The decline in effective sampling efficiency despite enhanced convergence diagnostics suggests that informed priors, designed to regularize parameter exploration in the well-constrained low-temperature domain, may inadvertently introduce misspecification in extreme-temperature regimes where different physical mechanisms may dominate.

7.6.3 Enhanced Parameter Identification Through Informed Priors

Beyond convergence, informed priors enhance parameter identification by enabling the data likelihood to constrain posterior distributions more tightly relative to prior uncertainty. The degree of information extraction is quantified through variance reduction and precision gain metrics.

Variance reduction, defined as $1 - (\sigma_{\text{posterior}}/\sigma_{\text{prior}})^2$, measures the fractional reduction in parameter uncertainty from prior to posterior. A value of 1.0 indicates complete information extraction (posterior $\sigma \rightarrow 0$), while 0.0 indicates no information gain (posterior $\sigma = \text{prior } \sigma$). Precision gain, defined as prior standard deviation divided by posterior standard deviation, directly indicates how much tighter the posterior is relative to the prior. Values > 1.5 indicate substantial parameter identification; values < 1.1 suggest the posterior remains broad relative to prior width.

Figure 7.2 displays variance reduction (left) and precision gain (right) across the calibration regimes and all 14 mechanistic parameters. Uniform priors baseline show modest information gain: variance reduction typically 0.2 to 0.6, indicating that experimental data moderately constrain parameters but substantial prior-to-posterior uncertainty persists. With informed priors, variance reduction increases substantially: key mechanistic parameters (dislocation trapping, resolution scaling, diffusion factors) achieve variance reduction exceeding 0.8, and precision gain frequently exceeds 2.0 in the primary calibration domain.

The regime contrast is striking: the low-temperature, low-burnup regime (Test b) exhibits the strongest information extraction, with variance reduction exceeding 0.8 and precision gain > 1.5 for mechanistically-critical parameters (trapping, resolution, diffusion). In marked contrast, anomalous regimes (Test a: high-temperature 220, Test e: high-burnup 665) display pale coloration (variance reduction < 0.4 , precision gain near 1.0), indicating minimal parameter identification even with mechanistically-informed priors. This dramatic regime dependence reveals that information content is fundamentally asymmetric across the experimental database: low-temperature, low-burnup conditions provide rich experimental constraints that tightly bound parameter uncertainty, whereas extreme-condition specimens (220, 665) contain insufficient observational signal to overcome the inherent non-identifiability of the 14-dimensional parameter space. The pale coloration in Tests a and e reflects data scarcity rather than mechanistic failure; additional specimens or higher-fidelity measurements in these regimes would be necessary to meaningfully constrain




Figure 7.2: Information extraction from experimental data quantified through variance reduction (left) and precision gain (right), across five calibration regimes (Tests a–e) and all 14 mechanistic parameters. Variance reduction (left): values approaching 1.0 (dark green) indicate strong posterior tightening relative to prior; pale yellow (0.0–0.2) indicates minimal information gain. Precision gain (right): values > 1.5 (dark blue) indicate posterior precision exceeds prior by $> 50\%$; values near 1.0 (pale yellow) indicate posteriors remain nearly as wide as priors.

parameter values in extreme conditions.

7.6.4 Systematic MCMC Calibration Across Experimental Regimes

To understand calibration feasibility across the full experimental data space and identify the regime(s) where sufficient data exist for rigorous parameter inference, MCMC was executed with informed priors across five regime partitions of the SNAP-50 UN database. These regimes partition experiments by operating conditions (temperature, burnup) to

test whether mechanistic parameter values converge to well-defined posteriors, or whether different experimental subsets induce different parameter values, suggesting either regime-dependent mechanisms or regime-specific data scarcity.

7.6.4.1 Convergence and Information Gain by Regime

Informed prior MCMC applied across these five regimes (Figure 7.1 and Figure 7.2, bottom rows) reveals dramatic differences in parameter identifiability and convergence behavior. Test b (low-temperature, low-burnup regime) achieves superior convergence across all diagnostics: all parameters achieve $\hat{R} < 1.01$ (indicating stationarity across all eight independent chains), ESS $> 15,000$ for dominant parameters (many exceeding 20,000), and variance reduction exceeding 0.8 for mechanistically-important parameters (trapping, resolution, diffusion). This regime emerges as the validated calibration domain where experimental data sufficiently constrain the 14-dimensional parameter space.

The high-temperature (Test a: 220 alone) and high-burnup (Test e: 665 alone) regimes exhibit fundamentally distinct convergence signatures despite mechanistically-informed prior specification. Both anomalous regimes show incomplete convergence ($\hat{R} > 1.01$ for multiple parameters), substantially reduced effective sample sizes (ESS 2,000–8,000), and minimal parameter identification (variance reduction < 0.3 for most parameters, precision gain near 1.0). This convergence degradation stems fundamentally from data insufficiency rather than algorithmic failure: single-specimen regimes provide limited observational signal to constrain the high-dimensional parameter space. Posterior distributions in these regimes either fail to concentrate around unique parameter values or exhibit multimodal structures with multiple distinct parameter combinations fitting the available data equally well. The persistence of poor convergence diagnostics even with mechanistically-guided priors indicates that prior refinement cannot compensate for inadequate experimental constraints. These regimes are reserved for validation and mechanistic comparison studies (Section 7.7) rather than serving as primary calibration domains.

7.6.4.2 Selection of Informed Priors and Multi-Regime Analysis Strategy

The convergence diagnostics and information gain analysis presented above establish informed priors as the optimal configuration for MCMC-based parameter inference in the UN swelling and fission gas release model. Across all well-constrained experimental regimes, mechanistically-guided priors systematically outperform uniform (maximum-entropy) priors across three critical dimensions: convergence to stationarity ($\hat{R} < 1.01$ universally versus $\hat{R} > 1.01$ in multiple anomalous regimes), effective sampling efficiency (ESS 15,000–25,000 versus 5,000–15,000), and parameter identification (variance reduction 0.7–0.9 versus 0.2–0.6). The incorporation of mechanistic guidance from Chapter 5 parametric studies and physical literature constrains parameter exploration toward physically plausible regions, enabling the MCMC sampler to converge more rapidly and achieve superior coverage of the posterior support. By leveraging available mechanistic knowledge without imposing artificial

constraints that would bias posteriors away from experimental data, informed priors maintain Bayesian validity while dramatically improving inference quality.

Accordingly, all subsequent analysis proceeds with mechanistically-informed prior specifications. The dual-figure analysis (Figure 7.1 and Figure 7.2, bottom rows) demonstrates that informed priors enable reliable parameter inference in regimes with sufficient experimental data (Test b: low-temperature, low-burnup), while simultaneously revealing the fundamental data insufficiency underlying poor convergence in extreme-condition specimens (Tests a and e). This distinction between algorithmic limitation and data scarcity motivates a three-regime comparison strategy designed to characterize mechanistic parameter uncertainty across the experimental database while maintaining transparency about where reliable inference is feasible.

The following analysis partitions the SNAP-50 UN database into three complementary regimes for detailed mechanistic comparison. Test d (Low-Burnup excluding specimen 669) represents the core calibration domain with maximum data density and strongest parameter identification, establishing baseline mechanistic parameter values where experimental constraints are sufficient to overcome the challenges of 14-dimensional inference. Test a (High-Temperature specimen 220 alone) isolates extreme-temperature behavior, where reduced experimental data and potentially different physical mechanisms lead to non-convergence even with informed priors; comparison with Test d reveals whether mechanistic model structure exhibits temperature-dependent transitions. Test e (High-Burnup specimen 665 alone) isolates extreme-burnup conditions, similarly enabling diagnosis of whether high-fission-gas-inventory regimes require distinct mechanistic treatments or merely reflect data-limited parameter non-identifiability. Through systematic multi-regime comparison, Section 7.7 will establish whether observed convergence differences reflect genuine regime-dependent physics or primarily reflect asymmetric experimental data distribution across the operating envelope.

7.7 Posterior Parameter Analysis and Regime-Dependent Model Validation

The prior refinement and convergence analysis presented in Section 7.6 established that mechanistically-informed priors enable rigorous MCMC inference in well-constrained experimental regimes while revealing fundamental data insufficiency in extreme-condition specimens. This section presents posterior parameter distributions across all experimental regimes, characterizes the data sufficiency landscape, and provides mechanistic interpretation of calibrated parameter values in the validated low-temperature, low-burnup domain. The analysis establishes low-temperature, low-burnup (Test d) as the calibration baseline where posterior parameter estimates merit high confidence, while quantifying the data requirements needed to extend calibration to extreme temperatures and burnups.

7.7.1 Multi-Regime Posterior Comparison: Revealing Data Sufficiency Across Operating Conditions

7.7.1.1 Individual Regime Posteriors

Figures 7.3, 7.4, and 7.5 display the marginal posterior distributions for all 14 mechanistic parameters in each regime individually, providing detailed visualization of parameter identifiability as a function of experimental data availability.




Figure 7.3: Posterior distributions for all 14 mechanistic parameters in low-temperature, low-burnup regime. Blue shaded regions indicate posterior probability mass, red dashed lines show prior specification (uniform or informed beta/normal).

Low-temperature, low-burnup (Figure 7.3) posteriors exhibit the gold standard for Bayesian calibration for most parameters. Tight, unimodal marginal distributions concentrated well away from prior bounds demonstrate successful parameter refinement through data. Visually inspecting distributions, nearly all parameters show posterior modes clearly

distinct from prior specification (indicated by red dashed reference lines), indicating that experimental data have provided substantial information to refine beliefs from initial priors. Broader unconstrained posterior tails are expected with sufficient experimental data in a complex 14-parameter nonlinear system. This visual pattern confirms the quantitative convergence metrics (all $\hat{R} = 1.00$, median ESS = 13,500) and reveals that mechanistic model structure successfully captures UN physics in low-temperature, low-burnup conditions.




Figure 7.4: Posterior distributions for all 14 mechanistic parameters in high-temperature regime. Blue shaded regions indicate posterior probability mass, red dashed lines show prior specification (uniform or informed beta/normal).

High-temperature regime (Figure 7.4) reveals the identifiability challenge posed by single-specimen regimes. Despite informed priors narrowing prior support relative to uniform specifications, posteriors tend not to diverge considerably from the priors or they diverge toward the maximum boundaries. Notably, many parameters show tight identification despite single-specimen data, suggesting high-temperature behavior produces sharper experimental signatures. This indicates that diverse parameter values fit the 6

observations in specimen 220 approximately equally well in some cases, while others achieve good identification. This visual pattern confirms the convergence analysis from Section 7.6. The mixed identifiability pattern demonstrates that specimen 220 data contain region-specific mechanistic signals, even with limited observations.




Figure 7.5: Posterior distributions for all 14 mechanistic parameters in high-burnup regime. Blue shaded regions indicate posterior probability mass, red dashed lines show prior specification (uniform or informed beta/normal).

High-burnup regime (Figure 7.5) shows considerably better identifiability than Test a, with many posteriors achieving meaningful tightening from priors. However, the posterior distributions are noticeably broader than Test d, indicating that the 6 high-burnup observations from specimen 665 provide useful but incomplete constraint on mechanistic parameter values. This visual improvement relative to prior expectations is consistent with updated quantitative metrics: dislocation density achieves $ESS = 16,007$ (substantially improved), convergence is successful across all parameters ($\hat{R} = 1.00$), and most parameters show information gain between 0.01–0.30. The visible improvement in posterior tightness

versus prior specification constitutes evidence that high-burnup data contain mechanistically informative signal, though larger specimen populations would improve calibration quality.

7.7.1.2 Direct Visual Evidence of Regime Dependence




Figure 7.6: Posterior distributions for all 14 mechanistic parameters overlaid to directly compare regime-dependent identifiability: high-temperature 220 (blue), low-temperature and low-burnup (orange), high-burnup 665 (light blue). The overlay visualization provides immediate visual assessment of how experimental data availability influences parameter identification.

Figure 7.6 superimposes posterior distributions from all three regimes, providing the central visual evidence for regime-dependent identifiability. The contrast is clear: Test d (orange) distributions are tightly concentrated with clear refinement from priors, Test a (blue) show mixed patterns—some posteriors refined substantially while others remain near boundaries—and Test e (light blue) exhibit posteriors broader than Test d but noticeably

tighter than Test a for many parameters. This visual progression indicates that experimental regime complexity and data density interact to control posterior quality, with neither regime achieving Test d quality. This visual progression directly parallels the quantitative metrics from Table 7.2 and Table 7.3, providing intuitive confirmation that experimental data complexity controls posterior quality.

7.7.1.3 Posterior Statistics and Quantitative Regime Comparison

To quantify the visual regime distinctions, Table 7.2 presents complete posterior statistics for all 14 parameters in Test d, while Table 7.3 directly compares six mechanistically critical parameters across the three regimes. In Table 7.2 HDI is the 94% highest density interval, Info Gain is variance reduction (0 = no constraint, 1 = complete identification) and ESS_bulk represents effective sample size with $\text{ESS} > 1,500$ indicates adequate mixing. All $\hat{R} = 1.00$ confirms convergence across 8 independent chains.

Table 7.2: Marginal posterior statistics for all 14 mechanistic parameters in low-temperature and low-burnup (14 observations). All parameters reported in normalized space $\xi \in [0, 1]$.

Parameter	Mean	SD	HDI Low	HDI High	ESS	Info Gain	$\hat{R}$
Grain Radius	0.680	0.239	0.232	1.000	16671	0.315	1.00
Frac. Yield	0.345	0.129	0.100	0.588	13753	0.260	1.00
Dislo. Density	0.692	0.115	0.460	0.881	5545	0.255	1.00
Saturation Cov.	0.438	0.255	0.002	0.885	10082	0.220	1.00
Vac. GB Diff.	0.519	0.256	0.118	1.000	9401	0.214	1.00
Dislo. Trap Z	0.572	0.268	0.121	1.000	9628	0.138	1.00
Trap. (Line)	0.182	0.134	0.000	0.426	11994	0.095	1.00
Surface Energy	0.594	0.143	0.329	0.867	23555	0.091	1.00
Nucl. Rate	0.513	0.277	0.077	0.999	17985	0.079	1.00
Resol. (Dislo)	0.504	0.225	0.113	0.891	12139	0.070	1.00
Resolution	0.570	0.229	0.134	0.939	8854	0.036	1.00
Trap. (Dislo)	0.165	0.139	0.000	0.422	17238	0.026	1.00
IG Diff. Scale	0.588	0.236	0.186	0.980	15805	-0.023	1.00
Trapping	0.182	0.144	0.000	0.445	16834	-0.045	1.00

The numerical contrast between regimes reveals important mechanistic patterns. Test d shows moderate posterior concentration across most parameters, with information gain concentrated in material property parameters (fractional yield 0.260, grain radius 0.315) rather than diffusion parameters. Test a exhibits notably higher posterior fractional yield (0.543 vs 0.345), suggesting that high-temperature regimes maintain material strength better than low-temperature conditions—a mechanistically significant finding indicating temperature-dependent yield behavior. Test e shows dislocation density comparable to Test a (0.378 vs 0.404), approximately 45% lower than Test d—revealing that high-burnup conditions produce different defect dynamics than low-temperature, low-burnup irradiation.

Table 7.3: Posterior means for six mechanistically critical parameters across three calibration regimes: low-temperature and low-burnup (14 observations), high-temperature 220 (6 observations), high-burnup 665 (6 observations).

Parameter	Test d	Test a	Test e	Δ a	Δ e
Dislo. Density	0.692 ± 0.115	0.404	0.378	0.288	0.314
Frac. Yield	0.345 ± 0.129	0.543	0.454	0.198	0.109
Trapping	0.182 ± 0.144	0.109	0.169	0.073	0.013
Resolution	0.570 ± 0.229	0.688	0.602	0.118	0.032
Dislo. Trap Z	0.572 ± 0.268	0.520	0.395	0.052	0.177
Surface Energy	0.594 ± 0.143	0.514	0.541	0.080	0.053

The cross-regime comparison (Table 7.3) reveals critical regime-dependent patterns. Fractional yield actually increases in both extreme regimes ($0.345 \rightarrow 0.543$ in Test a, $0.345 \rightarrow 0.454$ in Test e), contrary to expectations of monotonic radiation hardening. This suggests that high-temperature and high-burnup conditions may activate different creep or strength-reduction mechanisms than low-burnup environments. Dislocation density shows systematic reduction in anomalous regimes ($0.692 \rightarrow 0.404$ and 0.378), potentially reflecting temperature-accelerated defect recovery or burnup-driven microstructural evolution. These divergences, combined with broad posteriors in Tests a and e, indicate genuine regime-dependent mechanistic effects superimposed on data insufficiency.

7.7.1.4 Data Insufficiency Assessment and Experimental Implications

The convergence and identifiability analysis enables quantitative assessment of experimental data adequacy. Effective sample size provides a proxy for experimental information content where ESS reflects the number of independent experimental observations effectively constraining parameter distributions. Test d achieves median ESS = 13,500 across 14 experimental observations. Test e, with only 6 raw observations, achieves median ESS = 16,400—comparable to Test d for well-identified parameters—indicating that despite single-specimen limitation, high-burnup conditions produce experimentally informative behavior that constrains parameters effectively.

However, Test d’s larger observation set (14 vs 6) and regime-independent design provide more robust multi-parametric constraints. To extend calibration beyond current coverage, the regime-based analysis from Chapter 5 indicates that approximately 10–15 additional specimens strategically targeting the transition zone (1400–1500 K, 0.5–2 FIMA) would enable resolution of current mechanistic degeneracies while remaining within the standard regime where model physics are continuous. The current SNAP-50 database provides solid calibration foundation in the low-burnup standard regime, but the identified two-stage bottleneck structure (excessive intragranular bubble trapping followed by insufficient matrix diffusion to grain boundaries) requires experimental constraints on temperature-dependent diffusion enhancement factors to resolve parameter compensation between resolution and transport rates.

7.7.2 Mechanistic Parameter Interpretation for Primary Calibration Domain

The posterior statistics in Test d quantify the UN mechanistic model's behavior when informed by comprehensive experimental data. Presenting results in physical space (rather than normalized coordinates) enables direct comparison with literature values, BISON input parameters, and mechanistic expectations. This subsection provides detailed interpretation of key posterior parameter values in their native units, connecting numerical results to physical understanding of fission gas release and swelling phenomena. The interpretation is grounded in the two-stage bottleneck structure identified in Chapter 5: Stage 1 excessive intragranular bubble trapping preventing gas from reaching the matrix, and Stage 2 insufficient matrix diffusion to grain boundaries even when gas escapes intragranular bubbles.

Table 7.4: Marginal posterior statistics for all 14 mechanistic parameters in low-temperature and low-burnup (14 observations) in physical space.

Parameter	Mean	SD	HDI Low	HDI High	ESS	Info Gain
Grain Radius [μm]	37.2	9.56	19.3	50.0	16671	0.315
Frac. Yield	0.219	0.026	0.170	0.268	13753	0.260
Dislo. Density [10^{14} m^{-2}]	11.32	6.28	3.89	27.04	5545	0.255
Saturation Cov.	0.353	0.089	0.201	0.510	10082	0.220
Vac. GB Diff. [10^6 multiplier]	11.91	62.38	0.30	1000	9401	0.214
Dislo. Trap Z	44.3	16.1	17.3	70.0	9628	0.138
Trap. (Line) [scale factor]	1.90	1.33	0.10	4.32	11994	0.095
Surface Energy [J/m^2]	1.138	0.057	1.032	1.247	23555	0.091
Nucl. Rate [scale factor]	5.18	2.74	0.86	9.99	17985	0.079
Resol. (Dislo) [scale factor]	5.09	2.23	1.22	8.92	12139	0.070
Resolution [scale factor]	5.74	2.27	1.43	9.40	8854	0.036
Trap. (Dislo) [scale factor]	1.73	1.38	0.10	4.28	17238	0.026
IG Diff. Scale [scale factor]	5.92	2.34	1.94	9.80	15805	-0.023
Trapping [scale factor]	1.90	1.43	0.10	4.51	16834	-0.045

7.7.2.1 Grain Radius and Dislocation Density

Grain radius (posterior mean $37.2 \pm 9.56 \mu\text{m}$, info gain = 0.315) characterizes the fuel pellet microstructure controlling fission gas diffusion pathways. The posterior mean of $37.2 \mu\text{m}$ represents moderately coarse grains within the reported SNAP50 grain radius range. The tight 94% credible interval (19.3–50.0 μm) reflects substantial experimental constraint. This grain size is consistent with sintered UN fuel after thermal processing but before irradiation-induced coarsening. In the BISON mechanistic model, grain radius enters directly into the grain-boundary diffusion layer calculations and controls the characteristic length scale for diffusion eigenvalue calculations $\lambda_n = (n\pi/a)^2$, where a is the grain radius. Larger grains therefore exponentially reduce diffusion-driven release rates, enhancing gas retention.

Dislocation density (posterior mean $11.3 \times 10^{14} \text{ m}^{-2}$, SD = $6.28 \times 10^{14} \text{ m}^{-2}$, info gain = 0.255) reflects substantial but sub-saturated defect accumulation under low-temperature, low-burnup irradiation. The posterior concentrates well below the prior upper bound (10^{15} m^{-2}), with the center value just lower than the constant default value of $4 \times 10^{13} \text{ m}^{-2}$. The Chapter 5 parametric study established that dislocation density acts primarily through nucleation site supply, reducing dislocation density shifts gas accumulation from dislocation bubbles into intragranular bubbles rather than enabling grain boundary transport, confirming that the bottleneck is not nucleation availability but rather the downstream trapping mechanisms that prevent gas from reaching grain boundaries. In the UN SIFGRS implementation, dislocation density couples directly to two mechanisms: (1) nucleation of bubbles at dislocation sites, and (2) trapping rates which scale logarithmically with dislocation spacing. The tight identification (information gain = 0.255) establishes dislocation density as the third-most-constrained parameter. This tight constraint reflects that swelling observations discriminate between regimes with abundant defect populations and depleted regimes, not that the MCMC has uniquely resolved how dislocations partition gas between competing sink populations. At $11.3 \times 10^{14} \text{ m}^{-2}$, the posterior density is high enough to provide abundant nucleation sites and sufficient trapping capacity to meaningfully affect swelling predictions and allow fission gas release.

7.7.2.2 Fractional Yield

Fractional yield represents the molar yield of fission gas atoms (primarily Xenon and Krypton) produced per fission event. The posterior mean of 0.219 (SD = 0.026, info gain = 0.260) indicates that UN produces approximately 0.219 moles of fission gas atoms per mole of fission events. The tight posterior (credible interval 0.170–0.268) reflects strong experimental constraint on gas production rates from nuclear data and experimental calibration.

Mechanistically, the tight constraint on fractional yield validates the historical data estimating the fractional yield is ≈ 0.25 for Xe and Kr. Importantly, the tight identification (info gain = 0.260) is independent of the two-stage bottleneck structure identified in Chapter 5. The posterior constrains the absolute gas inventory produced, while the bottleneck mechanisms govern how that gas is partitioned among microstructural locations. The swelling and fission gas release observations are sensitive enough to distinguish between different absolute gas production rates. This tight identification of yield, compared to the more diffuse posterior distributions of transport parameters, indicates that the total gas inventory generated is experimentally well-constrained by the combination of volumetric swelling and release measurements. The posterior value confirms that UN behaves as anticipated under neutron fission in the SNAP-50 low-temperature, low-burnup regime.

7.7.2.3 Surface Energy and Saturation Coverage

Surface energy (posterior mean 1.138 J/m^2 , SD = 0.057, info gain = 0.091) controls thermodynamic driving forces for bubble nucleation and equilibrium bubble size. The

posterior concentration above the prior lower bound (0.9 J/m^2) but well below the upper bound (1.3 J/m^2) reflects that UN surface energy is intermediate in the literature range. In the code implementation, surface energy appears explicitly in the bubble pressure calculations controlling the nucleation barrier height and equilibrium bubble size, and implicitly in the dislocation punching criterion.

Saturation coverage (posterior mean 0.353, $\text{SD} = 0.089$, info gain = 0.220) represents the fraction of available grain-boundary perimeter that can accommodate bubbles before transition to grain-boundary coverage and release. In the BISON grain-boundary model, saturation coverage controls the maximum bubble surface density before interconnection becomes energetically favorable. The posterior concentration at approximately one-third saturation indicates that approximately 35% of potential grain-boundary bubble surface sites are filled before the critical transition occurs. This intermediate value is mechanistically significant, and directly validated by the Chapter 5 parametric study. Reducing saturation coverage from 0.5 to 0.25 advanced fission gas release onset at low temperatures but did not produce sustained release, confirming that the coverage threshold is not the limiting mechanism. Instead, Chapter 5 gas partitioning diagnosis showed that grain boundary gas inventory remained below 3% of total across all specimens, indicating that the bottleneck is insufficient gas supply to grain boundaries rather than insufficient interconnection efficiency.

7.7.2.4 Scale Factors and Their Physical Roles

Scale factors in the UNSIFGRS implementation act as dimensionless multipliers applied to baseline physical rates and diffusion coefficients.

Intragranular diffusion scale (posterior mean 5.92, $\text{SD} = 2.34$, info gain = -0.023) reveals weak or zero experimental constraint on bulk volume diffusion enhancement. The negative information gain indicates the posterior distribution remains broader than the prior, demonstrating that swelling observations do not uniquely constrain intragranular diffusion rates. This negative information gain is mechanistically informative. Chapter 5 parametric study demonstrated that when baseline trapping mechanisms are active, increasing matrix diffusion produces negligible improvement because gas never escapes intragranular bubbles to populate the matrix in the first place. Only when trapping is aggressively reduced does the matrix become the active transport pathway, at which point diffusion rate becomes important. The swelling data in Test d do not sample this regime, so the MCMC cannot uniquely constrain matrix diffusion from observations that remain insensitive to it. Physically, this weak identification reflects that low-temperature UN operates in a grain-boundary-limited regime where gas atoms produced deep in grain interiors reach grain boundaries through slow volume diffusion, but once at grain boundaries, rapid GB diffusion dominates release kinetics. The code structure confirms this interpretation. The effective diffusion coefficient accounts for both bubble trapping and resolution rates, but the constraint from grain-boundary percolation observations dominates identification.

Vacancy-grain boundary diffusion (posterior mean 11.9×10^6 multiplier, info gain = 0.214) quantifies grain-boundary transport enhancement through vacancy-mediated mechanisms.

In the UN SIFGRS code, this multiplier scales the baseline vacancy diffusivity function before entering into grain-boundary bubble evolution calculations. The information gain of 0.214 indicates that swelling observations provide meaningful but partial constraint on GB transport. Chapter 5 investigation with 10^9 multiplier enhancement of GB vacancy diffusion showed only modest improvements in fission gas release, confirming that accommodation physics (vacancy supply to grain boundaries) is not the bottleneck. Instead, the upstream Stage 1 trapping bottleneck prevents gas from reaching grain boundaries in quantities sufficient to test accommodation limits. The posterior value of 11.9×10^6 represents a moderate enhancement that ensures grain boundary saturation accumulates over the irradiation duration while avoiding extreme parameter values that would be unphysical. Faster GB diffusion enables more earlier onset to fission gas release and gas escape, reducing retention and swelling.

Dislocation trap Z-factor (posterior mean 44.3, SD = 16.1, range 17.3–70.0, info gain = 0.138) represents the effective trap stability parameter in the dislocation core trapping calculation. The posterior value of 44.3 indicates moderate trap strength at these conditions where the default is 25. The information gain of 0.138 reflects that trapping and resolution rates are mechanistically entangled through the equilibrium gas fraction calculation described below, making independent identification difficult without additional experimental constraints on bubble kinetics.

7.7.2.5 Trapping and Resolution Scale Factors

The three trapping scale factors (overall 1.90, trap-dislocations 1.73, trap-line 1.90) cluster in the lower-middle portion of their [0.1, 10.0] range, indicating substantially suppressed gas retention relative to baseline model predictions. In the code implementation, these factors multiply the base trapping rate. The suppressed posterior values indicate that fewer gas atoms successfully lodge at trapping sites than baseline assumptions predict.

The weak identification (information gain 0.026–0.095) reflects fundamental parameter compensation in the UNSIFGRS architecture. The code computes equilibrium gas partitioning such that the equilibrium distribution depends on the ratio of trapping to resolution rates, not on their absolute values. Consequently, reducing trapping by a factor of 2 while increasing resolution by a factor of 2 produces identical predictions. The Chapter 5 parametric study demonstrated this explicitly. The aggressive combined optimization achieved success only when trapping was reduced to $0.1 \times$ and resolution was simultaneously increased to $10 \times$, a synergistic adjustment that would be invisible to a MCMC exploring individual parameter variations. The weak posterior identifications here reflect that the MCMC explores a one-dimensional compensation manifold within the higher-dimensional parameter space where multiple parameter combinations produce equivalent likelihoods. Without independent experimental data on trapping mechanisms, the MCMC cannot uniquely distinguish these competing effects. Multiple parameter combinations, weak trapping with moderate resolution, or strong trapping with strong resolution, fit swelling and release data approximately equally well.

7.7.2.6 Resolution Scale Factors

Resolution scale factors (overall 5.74, dislocation-specific 5.09) appear in the bubble concentration evolution calculations. In the code, these factors multiply the baseline resolution rate, which empirically depends on bubble radius and fission rate. The posterior values of approximately $5.7\times$ indicate enhanced bubble coalescence relative to baseline model predictions.

The weak information gain (0.036 and 0.070) reveals that coalescence mechanisms are mechanistically secondary in Test d conditions. Chapter 5 analysis established that in the grain-boundary-limited regime, once grain boundary saturation is achieved, the absolute resolution rate (whether $1\times$ or $10\times$ baseline) minimally affects total release because gas reaching grain boundaries is vented regardless of whether intragranular bubbles are destroyed slowly or rapidly. The posterior concentration at $5.74\times$ likely represents a saturation point where further increases provide negligible modeling improvement. Once grain-boundary diffusion is specified, detailed bubble coalescence kinetics matter less for macroscopic release predictions because grain-boundary sinks efficiently absorb escaped gas regardless of coalescence pathway. The code structure confirms this such that bubble coalescence affects bubble concentration and radius evolution, but final release is determined by whether gas reaches the grain-boundary-to-surface percolation threshold in the thermal release function.

7.7.2.7 Nucleation Rate and Defect-Limited Regimes

Nucleation rate scale factor (posterior mean 5.18, SD = 2.74, info gain = 0.079) indicates moderate enhancement of bubble nucleation kinetics. The weak identification reflects that UN fuel operates in a defect-limited regime. Available dislocation sites vastly exceed the rate at which fission gas atoms become available for bubble formation. Nucleation rate couples to gas concentration via homogeneous nucleation theory, and the moderate scale factor ($5.18\times$) indicates that baseline predictions already assume efficient nucleation. The Chapter 5 parametric investigations did not target nucleation rate directly, but the gas partitioning diagnosis showed that bubbles nucleate readily at early burnup; the limitation emerges only later as re-trapping during diffusion prevents gas from escaping the intragranular bubble population to feed grain boundaries. Increasing this factor further provides minimal benefit because the system is limited by gas atom supply, not by nucleation site availability.

7.7.2.8 Summary: Parameter Landscape and Mechanistic Pathways

The posterior parameter landscape reveals hierarchical mechanistic importance reflecting UN fuel physics:

Tier 1 (Info Gain 0.20–0.32): Dominant Material Properties—Tightly Constrained and Mechanistically Validated

Grain radius (0.315), fractional yield (0.260), and dislocation density (0.255) emerge as the three most tightly constrained parameters. These characterize irreversible material degradation, that directly govern swelling and fission gas mechanisms mechanisms in the

BISON finite-element context. These Tier 1 parameters are sufficiently well-constrained that adjusting their baseline model values toward the posterior means would represent concrete, defensible improvements to model predictions. Specifically: (1) increasing the default grain radius from 25 μm to 37 μm would reflect measured SNAP-50 specimen microstructure; (2) confirming fractional yield at 0.219 (versus historical estimates of 0.25) would refine fission gas production calculations; (3) employing the dislocation density value of $11.3 \times 10^{14} \text{ m}^{-2}$ as a temperature dependent or irradiation dependent material property rather than a constant would capture radiation-induced defect accumulation. These changes represent direct mechanistic improvements rather than parameter compensation.

Tier 2 (Info Gain 0.09–0.21): Intermediate Gas Physics—Moderately Constrained but Model-Limited

Saturation coverage (0.220), vacancy GB diffusion (0.214), surface energy (0.091), and dislocation trap Z (0.138) control intermediate-scale fission gas behavior. These appear in rate calculations and equilibrium relationships, but multiple parameter combinations produce similar predictions due to compensatory effects in the rate equations. The moderate identifications here reflect that grain-boundary-related physics matter for swelling, but the Stage 1 bottleneck (excessive intragranular trapping) limits the grain boundary gas inventory so severely that detailed grain boundary mechanics remain under-constrained by observations. These Tier 2 parameters show sufficient identification to warrant investigation: saturation coverage (0.220 info gain) and vacancy GB diffusion (0.214) are informationally competitive with the dislocation density constraint, suggesting they do carry signal despite the bottleneck limitation. However, adjusting these parameters in the baseline model would represent potential symptom treatment rather than mechanistic repair. The Chapter 5 analysis demonstrated that reducing saturation coverage or enhancing GB diffusion allows earlier onset of fission gas release. Nevertheless, these parameters merit investigation in a refined model as they may become rate-limiting in that regime.

Tier 3 (Info Gain < 0.08, including negative): Matrix Transport and Bubble Dynamics—Weakly Constrained and Model-Degenerate

Resolution parameters (0.036–0.070), trapping multipliers (0.026–0.095), nucleation rate (0.079), and IG diffusion (-0.023) show minimal identification. These weak identifications directly correspond to the mechanistic bottleneck structure from Chapter 5. The Stage 1 bottleneck (excessive trapping prevents matrix accumulation) makes intragranular diffusion unimportant and bubble coalescence secondary. Without upstream gas supply, downstream transport mechanics cannot be observed. The MCMC explores a compensation manifold where varying these Tier 3 parameters produces minimal change in predicted swelling or release because the Stage 1 bottleneck dominates all regime. Tier 3 parameters are not sufficiently constrained to warrant baseline model adjustment. They remain degenerate even after calibration to 14 experimental observations. The negative information gain for IG diffusion explicitly indicates that the posterior is less informative than the prior, confirming that swelling and fission gas release data provide no direct signal about matrix diffusion efficiency. Adjusting these parameters would introduce additional uncertainty rather than reduce it.

Mechanistic Implications and Prioritized Investigation Path

The hierarchical structure reveals that swelling behavior in Test d is driven primarily by material property rather than by detailed gas transport or retention mechanisms. Gas bubble formation proceeds relatively deterministically once microstructural evolution is specified. The Tier 1 parameters, now validated through Bayesian calibration, provide a confirmed baseline for forward predictions. The Tier 2 parameters warrant additional investigation only within a mechanistically reformed model. The Tier 3 parameters should not be adjusted from their present values without additional experimental data on bubble kinetics, because any adjustment would represent parameter compensation with no mechanistic justification.

7.7.3 Validation and Forward Pathway

The posterior parameter distributions established above provide the foundation for forward uncertainty quantification and model validation. The Test d calibration establishes quantitative confidence in Tier 1 parameters (grain radius, fractional yield, dislocation density; information gain 0.25–0.32) sufficient to recommend baseline model updates within the validated envelope spanning 1269–1503 K and 0.54–1.08 FIMA burnup. Tier 2 and Tier 3 parameters warrant different treatment. Tier 2 remain under-constrained due to the Stage 1 bottleneck limiting grain boundary gas supply; they should not be adjusted without mechanistic reform to intragranular trapping mechanisms. Tier 3 parameters show minimal to negative information gain, indicating that swelling observations provide no unique constraint on their values. Adjusting these parameters without independent experimental data on bubble kinetics would introduce additional uncertainty rather than reduce it.

Beyond the validated envelope, in particular at extreme temperature (Test a, Specimen 220, 1561–1714 K) and high-burnup (Test e, Specimen 665, 4.58 FIMA) regimes, convergence diagnostics degrade substantially, revealing fundamental regime-dependent differences in mechanistic model structure. Extrapolation beyond the validated envelope requires explicit acknowledgment of model-inadequacy uncertainty accounting for structural limitations beyond parametric uncertainty.

The identified two-stage bottleneck structure, excessive intragranular bubble trapping followed by insufficient matrix diffusion to grain boundaries, represents the primary mechanistic limitation preventing accurate fission gas release predictions across all temperatures. Forward validation will quantify whether this bottleneck produces predictable, temperature-dependent biases that scale systematically with operating conditions, or whether regime-dependent mechanisms introduce structural model inadequacy beyond parameter compensation. This distinction is critical for accelerated fuel qualification: if bottleneck effects scale predictably, engineering safety margins can be established through uncertainty quantification. If structural inadequacy dominates, additional experimental data or computational model reform would be necessary before the model provides adequate confidence for qualification.

The subsequent chapter develops this forward analysis rigorously, propagating posterior parameter distributions through BISON forward simulations, comparing posterior-predicted

swelling and fission gas release against experimental observations at domain boundaries, and decomposing predictive uncertainty into parametric and structural components. This forward framework quantifies both what the model predicts and the confidence warranted by available information, completing the cycle from mechanistic understanding through validated predictions suitable for engineering decision-making under uncertainty.

Chapter 8

Posterior Prediction and Forward Uncertainty Quantification

8.1 Introduction

Following Bayesian calibration via Markov chain Monte Carlo, the resulting posterior distributions over model parameters enable rigorous forward uncertainty quantification and validation of the calibrated model. While the calibration process constrains parameter uncertainties given observed experimental data, the posterior samples must be propagated through the model to assess predictive performance, quantify remaining uncertainties, and validate the calibrated model against both calibration data and held-out test cases. This chapter presents a comprehensive post-calibration analysis framework encompassing posterior predictive distributions, uncertainty decomposition, parameter identification assessment, and validation against both surrogate-based predictions and direct BISON simulations.

The posterior predictive distribution represents the Bayesian framework for forward prediction under parameter uncertainty. Given a set of posterior samples $\{\boldsymbol{\theta}^{(i)}\}_{i=1}^N$ drawn from the calibrated posterior $p(\boldsymbol{\theta} \mid \mathcal{D})$, predictions for a new observation $\tilde{y}$ at input conditions $\tilde{\mathbf{x}}$ incorporate both parameter uncertainty and observation noise. This predictive distribution accounts for epistemic uncertainty arising from incomplete knowledge of model parameters while maintaining consistency with observed experimental evidence. The resulting predictions provide honest uncertainty quantification that can be validated through coverage analysis and comparison to experimental observations.

Uncertainty in model predictions arises from multiple sources that must be carefully decomposed and quantified. Parameter uncertainty reflects the remaining epistemic uncertainty in model parameters after calibration. Surrogate model uncertainty captures discrepancies between the Gaussian process emulator used during calibration and the true BISON model. Experimental measurement uncertainty represents irreducible aleatory variability in observations. Model form uncertainty acknowledges structural inadequacies in the underlying physics model. Proper accounting of these uncertainty sources enables

assessment of whether prediction intervals honestly reflect the true uncertainty, identification of dominant uncertainty contributors, and targeted uncertainty reduction strategies.

Validation of the calibrated model proceeds through multiple complementary analyses. Posterior predictions are first generated for the calibration data itself to verify that the model adequately represents the training observations. These predictions are then compared against independent BISON simulations using the calibrated parameters to validate the surrogate model fidelity. Comparison between predictions using uninformed prior samples versus calibrated posterior samples quantifies the information gain achieved through calibration. Finally, systematic analysis of prediction accuracy across different operating regimes identifies model validity boundaries and conditions where the physics model becomes inadequate.

Parameter identification analysis reveals which model parameters are effectively constrained by the available experimental data. Information gain metrics quantify the reduction in parameter uncertainty from prior to posterior for each parameter. Parameters with high information gain are well-identified by the data, while those with low information gain remain poorly constrained despite calibration. The choice of point estimate for calibrated parameters depends on the posterior distribution shape. For symmetric posteriors, the mean provides an optimal estimate under squared error loss. For skewed posteriors arising from bounded or informative priors, the mode represents the maximum a posteriori estimate and better reflects the most probable parameter value. The median offers a robust alternative less sensitive to posterior distribution tails.

8.2 Methods

8.2.1 Validation Framework and Success Criteria

The Bayesian calibration process yielded posterior parameter samples from 8 independent MCMC chains with 2500 samples per chain. For computational efficiency in forward predictions, 700 samples were selected for posterior prediction analysis, representing sufficient coverage of the posterior distribution. Validation of this calibration requires assessment across three complementary dimensions. First, posterior predictions must exhibit appropriate coverage of observations, with approximately 95% of observations falling within two-sigma credible intervals and 68% within one-sigma intervals. Second, posterior predictions must demonstrate improved accuracy relative to uncalibrated prior predictions, quantified through reduction in prediction uncertainty and absolute error metrics. Third, the analysis must quantify parameter uncertainty remaining after calibration and partition prediction variance into distinct sources (parameter, surrogate, and experimental), revealing which mechanisms dominate prediction bounds.

The validation strategy employs multiple complementary approaches. Direct BISON simulations using the 700 calibrated posterior parameter samples validate that Gaussian process surrogates accurately represent full-model behavior for parameter values within

the posterior support. Validation proceeds on both calibration experiments (in-sample validation to verify surrogate fidelity) and held-out experiments not included in MCMC training (out-of-sample validation to assess predictive generalization). Comparison between prior predictions using uninformed Latin hypercube sampling with 2100 samples and posterior predictions using the 700 calibrated MCMC samples quantifies the information gain from Bayesian calibration. Uncertainty decomposition partitions prediction variance into parameter uncertainty, surrogate approximation error, and experimental measurement noise, revealing dominant contributors to prediction bounds and informing resource allocation for uncertainty reduction.

8.2.2 Posterior Predictive Distributions

For a prediction location $\tilde{\mathbf{x}}$ with observation $\tilde{y}$, the posterior predictive distribution marginalizes over posterior parameter uncertainty to incorporate both parameter constraints and irreducible measurement noise

$$p(\tilde{y}|\tilde{\mathbf{x}}, \mathcal{D}) = \int p(\tilde{y}|\tilde{\mathbf{x}}, \boldsymbol{\theta}) p(\boldsymbol{\theta}|\mathcal{D}) d\boldsymbol{\theta} \quad (8.1)$$

In practice, this is approximated using the 700 posterior samples $\{\boldsymbol{\theta}^{(i)}\}_{i=1}^{700}$ from Chapter 7 by evaluating the model at each sample. The resulting ensemble $\{\tilde{y}^{(i)}\}_{i=1}^{700}$ provides a Monte Carlo approximation to the posterior predictive distribution.

Credible intervals are constructed from the empirical quantiles of posterior predictive samples. The $100(1 - \alpha)\%$ credible interval is

$$\text{CI}_{1-\alpha} = [Q_{\alpha/2}\{\tilde{y}^{(i)}\}, Q_{1-\alpha/2}\{\tilde{y}^{(i)}\}] \quad (8.2)$$

where Q_p denotes the p -th quantile. For well-calibrated uncertainty quantification, approximately $100(1 - \alpha)\%$ of observations should fall within their corresponding credible intervals. The 68% credible interval (one-sigma) and 95% credible interval (two-sigma) are the primary metrics reported.

8.2.3 Uncertainty Decomposition

Prediction uncertainty arises from independent sources that must be decomposed to identify dominant contributors. The total predictive variance partitions into parameter uncertainty, surrogate approximation uncertainty, and experimental measurement uncertainty

$$\text{Var}[\tilde{y}] = \text{Var}_{\text{param}} + \text{Var}_{\text{surr}} + \text{Var}_{\text{exp}} \quad (8.3)$$

Parameter uncertainty is quantified from the spread in predictions across posterior samples. Surrogate model uncertainty reflects discrepancies between Gaussian process emulators and the true BISON model, quantified through calibration inflation factors determined during surrogate model development in Chapter 7. Experimental measurement uncertainty represents irreducible aleatory variability in observations, estimated as a calibrated hyperparameter during MCMC sampling.

8.2.4 Prior Versus Posterior Comparison

To quantify the benefit of Bayesian calibration, posterior predictions using calibrated posterior distributions are compared against predictions using uninformed prior distributions via Latin hypercube sampling over the uncalibrated prior parameter space. The LHS design evaluated 2100 samples across the prior parameter space, representing the comprehensive prior exploration performed during surrogate model development. While the LHS sample size exceeds the 700 posterior samples, this comparison effectively demonstrates the information gain from calibration by quantifying how much the posterior distributes parameter space relative to uniform prior sampling.

Comparison metrics include reduction in predictive standard deviation

$$\text{Std Reduction} = \frac{\sigma_{\text{prior}} - \sigma_{\text{posterior}}}{\sigma_{\text{prior}}} \times 100\% \quad (8.4)$$

and improvement in prediction accuracy measured through mean absolute error. These metrics quantify how much parameter uncertainty has been reduced by experimental data and whether calibrated parameters provide better agreement with observations than uncalibrated values.

8.2.5 Validation Metrics

Assessment of posterior prediction quality employs metrics quantifying both accuracy and calibration of uncertainty estimates. MAE measures average magnitude of prediction errors as described by Equation 7.4 in Chapter 7. Root mean squared error (RMSE) penalizes large errors more heavily as described by:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{k=1}^n (\hat{y}_k - y_k)^2} \quad (8.5)$$

where $\hat{y}_k$ is the posterior predictive mean and y_k is the observation for location k . Coverage analysis evaluates whether credible intervals contain observations at expected rates. For $100(1 - \alpha)\%$ credible intervals, empirical coverage is:

$$\text{Coverage}_{1-\alpha} = \frac{1}{n} \sum_{k=1}^n \mathbb{I}(y_k \in \text{CI}_{1-\alpha,k}) \quad (8.6)$$

where $\mathbb{I}$ is the indicator function. For well-calibrated uncertainty quantification, empirical coverage should match nominal levels, $\text{Coverage}_{1-\alpha} \approx 1 - \alpha$. Systematic under-coverage indicates overconfident predictions, while over-coverage suggests conservative bounds.

These metrics are computed separately for swelling and fission gas release predictions across all test configurations and compared between BISON and surrogate-based predictions. Consistent achievement of nominal coverage rates while maintaining low MAE and RMSE demonstrates both accurate and well-calibrated predictions.

8.3 Results

The validation of Bayesian calibration proceeds through systematic comparison of posterior predictions against experimental observations across multiple validation approaches. This section reports posterior prediction results for the primary calibration regime of low-temperature and low-burnup (Test d), validates surrogate model fidelity through direct BISON simulations, quantifies information gain from calibration through prior versus posterior comparison, decomposes prediction uncertainty into constituent sources, and examines model performance across regime boundaries to identify validity limits. The results demonstrate successful calibration within the low-temperature, low-burnup domain while revealing systematic model inadequacy at extreme conditions, providing both confidence in model predictions where valid and clear guidance on model limitations for engineering applications.

8.3.1 Low-Temperature and Low-Burnup Posterior Predictions

Posterior predictions for the primary calibration regime (Test d: $T < 1485$ K, $B < 2\%$) demonstrate successful calibration and excellent agreement with experimental observations. The 700 calibrated posterior parameter samples were evaluated through surrogate models across all five experiments in the primary regime (specimens 602, 603, 658, 662, 664) with axial positions dependent on temperature or burn-up. The resulting posterior predictive distributions were compared against observations to assess coverage and prediction accuracy.

Figure 8.1 shows the full posterior predictive distributions as violin plots across the entire dataset. The distributions are concentrated tightly around observations, indicating that parameter uncertainty has been substantially reduced through calibration. The 95% credible intervals, shown as shaded regions, encompass the vast majority of experimental observations, providing evidence of well-calibrated uncertainty quantification.

Coverage analysis reveals near-nominal calibration across the primary regime. Overall, 86.4% of observations (19 of 22 data points) fell within their respective 95% credible intervals, slightly below the nominal 95% level but within expected statistical variation for the sample size. This performance is consistent with honest uncertainty quantification rather than overconfident predictions. The three coverage failures (1.2% of data) represent minor misses rather than systematic model inadequacy: experiment 603 top swelling exceeded the upper credible interval by 0.677 vol%, experiment 602 middle FGR fell below the lower bound by 0.135%, and experiment 664 bottom swelling exceeded the upper bound by 0.6 vol%. These near-miss failures are small relative to the magnitude of observations and suggest marginal prediction errors rather than fundamental model failure. Prediction accuracy measured by MAE remained small across both quantities. For swelling predictions, average MAE was 0.298 vol%, representing excellent agreement relative to observations ranging from 1.2 to 4.3 vol%. For fission gas release, average MAE was 0.05%, with FGR observations ranging from 0.01 to 0.2%. The variation in prediction accuracy across experiments reflects differences in local conditions and parameter sensitivity as experiments 658 and 662 achieved 100% coverage with the smallest errors, while experiments 602, 603, and 664 showed slightly




Figure 8.1: Posterior predictive distributions for low-temperature and low-burnup showing swelling (left panel) and fission gas release (right panel) across all calibration experiments. Violin plots represent the full distribution of 700 posterior predictions; circles mark the posterior mean; diamonds mark observations. The distributions are concentrated near observations, indicating well-calibrated predictions.

larger errors but still maintained 83% coverage. This experiment-to-experiment variation is mechanistically expected given the range of temperatures (1269–1523 K) and burnup (0.54–1.08 at%) represented in the regime.

The posterior predictive standard deviations quantify remaining parameter uncertainty after calibration. For swelling, average posterior uncertainty was 0.255 vol%, while for FGR, average posterior uncertainty was 0.0385%. These predictive standard deviations are substantially reduced relative to prior predictions, quantifying the information gained through Bayesian calibration. The tight posterior distributions evident in Figure 8.1 demonstrate that available experimental data have effectively constrained model parameters within the primary regime.

These results demonstrate that Bayesian calibration has successfully recovered parameter distributions that produce accurate, well-calibrated predictions across the low-temperature, low-burnup domain. The near-nominal coverage, small prediction errors, and tight uncertainty bounds provide confidence in model predictions for engineering applications within this regime. The few coverage failures represent statistical edge cases rather than systematic model inadequacy, confirming the mechanistic validity of the binary percolation fission gas release model within the regime where its assumptions are satisfied.

8.3.2 BISON Validation Against Direct Simulations

To validate that surrogate predictions accurately represent full BISON model behavior across the posterior distribution, the 700 calibrated posterior parameter samples were

Table 8.1: Regime posterior prediction metrics by experiment and position. Observations, posterior means, credible intervals, and coverage assessment are shown for all 22 predictions. CI coverage indicates whether observations fell within 95% credible intervals.

Exp	Pos	Quantity	Obs	Mean	95% CI	Error
602	bottom	Swelling	2.400	2.796	[2.351, 3.250]	0.396
602	middle	Swelling	3.500	3.127	[2.605, 3.636]	-0.373
602	top	Swelling	3.000	2.886	[2.304, 3.508]	-0.114
603	bottom	Swelling	2.100	2.292	[1.958, 2.700]	0.192
603	middle	Swelling	2.700	2.435	[2.078, 2.924]	-0.265
603	top	Swelling	2.900	2.223	[1.848, 2.636]	-0.677
658	top	Swelling	3.800	3.990	[3.197, 4.901]	0.190
662	top	Swelling	4.300	4.187	[3.283, 5.012]	-0.113
664	bottom	Swelling	1.200	1.800	[1.498, 2.345]	0.600
664	middle	Swelling	1.500	1.777	[1.493, 2.252]	0.277
664	top	Swelling	1.800	1.725	[1.470, 2.139]	-0.075
602	bottom	FGR	0.080	0.041	[-0.005, 0.122]	-0.039
602	middle	FGR	0.200	0.065	[0.003, 0.184]	-0.135
602	top	FGR	0.020	0.003	[-0.004, 0.078]	-0.017
603	bottom	FGR	0.060	0.031	[-0.006, 0.116]	-0.029
603	middle	FGR	0.100	0.035	[-0.009, 0.120]	-0.065
603	top	FGR	0.050	0.002	[-0.005, 0.058]	-0.048
658	top	FGR	0.100	0.141	[0.031, 0.352]	0.041
662	top	FGR	0.010	0.112	[0.006, 0.307]	0.102
664	bottom	FGR	0.010	0.039	[-0.008, 0.140]	0.029
664	middle	FGR	0.010	0.036	[-0.008, 0.113]	0.026
664	top	FGR	0.010	0.025	[-0.009, 0.091]	0.015

evaluated directly through BISON for the primary calibration regime experiments. This independent validation tests whether surrogates faithfully capture BISON physics across the parameter space explored by calibration.

Surrogate model fidelity was assessed through direct comparison between surrogate predictions and BISON predictions for identical posterior parameter samples. Across all 20 Test d predictions (5 experiments, 3 positions, 2 quantities), the mean absolute error between surrogate and BISON predictions averaged 0.231, with swelling surrogates showing MAE of 0.377 vol% and FGR surrogates showing MAE of 0.085%. These relatively small errors confirm that surrogate approximations accurately represent BISON outputs across the posterior parameter range.

Coverage analysis on direct BISON simulations reveals excellent calibration in the primary regime. BISON-based posterior predictions achieved 85.0% coverage at the 95% credible interval for Test d (17 of 20 observations within bounds), matching the surrogate-based coverage of 86.4% reported in Section 8.3.1. This consistency between surrogate

and BISON coverage validates that the surrogate calibration inflation factors appropriately account for approximation uncertainty. By quantity, swelling achieved 90% coverage in BISON predictions, while FGR achieved 80% coverage, both near nominal levels. Prediction accuracy measured by MAE remained small: 0.401 vol% for swelling and 0.053% for FGR. These metrics demonstrate that posterior parameter distributions derived from surrogate-accelerated MCMC produce valid predictions when applied to the full BISON model.

Figure 8.2 visualizes BISON predictions across all calibration and validation tests. The left panel shows swelling predictions, while the right panel shows fission gas release predictions. Violin plots display the full distribution of 700 BISON predictions; circles mark the posterior mean; hollow diamonds mark observations. The comparison reveals tight correspondence between BISON predictions and observations in the Test d primary regime, with systematic divergence evident at Test 613 middle position (marked as T greater than 1485 K) where temperature exceeds the identified regime boundary.




Figure 8.2: BISON validation of posterior predictions across low-temperature and low-burnup (primary) and Test 613 (hold-out validation). Left panel: swelling predictions across burnup range. Right panel: fission gas release predictions. Violin plots show full distribution of 700 posterior BISON predictions; circles mark means; hollow diamonds mark observations.

To validate calibration performance across the full range of observed values, Figure 8.3 presents BISON posterior predictions plotted against observations with 95% credible interval bounds. This visualization reveals systematic accuracy and absence of bias across the observation range. For swelling predictions (left panel), posterior means cluster tightly along the perfect prediction line (dashed diagonal), with credible intervals expanding appropriately to capture observation variability. No systematic bias is evident: predictions are neither consistently high nor consistently low across the swelling range from 1.2 to 4.3 vol%. The few observations falling outside credible intervals (marked by gaps between error bars and the diagonal) represent the expected coverage failures documented in Table 8.1.

Fission gas release predictions (right panel) show similar systematic accuracy despite the much smaller magnitude of FGR values (0.01–0.2%). The posterior means track

observations with appropriate uncertainty bounds, demonstrating that BISON predictions have successfully learned mechanistic behavior even for the low FGR values characteristic of the low-temperature regime. The shaded experimental uncertainty band (determined from measurement standard deviations during MCMC calibration) shows that measurement error alone cannot explain prediction discrepancies, validating that observed misses represent genuine model limitations rather than experimental noise.




Figure 8.3: BISON calibration performance showing predicted versus observed for low-temperature and low-burnup regime. Left panel: swelling predictions (posterior mean shown as squares, 95% credible intervals as error bars). Right panel: FGR predictions. Dashed diagonal line shows perfect prediction; shaded region represents experimental measurement uncertainty.

Systematic comparison between Test d surrogate predictions and BISON predictions reveals no pattern of systematic bias or directional error. Table 8.2 summarizes surrogate fidelity metrics for each experiment-position-quantity combination. Mean absolute errors range from 0.019 (FGR, 664 top) to 0.653 (swelling, 664 bottom), with no systematic pattern across experiments. Correlation between surrogate and BISON predictions reflects that both methods are driven by the same 700 posterior samples but BISON simulations incorporate full nonlinear finite element physics while surrogates provide interpolated approximations. The absence of systematic bias in surrogate residuals confirms that Gaussian Process surrogates neither systematically over- nor under-predict BISON outputs.

These validation results confirm that surrogate-accelerated MCMC has successfully recovered posterior distributions that accurately represent full BISON model behavior. The near-identical coverage and comparable prediction accuracy between surrogate and BISON predictions, combined with systematic accuracy across the full observation range demonstrated in Figure 8.3, provide high confidence that the computational acceleration

Table 8.2: Surrogate fidelity metrics comparing Gaussian Process predictions to direct BISON simulations for identical posterior parameter samples. MAE quantifies average prediction discrepancy; RSME penalizes larger errors. Low values across all experiments confirm surrogate accuracy suitable for MCMC acceleration.

Experiment	Quantity	MAE	RMSE	Correlation	Param Unc
602	Swelling	0.442	0.545	0.454	0.516
602	FGR	0.085	0.242	0.149	0.244
603	Swelling	0.269	0.328	0.441	0.370
603	FGR	0.046	0.153	0.154	0.182
664	Swelling	0.286	0.335	0.484	0.288
664	FGR	0.056	0.151	0.342	0.191
658	Swelling	0.327	0.404	0.440	0.832
658	FGR	0.180	0.235	0.444	0.464
Average	—	0.231	0.339	0.349	—

achieved through surrogates has not compromised the statistical rigor or mechanistic validity of the Bayesian calibration.

8.3.3 Out-of-Sample Validation: Held-Out Experiment 613

To test whether posterior parameters learned from Test d generalize to independent data from the same regime, the 700 calibrated posterior samples were evaluated through direct BISON for experiment 613, which was not included in MCMC training. This out-of-sample validation reveals whether the posterior distribution captures mechanistic behavior beyond the specific experiments used for calibration or whether predictions degrade on unseen data from regime boundaries.

Experiment 613 represents conditions nominally within the mechanistic regime where the model was developed (low temperature, low burnup) but with independent experimental measurements not used in calibration. The experiment encompasses three axial positions with varying temperatures and burnup, providing six independent validation predictions (three positions, two quantities). Notably, Test 613 middle position operates at $T = 1505$ K, exceeding the identified regime boundary of 1485 K, providing a critical test of whether the model exhibits graceful degradation or catastrophic failure at regime boundaries.

Out-of-sample results reveal regime-dependent behavior with mechanistically interpretable patterns. For Test 613 predictions overall, 50% of observations fell within 95% credible intervals (3 of 6 observations), lower than the 85% achieved on Test d calibration data. However, this degradation is not uniformly distributed across positions. Table 8.3 presents detailed results segregated by position and quantity.

Test 613 bottom and top positions, with temperatures of 1408 K and 1394 K respectively, fall safely within the validated regime. Swelling predictions at these positions show good coverage: both bottom and top achieved coverage within the credible intervals with small prediction errors (underprediction of 0.28 vol% and 0.30 vol% respectively). These results

match Test d performance characteristics, demonstrating that the posterior distribution successfully generalizes to independent data from within the valid regime.

Fission gas release predictions at Test 613 show systematic underprediction across all three positions. At bottom position ($T = 1408$ K), the observation is 0.03% while the posterior mean predicts 0.007%, falling outside the narrow credible interval of [0.000, 0.001%]. At middle position ($T = 1505$ K, above regime boundary), the observation is 0.06% with posterior mean 0.018% and credible interval [0.000, 0.134%]. At top position ($T = 1394$ K), the observation is 0.04% versus predicted mean 0.005% with credible interval [0.000, 0.000%]. The systematic underprediction across all positions suggests that FGR mechanisms at these temperatures are not fully captured by the posterior parameter distribution learned from Test d.

The mechanism underlying this FGR underprediction reflects the physics captured during calibration. Test d, where FGR is small (0.01–0.2%), calibrated the posterior distribution to predict minimal FGR. The posterior samples learned that fission gas generation and transport are limited through low diffusivity and efficient trapping mechanisms appropriate to the low-temperature regime. When extrapolating to Test 613 at higher (though still moderately elevated) temperatures, the posterior predicts based on these calibrated parameters: low FGR. However, experimental observations show higher FGR, indicating that either diffusivity is enhanced or trapping is less efficient at these temperatures relative to model predictions. This mechanistically interpretable underprediction pattern demonstrates that the model is not breaking down catastrophically but rather is failing to capture temperature-dependent mechanisms that become active above approximately 1400 K.

Test 613 middle position at $T = 1505$ K, which exceeds the identified regime boundary, shows swelling coverage failure (observation 2.28% falls below credible interval lower bound of 1.28%) while FGR achieves coverage due to the wider credible interval. This mixed coverage pattern at the regime boundary provides evidence that different physical mechanisms transition at different rates: swelling (driven by void nucleation and growth) shows sharp coverage failure, while FGR (driven by diffusion and bubble coalescence) exhibits gradual underprediction across the temperature range.

Table 8.3: Out-of-sample validation results for held-out experiment 613 not included in MCMC training.

Position	Quantity	T (K)	Obs	Mean	95% CI	In CI?
Bottom	Swelling	1408	2.080	1.799	1.196–2.232	Yes
Bottom	FGR	1408	0.030	0.007	0.000–0.001	No
Middle	Swelling	1505	2.280	1.764	1.284–2.223	No
Middle	FGR	1505	0.060	0.018	0.000–0.134	Yes
Top	Swelling	1394	1.930	1.632	1.092–2.023	Yes
Top	FGR	1394	0.040	0.005	0.000–0.000	No

These results demonstrate that the posterior distribution successfully captures mechanistic behavior within the calibrated regime and exhibits graceful, mechanistically interpretable

degradation at regime boundaries. The out-of-sample validation confirms both the strengths and the limitations of the calibrated model: swelling predictions generalize well to new data within the valid regime, while FGR predictions systematically underestimate at temperatures above approximately 1400 K, and swelling predictions show coverage failure above 1485 K. These patterns validate the regime hypothesis and provide guidance for model improvement: enhanced understanding of temperature-dependent diffusivity and trapping mechanisms is required for accurate FGR predictions outside the calibrated regime.

8.3.4 Prior Versus Posterior Comparison: Quantifying Information Gain

To quantify the benefit of Bayesian calibration, posterior predictions using calibrated MCMC distributions are compared against prior predictions using uninformed Latin hypercube sampling over the uncalibrated parameter space. This comparison demonstrates the magnitude of parameter constraint achieved through experimental calibration and validates that Bayesian inference has extracted genuine mechanistic information from observational data rather than merely reducing uncertainty through posterior collapse or overfitting.

The comparison encompasses 2100 prior samples from the uninformed Latin hypercube design and 700 posterior samples from the calibrated MCMC chains, both evaluated through the same surrogate models across all Test d experiment-position-quantity combinations. This prior-posterior contrast reveals how much the posterior has concentrated relative to the prior and whether predictions have improved in accuracy relative to observations.

8.3.4.1 Uncertainty Reduction

Posterior predictions show dramatic uncertainty reduction relative to prior predictions across both quantities. For swelling predictions, the prior (LHS) standard deviation averaged 3.053 vol%, while the posterior (MCMC) standard deviation reduced to 0.520 vol%, representing an average reduction of 82.4%. This uncertainty collapse is not uniform across experiments but is systematic: uncertainty reduction ranges from 72.98% to 87.40%, with higher reductions at lower temperatures where data constraints are strongest. The magnitude of this reduction is substantial, indicating that the 2100 prior samples explored a much wider region of parameter space than is consistent with experimental observations.

Fission gas release predictions show even more dramatic uncertainty reduction. The prior standard deviation averaged 2.353%, while the posterior reduced to 0.284%, yielding an average reduction of 88.7% with a range of 81.07% to 93.31%. The greater uncertainty reduction in FGR relative to swelling reflects that FGR is more sensitive to parameter variations and thus more strongly constrained by observations. The consistent range of uncertainty reduction across all experiments (all above 80%) indicates robust parameter identification: the posterior has not concentrated in isolated parameter subspaces but rather across multiple regimes.

8.3.4.2 Accuracy Improvement

Beyond uncertainty reduction, posterior predictions demonstrate improved accuracy relative to prior predictions. For swelling, the prior mean absolute error against observations was 2.293 vol%, while posterior MAE reduced to 0.491 vol%, an improvement of 78.6%. This accuracy gain demonstrates that Bayesian calibration has not only narrowed prediction distributions but also shifted the mean predictions toward observations, indicating that the experimental data have identified parameter regions that produce more accurate predictions.

Fission gas release shows substantially larger accuracy improvements. Prior MAE was 1.551% while posterior MAE reduced to 0.115%, yielding a 92.6% improvement. This exceptionally large improvement reflects that the prior distribution centered FGR predictions far from observations (the uninformed parameter space permits extreme values), while the posterior has concentrated on parameter values producing accurate FGR predictions. The magnitude of this improvement provides compelling evidence that the MCMC has converged to mechanistically sensible parameter distributions rather than merely tightening uncertainty bounds.

8.3.4.3 Calibration Journey: Representative Example

Figure 8.4 demonstrates experiment 603 middle ($T = 1450$ K, well within the validated regime), showing textbook Bayesian calibration, broad, uncertain prior (swelling mean = 4.28 vol%, std = 2.38 vol%) concentrating to a narrow, accurate posterior (mean = 2.63 vol%, std = 0.40 vol%) that encapsulates the observation of 2.70 vol%. FGR shows the characteristic pattern of FGR calibration with a prior with significant probability mass at high values (mean = 1.56%, std = 2.44%) collapsing to a concentrated posterior near zero (mean = 0.056%, std = 0.223%), matching the observed 0.100% FGR.




Figure 8.4: Calibration for experiment 603 middle ($T = 1450$ K). Swelling and FGR both show standard Bayesian calibration with substantial uncertainty reduction (83.2% and 90.8% respectively) and posterior distributions encapsulating observations with appropriate confidence intervals.

8.3.4.4 Information Gain Summary

Across all 24 predictions in Test d (12 experiments spanning 5 specimens, 3 positions, 2 quantities), Bayesian calibration has achieved an average uncertainty reduction of 86%, with swelling at 82.4% and FGR at 88.7%. This magnitude of uncertainty reduction, combined with 78.6% and 92.6% accuracy improvements respectively, demonstrates that experimental observations contain substantial information about model parameters. The consistent reduction across diverse experiment conditions (varying temperature, burnup, and position) validates that the posterior has learned robust, mechanistically sound parameter distributions applicable across the regime rather than specialized to particular experiments.

The observed pattern of larger uncertainty reduction in FGR (89%) compared to swelling (82%) reflects the physical structure of the inverse problem: FGR is more sensitive to diffusivity and trapping parameters, which are more strongly constrained by observations, while swelling depends on threshold phenomena (percolation) less sensitive to small parameter variations. This differential information gain across quantities demonstrates that the Bayesian analysis has revealed the mechanistic structure of parameter influence, not merely reduced uncertainty uniformly.

8.3.5 Cross-Regime Validation Summary

The calibrated posterior parameter distribution from Test d was applied to predict fuel performance across five distinct experimental regimes spanning a broader range of temperatures and burnup levels. These cross-regime predictions assess the generalization of calibration beyond the primary regime and identify the operational boundaries within which the mechanistic model provides predictive capability.

The five regime configurations are defined by thermal and irradiation conditions. Test a represents high-temperature conditions (specimen 220) that explore behavior at elevated temperatures. Test b encompasses low-temperature ($T < 1485$ K), low-burnup ($B < 1\%$) experiments including diverse specimens across the thermal range, spanning the validated domain with multiple data points. Test c combines high-temperature ($T > 1485$ K) with high-burnup ($B > 1\%$) conditions, reaching mechanistic extremes. Test d, the primary calibration regime, includes low-temperature, low-burnup experiments similar to Test b but excludes specimen 669, focusing on a cohesive set of intermediate conditions. Test e, designated as high burnup data, remains entirely separate from calibration and serves to assess predictive performance.

Table 8.4 summarizes posterior prediction performance across all five regimes, evaluated using the same metrics as primary validation: credible interval coverage at the 95% level, mean absolute error, root mean square error, and mean relative error. The results reveal a clear progression in model performance as conditions move systematically across the temperature and burnup space.

Table 8.4: Cross-regime validation metrics for posterior predictions. Coverage percentage indicates the fraction of observations falling within 95% credible intervals. MAE and RMSE quantify prediction accuracy; Mean relative error shows signed bias.

Test	Quantity	N	Coverage %	MAE	RMSE	Mean Rel Error %
a	Swelling	3	100.0	0.956	0.985	−3.56
b	Swelling	14	71.4	0.443	0.561	3.80
c	Swelling	6	66.7	0.527	0.547	2.25
d	Swelling	11	81.8	0.297	0.352	4.06
e	Swelling	3	100.0	0.811	0.839	9.88
a	FGR	3	66.7	2.357	2.954	−29.35
b	FGR	14	71.4	0.049	0.057	93.16
c	FGR	6	0.0	1.213	1.287	2400.39
d	FGR	11	90.9	0.050	0.061	123.86
e	FGR	3	100.0	1.122	1.537	−5.55

Performance in Test a (high temperature) demonstrates robust swelling predictions with 100% coverage and modest error (MAE 0.956 vol%), confirming that calibrated parameters extrapolate accurately to elevated temperatures. Fission gas release predictions in this regime show reduced coverage (66.7%) with larger error (2.357%), suggesting that at high temperatures where the calibration data provides limited FGR constraints, model predictions encounter greater relative uncertainty. Nevertheless, the 100% swelling coverage indicates mechanistic validity at these higher temperatures.

Test b results reveal nominal calibration performance spanning the low-temperature, low-burnup domain. Swelling achieves 71.4% coverage with MAE of 0.443 vol%, while FGR maintains 71.4% coverage with very small absolute errors (0.049%). The consistency between Test b and Test d metrics suggests that calibration has captured mechanistic behavior applicable across the validated low-temperature regime regardless of specific specimens, with the larger Test b sample size (14 predictions versus 11 in Test d) providing broader evidence of generalization.

Test c performance degrades significantly, particularly for fission gas release. Swelling coverage remains acceptable at 66.7%, but FGR coverage collapses to 0% with catastrophic error (MAE 1.213%, mean relative error 2400%). This dramatic failure indicates that at the extreme condition combination of high temperature and high burnup represented by Test c, fundamental model assumptions break down. The binary percolation mechanism does not capture the dramatically enhanced FGR observed in this regime, likely due to mechanistic transitions such as increased grain boundary bubble mobility, accelerated coalescence, or activation of alternative gas release pathways at extreme conditions. The swelling predictions remain reasonable, suggesting the swelling mechanism retains some validity while FGR mechanisms have fundamentally shifted.

Test d (primary calibration regime) maintains the strong validation performance documented in Section 8.3.2, with swelling at 81.8% coverage (MAE 0.297 vol%) and FGR at

90.9% coverage (MAE 0.050%). These metrics confirm that calibration was effective within the carefully selected low-temperature, low-burnup domain that defines the mechanistic regime of validity.

Test e (high burnup hold-out validation) demonstrates the most encouraging result: 100% coverage for both swelling and FGR with reasonable accuracy (0.811 vol% and 1.122%, respectively). This perfect coverage on entirely held-out data validates that the posterior parameter distribution captures genuine mechanistic information and has not overfit to calibration data. The high-burnup conditions in Test e differ from calibration, yet predictions maintain 100% coverage, indicating that despite increased FGR magnitudes at high burnup, the calibrated parameters provide appropriately uncertain predictions. The Test e FGR predictions exhibit larger absolute error (1.122%) but maintain correct coverage properties, indicating well-calibrated uncertainty quantification even when extrapolating to higher burnup.

The cross-regime analysis reveals three operationally distinct regimes for mechanistic prediction. The low-temperature, low-burnup domain (Tests b and d) exhibits strong calibration with nominal coverage (71%–82%) and small prediction errors across both quantities. Test a at elevated temperature maintains excellent swelling predictions while showing FGR degradation, suggesting temperature-dependent mechanistic effects emerge above calibration conditions. The high-temperature, high-burnup extreme (Test c) falls completely outside the mechanistic validity of the binary percolation model, with FGR predictions failing catastrophically while swelling remains partially predictable. Test e demonstrates that high-burnup conditions alone do not break the model; rather, the combination of high temperature and high burnup at Test c exceeds mechanistic bounds.

These results motivate the accelerated fuel qualification framework: calibration has successfully constrained parameters within a well-defined mechanistic regime of low-temperature ($T < 1485$ K), low-burnup ($B < 1\%$) conditions. Extending the validated domain to higher temperatures and burnups requires either mechanistic model enhancement (incorporating temperature-dependent diffusion, alternative gas release pathways, and burnup-dependent phenomena) or targeted experiments at regime boundaries (intermediate-temperature high-burnup conditions) that would provide the data necessary to calibrate new mechanistic components. The clear performance degradation at extreme conditions provides diagnostic evidence for model form limitations rather than parameter uncertainty, directly informing where experimental and modeling effort should be directed in subsequent acceleration phases.

8.4 Conclusions: Validation and Mechanistic Constraints

The Bayesian calibration and validation methodology has successfully recovered posterior parameter distributions that produce accurate, well-calibrated predictions across a mecha-

nistically consistent regime. The comprehensive validation campaign demonstrates that surrogate-accelerated MCMC has not only reduced computational cost but has maintained statistical rigor and mechanistic fidelity throughout the inference process.

The posterior predictive distributions in the primary regime (Test d) achieved near-nominal coverage at 86.4% (19 of 22 observations within 95% credible intervals), with average prediction errors of 0.298 vol% for swelling and 0.050% for fission gas release. These metrics demonstrate accurate point predictions coupled with appropriate uncertainty quantification, validating that the inverse problem has recovered parameter values consistent with experimental observations. The tight correspondence between predicted and observed values across the full observation range, evidenced by the predicted-versus-observed scatter plot, confirms systematic accuracy without systematic bias.

Surrogate model fidelity validation confirmed that Gaussian Process surrogates accurately represent full BISON finite element physics across the posterior parameter space. The mean absolute error between surrogate and BISON predictions averaged 0.231, with surrogate-based and BISON-based coverage metrics matching at 86.4% and 85.0% respectively. This consistency validates that computational acceleration through surrogate approximation has not compromised mechanistic accuracy or introduced artificial uncertainty.

Prior-versus-posterior comparison quantified the information extracted from experimental data. Bayesian calibration achieved 82.4% uncertainty reduction in swelling parameters and 88.7% uncertainty reduction in fission gas release parameters, with 78.6% and 92.6% improvements in prediction accuracy respectively. These improvements demonstrate that the experimental observations contained substantial mechanistic information, which the MCMC algorithm successfully identified and incorporated into the posterior distribution. The consistent reduction magnitude across diverse experiments and positions validates that parameter identification was robust rather than specialized to particular test conditions.

Hold-out validation on Test e data confirmed that the posterior distribution generalizes beyond calibration data. The hold-out test achieved perfect coverage (100% for both quantities) with reasonable prediction errors (0.811 vol% swelling, 1.122% FGR), providing independent evidence that the calibration has recovered mechanistically valid parameters rather than achieving accuracy through overfitting.

Out-of-sample validation at regime boundaries identified mechanistic limits. Test 613 at temperatures exceeding 1485 K exhibited systematic underprediction in both swelling (0.4–1.1 vol% misses) and fission gas release (0.036–0.125% underprediction), despite tight posterior distributions. This pattern of confident incorrect predictions indicates model form inadequacy rather than parameter uncertainty, directly revealing the temperature regime where additional mechanistic components become necessary.

Cross-regime analysis across Tests a through e established three operationally distinct performance zones. Tests a, b, d, and e (spanning 0.5–1.1 at% burnup and 1269–1505 K temperature) maintain nominal to excellent calibration with average coverage above 70% and prediction errors remaining small relative to observation magnitudes. Test c at extreme conditions (high temperature, high burnup) exhibited catastrophic failure in fission gas release (0% coverage, 2400% relative error), demonstrating unambiguous model form

breakdown requiring mechanistic enhancement or specialized experiments.

The complete validation campaign validates four critical properties of the Bayesian calibration system:

Accuracy: Posterior predictions match observations to within measurement noise across the primary regime, with MAE values (0.297 vol% swelling, 0.050% FGR) substantially smaller than observation ranges.

Calibration: Credible intervals provide honest uncertainty quantification, with coverage metrics near nominal levels (86.4%) rather than artificially tight or loose intervals. The balance between coverage and accuracy indicates that posterior distributions have captured genuine parameter constraints without overconfidence.

Fidelity: Surrogate approximation error (average MAE 0.231) is small relative to both parameter uncertainty and total prediction uncertainty, validating that computational acceleration has not sacrificed mechanistic accuracy.

Generalization: Hold-out validation and cross-regime testing demonstrate that posterior parameters reflect mechanistic relationships applicable across diverse conditions, rather than overfitted specialization to calibration data.

These results establish the quantitative foundation for accelerated fuel qualification. Parameter uncertainty has been rigorously characterized, prediction uncertainty has been properly decomposed and quantified, and mechanistic regime boundaries have been identified through systematic comparison of prediction against observation. The Bayesian framework has successfully translated experimental observations into constrained parameter distributions that enable confident fuel performance prediction within the validated domain and diagnostic indication of where enhanced mechanistic models or targeted experiments are required to extend that domain.

Chapter 9

Conclusions and Future Directions

9.1 Central Hypothesis and Research Objectives

This dissertation investigated a central hypothesis fundamental to accelerated fuel qualification. The hypothesis says that mechanistic models grounded in first-principles physics, systematically validated against experimental data across multiple operating regimes, and calibrated through rigorous Bayesian uncertainty quantification can substantially reduce epistemic uncertainty in nuclear fuel performance predictions while establishing defensible boundaries for qualified operational domains.

The overarching research question driving this work was concerned whether integrating advanced mechanistic modeling with data science techniques could compress fuel qualification timelines while maintaining regulatory defensibility via targeted planning of experiments. For UN specifically, the question transformed into whether sufficient mechanistic understanding could be developed and remaining uncertainty quantified such that UN fuel could be confidently qualified for advanced reactor applications, despite historical data limitations and significant scatter in the available experimental database.

The research was organized around four principal objectives derived from the AFQ framework. First, the development of a physics-based mechanistic model for fission gas release and swelling in UN that explicitly captures temperature-dependent diffusion mechanisms and regime transitions. Second, rigorous validation of this model against historical irradiation data and identification of regime-specific performance boundaries. Third, global sensitivity analysis to identify which model parameters most strongly influence predictions and thus warrant calibration priority. Fourth, Bayesian inverse uncertainty quantification to constrain model parameters using available experimental data and quantify remaining uncertainty for use in model refinement and necessary additional targeted testing.

This chapter synthesizes dissertation findings, validates the central hypothesis, and articulates implications for accelerated UN fuel qualification and broader mechanistic modeling methodology.

9.2 Key Findings and Hypothesis Validation

9.2.1 Material Property Foundation

Compilation and characterization of UN material properties from six decades of literature spanning 1960 through 2020 established the baseline epistemic uncertainties in mechanical and thermal inputs to mechanistic models. Development of temperature-dependent correlations with quantified uncertainty bounds for elastic modulus, thermal conductivity, thermal expansion, and diffusion coefficients demonstrated that material property uncertainties are moderate relative to subsequent operational uncertainties. Most correlations achieved residual scatter of 10 to 20% around fitted trends. These baseline uncertainties propagate through coupled thermo-mechanical analyses but remain secondary relative to parameter uncertainties in mechanistic fission gas release and swelling models.

A key finding emerged from this analysis. Material property uncertainties are well characterized and insufficient to explain the order-of-magnitude scatter observed in historical UN irradiation data. Fission gas release and swelling behavior must be governed by mechanistic processes whose parameters remain poorly constrained by legacy data. This observation directly motivated the mechanistic modeling efforts.

9.2.2 Historical Experimental Database Synthesis

Comprehensive synthesis of historical UN irradiation programs spanning 1960 through 2000 revealed fundamental physical transitions in fuel behavior as functions of temperature, power density, and burnup. The SNAP-50 program at CANCEL and ORNL, the BR-10 fast reactor testing in Russia, the SP-100 program at EBR-II and FFTF, NASA Advanced Power Reactor testing, and Battelle high-temperature campaigns collectively demonstrated clear mechanistic transitions.

Volumetric swelling exhibits qualitatively different behavior below and above approximately 1400 to 1500 K. Low-temperature regimes show approximately linear swelling accumulation with burnup whereas high-temperature regimes show rapid saturation followed by plateau behavior. Fission gas release remains suppressed at less than one percent at temperatures below 1400 K but increases nonlinearly above this threshold, reflecting a fundamental change in gas transport mechanisms.

Despite identical macroscopic operating conditions of temperature, burnup, and power density, different irradiation programs exhibit 20 to 30% variation in swelling and 2 to 3 times variation in gas release. This scatter is attributed to fabrication history, fuel grain size, and microstructural evolution. Furthermore, high-temperature data at greater than 1600 K are limited to two SNAP-50 specimens, and high-burnup data at greater than 3%FIMA are sparse, reducing confidence in extrapolated predictions.

A key finding emerged regarding empirical approaches. Empirical correlations developed from historical data cannot perfectly explain program-to-program scatter and lack defensible

extrapolation capability beyond the limited experimental domain. Mechanistic models that capture underlying physics become essential for qualification of novel operating regimes.

9.2.3 Mechanistic Model Implementation

Development of the SIFGRS mechanistic model implemented in BISON introduced explicit physics-based representations of fission gas and swelling phenomena. The model employs a three-regime diffusivity decomposition wherein the total effective fission gas diffusivity decomposes into three additive components. Thermal equilibrium diffusivity D_1 represents thermally activated diffusion through the perfect lattice. Irradiation-enhanced defect-assisted diffusivity D_2 captures enhanced migration pathways created by excess point defects produced during irradiation. Athermal fission-fragment-driven diffusivity D_3 represents athermal transport driven by collision cascades from high-energy fission fragments.

This decomposition enables capture of temperature-dependent transitions in diffusion mechanisms that empirical power-law fits cannot represent. The model explicitly differentiates bubble populations by nucleation site including bulk bubbles in undisturbed lattice, dislocation bubbles nucleating on the dislocation network, and grain boundary bubbles accumulating at interfaces. This population-based approach captured selective migration pathways and residence times that determine gas partitioning between retention and release.

A binary percolation model for grain boundary interconnection provides mechanistic basis for understanding saturation transitions. The model captures the physical transition where rapid swelling at low burnup transitions to quasi-equilibrium between gas production and release at high burnup when grain boundary bubble networks achieve interconnectedness.

A key finding emerged regarding model structure. Mechanistic model structure enables capture of qualitatively distinct physical regimes observed in historical data. However, the model's predictive success depends critically on accurate parameter values that must be extracted from data through inverse calibration rather than guessed from physics alone.

9.2.4 Validation and Regime Identification

Rigorous validation of SIFGRS against SNAP-50 irradiation data revealed clear regime-dependent model performance. The standard low-temperature, low-burnup regime defined by fuel temperatures below 1485 K and burnup below 1%FIMA showed model predictions of fission gas release with systematic bias and swelling with robust agreement across all operating conditions. The baseline model underpredicts fission gas release by an average factor of 2 to 3, attributed to excessive intragranular bubble trapping limiting gas transport to grain boundaries. Volumetric swelling predictions remain robust across parameter variations, indicating that swelling mechanisms are mechanistically sound even when gas transport pathways are underpredicted.

The high-temperature regime above 1600 K with single specimen 220 exhibits anomalously high fission gas release at 12% relative to predictions of 4%. This suggests enhanced

diffusion or reduced trapping at elevated temperatures not captured by model parameterization. The temperature-dependent diffusivity model may underestimate activation coefficients or neglect temperature-dependent transitions in defect mechanisms.

The high-burnup regime at temperatures of 1300 to 1500 K with burnup exceeding 3%FIMA shows specimen 665 with fission gas release of approximately 6% versus predicted less than 1%. This suggests mechanistic limitations at extended burnup where bubble coalescence and interlinkage processes may differ from low-burnup assumptions. Swelling predictions remain qualitatively reasonable but quantitatively undershoot observations, consistent with underestimated grain boundary coalescence rates.

Specimens straddling the identified 1485 K boundary show transitional behavior with partial coverage failures, indicating that this boundary reflects a real mechanistic transition rather than an arbitrary threshold.

A key finding emerged regarding model performance. The mechanistic model successfully captures dominant physical behavior within the low-temperature, low-burnup regime while exhibiting interpretable degradation at regime boundaries. Model failures are mechanistically meaningful rather than numerical artifacts, pinpointing specific physical processes requiring model enhancement including temperature-dependent diffusivity and high-burnup bubble coalescence mechanisms.

9.2.5 Global Sensitivity Analysis

Morris elementary effects analysis across the full 14-dimensional mechanistic parameter space revealed parameter importance hierarchy. Eight parameters exhibit substantial influence with mean absolute elementary effects greater than 2 on fission gas release predictions. These parameters ranked by importance are grain radius with mean absolute elementary effect of 11.0, saturation coverage at 7.3, dislocation trap strength at 6.6, dislocation density at 6.1, resolution rate at 4.0, trapping rate at 2.7, and intragranular diffusion at 2.5.

Simultaneous high sensitivity to both dislocation trapping parameters and diffusivity coefficients indicates that fission gas release is limited by two coupled stages. Intragranular and dislocation trapping and matrix diffusion interact such that gas cannot be released by adjusting either stage alone. The top parameters maintain consistent ranking across the full operating envelope spanning temperatures of 1200 to 1750 K, burnups of 0.3 to 4.6%FIMA, and power densities of 0.3 to 1.5 kW/cm³. This indicates that the same underlying physical mechanisms control fission gas behavior across regimes, though optimal parameter values may vary.

A key finding emerged regarding model structure. Parameter sensitivity analysis validated the mechanistic model parametric study conducted such that calibration should focus on the 14 identified parameters while other mechanisms can be fixed. The consistency of parameter importance across regimes enables unified Bayesian calibration strategy rather than requiring fundamentally different models for each regime.

9.2.6 Bayesian Parameter Calibration and Information Gain

Surrogate-accelerated Bayesian MCMC calibration using 2100 BISON simulations to train Gaussian Process emulators, followed by NUTS sampling for parameter inference, achieved robust posterior convergence across multiple experimental regimes. Within the primary low-temperature, low-burnup calibration domain, posterior parameter distributions achieved full convergence with Gelman-Rubin $\hat{R}$ less than 1.01 for all parameters. Effective sample size exceeded 1500 for dominant parameters, enabling reliable posterior quantiles and credible intervals.

Parameter uncertainty reduction from prior to posterior averaged 82 to 89% across model outputs, demonstrating substantial information gain from historical experimental data. Specific parameters showed the following posterior behavior. Grain radius and dislocation density were constrained to microstructurally realistic values of 37 μm grain radius and $11 \times 10^{14} \text{ m}^{-2}$ dislocation density in the primary regime by SNAP-50 measurements, with posteriors reflecting these constraints rather than data-driven identification.

Saturation coverage and vacancy grain boundary diffusion showed primary regime calibration recovering posterior means of 0.35 saturation coverage and 11×10^7 grain boundary vacancy diffusion. Based on the parameteric study, this indicates that grain boundary saturation occurs at slightly lower coverage and higher diffusion than the baseline assumption. Posterior diffusivity multiplier concentrated near 6 times the baseline, indicating that diffusivity estimates from lower-length-scale calculations are slightly under predicted within the primary regime. Posterior resolution rate scaled to 5 times baseline, modestly higher than original parameterization, confirming that fission fragment resolution is not the primary limitation on gas release. Overall, information gain is a key metric to ensuring parameter refinement. While the parameter mean values shifted away from the baseline, additional experimental data and analysis is required to better define these values. However, the MCMC results identified the same trends as seen the parameteric study; trapping tends to be too high, diffusion and resolution tend to be too low and an earlier onset of FGR is required.

Secondary regimes showed data insufficiency. The high-temperature test containing single specimen 220 proved insufficient to constrain the 14-dimensional parameter space, with posterior distributions remaining wide at effective sample size less than 500 despite achieving good data fit. This indicates that one high-temperature experiment cannot uniquely identify mechanistic parameter values. The high-burnup test containing single specimen 665 similarly proved insufficient for parameter identification, with posterior spread reflecting mechanistic ambiguity rather than experimental noise. The moderate-temperature regime straddling the transient zone showed multiple specimens within the regime boundary displaying partial parameter convergence with elevated Gelman-Rubin values of 1.01 to 1.05, indicating that data from the regime boundary cannot uniquely constrain parameters suited to that boundary.

A key finding emerged regarding data requirements. Bayesian calibration successfully identified parameter distributions for the low-temperature, low-burnup regime but revealed

that existing historical data are insufficient for unique parameter identification in secondary regimes. Additionally, the low-temperature low-burnup regime struggles to clearly identify low temperature FGR migration parameters as seen by the overall information gain. This demonstrates precisely where targeted experimental efforts should focus. Low temperature FGR migration, High-temperature domains above 1600 K, high-burnup domains of 3 to 5%FIMA, and moderate-temperature transitional regimes require new experiments to enable regime-specific calibration.

9.2.7 Posterior Predictions and Forward Uncertainty Quantification

Forward propagation of calibrated posterior parameter distributions through BISON produced well-calibrated predictions for the primary regime. Within the low-temperature, low-burnup domain, 86.4% of experimental observations, specifically 19 of 22 data points, fell within 95% credible intervals. This is slightly below the nominal 95% level but within expected statistical variation for the sample size and consistent with well-calibrated uncertainty quantification. Volumetric swelling predictions achieved mean absolute error of 0.30 vol% against observations ranging 1.2 to 4.3 vol%. Fission gas release achieved mean absolute error of 0.05 % against observations ranging 0.01 to 0.2 %.

Direct BISON simulations using posterior parameters achieved 85% coverage, confirming that Gaussian Process surrogates accurately represent full model behavior. The mean surrogate-BISON error was 0.23 vol% for swelling and 0.09 % for FGR.

Quantification of information gain comparing prior versus posterior predictions revealed substantial uncertainty reduction. Posterior standard deviation reduced by 82.4% for swelling, decreasing from prior 3.05 vol% to posterior 0.52 vol%, and 88.7% for fission gas release, decreasing from prior 2.35 % to posterior 0.28 %. Prior predictions exhibited mean absolute error of 2.29 vol% for swelling and 1.55 % for FGR. Posterior predictions improved to 0.49 vol% and 0.12 %, respectively, representing 78.6% and 92.6% error reduction.

The magnitude of information gain validates that SNAP-50 experiments contain sufficient information content to constrain the primary mechanistic parameters. The information gain is not artificial posterior collapse but rather reflects genuine recovery of parameter values supported by multiple experimental observations.

Out-of-sample validation using held-out experiment 613 confirmed prediction generalization within the calibrated regime. Within-regime validation at temperatures below 1450 K showed good generalization with small underprediction of 0.28 to 0.30 vol% for swelling. Regime-boundary validation at temperature 1505 K showed expected failure with swelling coverage failure while fission gas release remained within credible intervals, demonstrating mechanistically interpretable degradation rather than catastrophic model breakdown. Fission gas release systematic underprediction across all Test 613 positions indicates that temperature-dependent mechanisms including enhanced diffusion or reduced trapping become active at approximately 1400 K, above the validated regime.

A key finding emerged regarding calibration success. Bayesian calibration has succeeded in extracting maximum information from historical SNAP-50 data, producing predictions with 86% experimental coverage, small prediction errors, and parameter uncertainty reduced by 80 to 90%. The out-of-sample validation confirms that predictions generalize within the calibrated regime while exhibiting expected failures at regime boundaries, validating both the mechanistic model structure and the identified regime boundaries.

9.3 Validation of Central Hypothesis

The central hypothesis that mechanistic models calibrated through Bayesian inference can substantially reduce epistemic uncertainty in fuel performance predictions is validated within the low-temperature, low-burnup regime with the following supporting evidence. Mechanistic models capture qualitatively distinct physical regimes including saturation transitions and temperature-dependent diffusivity. The hypothesis element asserting that physics-based models reduce uncertainty is confirmed.

Systematic validation has been achieved through Chapter 5 validation against the SNAP-50 database across 39 individual test cases. This demonstrates mechanistic model capability to capture dominant behavior. Regime-stratified analysis identifies specific regime boundaries rather than claiming universal applicability. The hypothesis element asserting that rigorous validation establishes defensible domains is confirmed.

Parameter identification through data has been demonstrated through Chapter 7 Bayesian calibration achieving 80 to 90% parameter uncertainty reduction within the primary regime. This directly demonstrates that experimental data contain sufficient information to constrain mechanistic parameters. The hypothesis element asserting that Bayesian quantification establishes realistic bounds is confirmed.

Uncertainty quantification enabling decision making has been demonstrated through Chapter 8 posterior predictions with 86% coverage and quantified regime boundaries providing regulatory-defensible basis for qualified operational envelope. The identification of data insufficiency in secondary regimes demonstrates how mechanistic modeling with uncertainty quantification enables directed experimental planning. This is precisely the capability required for accelerated fuel qualification programs. The hypothesis element asserting that uncertainty quantification enables qualified extrapolation is confirmed.

However, the hypothesis is valid only within specific regime boundaries. The low-temperature, low-burnup domain defined by temperature less than 1485 K and burnup approximately less than 1 %glsFIMA represents the region where model assumptions are satisfied and experimental data provide sufficient constraints. Beyond these boundaries, the model exhibits interpretable but uncalibrated predictions, requiring additional physics enhancement or experimental data.

9.4 Summary and Future Directions

9.4.1 Computational Contributions

This dissertation demonstrates that mechanistic modeling combined with Bayesian calibration can systematically reduce parameter uncertainties in complex fuel performance codes—but only where sufficient data exists and model assumptions remain valid. Three key methodological innovations proved essential. Regime-stratified validation that explicitly identifies where model assumptions hold rather than claiming universal validity, surrogate-accelerated inference achieving significant computational speedup while maintaining statistical rigor through careful uncertainty calibration, and mechanistically-informed priors that improved parameter identifiability by 30-50% compared to uninformed approaches.

The regime-stratified approach proved particularly valuable. Rather than forcing a single model across all conditions, explicit identification of regime boundaries (e.g., the 1485 K transition) enables transparent discussion of validity limits and prioritizes where physics enhancements are needed. This framework generalizes beyond fuel performance to any complex multi-physics system where dominant mechanisms shift across operating conditions.

9.4.2 What Was Achieved

Within the low-temperature, low-burnup regime (1213–1485 K, 0–3%FIMA), Bayesian calibration successfully reduced parameter uncertainties by 80–90%, compressing epistemic uncertainty associated with mechanistic model parameters. Validation predictions achieved 86.4% experimental coverage with swelling predictions accurate to ± 0.30 vol% mean absolute error and fission gas release predictions accurate to $\pm 0.05\%$ mean absolute error. The mechanistic understanding in this regime—particularly the decomposition of fission gas diffusion into temperature-dependent components, multi-population bubble nucleation and growth dynamics, and grain boundary coalescence transitions—proved sufficient for characterizing uranium mononitride fuel behavior with rigorously quantified confidence intervals.

Comparison of prior versus posterior predictions quantified information gain from the Bayesian calibration: swelling prediction accuracy improved 78.6% and fission gas release prediction accuracy improved 92.6%, demonstrating that systematic calibration against historical irradiation data substantially reduced model-form uncertainty. The posterior parameter distributions converged with Gelman–Rubin statistics $\hat{R} < 1.01$, confirming that MCMC chains achieved full mixing and Bayesian validity.

These results establish a qualified operational envelope directly relevant to micro-reactor deployment, spanning temperature and burnup conditions representative of current and planned advanced reactor designs. The quantified uncertainties, mechanistic basis, and demonstrated predictive capability provide a defensible foundation for fuel qualification that would be impossible to justify through purely empirical correlations or limited historical

data alone. This approach thus demonstrates the viability of accelerated fuel qualification methodology for next-generation fuel systems.

9.4.3 What Remains Uncertain

The framework also revealed significant limitations. High-temperature specimens (>1500 K) showed fission gas release $2\text{--}3\times$ higher than predictions, and high-burnup specimens ($>3\%$ FIMA) exhibited swelling $20\text{--}30\%$ higher than model estimates. In both cases, limited data availability prevented effective parameter calibration, and posterior distributions remained wide. These discrepancies suggest either missing physics (enhanced diffusion pathways, altered bubble coalescence kinetics) or inadequate parameterization of known mechanisms.

The transition region between identified regimes ($1400\text{--}1500$ K) showed mixed performance, indicating the three-regime decomposition may oversimplify gradual mechanistic transitions. Additionally, structural model limitations remain. The SIFGRS model assumes uniform grain size and dislocation density, neglects spatial heterogeneity in microstructure, and treats grain boundary coalescence as a threshold rather than continuous transition. These simplifications are explicitly acknowledged rather than hidden within empirical fits.

9.4.4 Priorities for Future Investigation

Near-term efforts should focus on three areas. First, a targeted sensitivity study on gas and vacancy diffusion including atomistic calculation uncertainty could assist in low temperature diffusion and regime transitions. Second, modern post-irradiation examination of UN using focused evaluation of fission gas migration would directly measure grain boundary bubble coverage and interconnectedness. Third, targeted irradiation campaigns at high temperature ($1600\text{--}1700$ K) and extended burnup ($3\text{--}4\%$ FIMA) would provide data for regime-specific calibration.

Longer-term integration with accelerated fuel qualification programs requires compiling mechanistic modeling results, Bayesian analysis, and regime-stratified validation into regulatory-grade documentation with explicit uncertainty margins and planned refinements. As new experimental data becomes available, iterative Bayesian calibration cycles can progressively expand qualified operational envelopes with quantified confidence.

References

- [1] Z. A. Miller et al. *Uranium Mononitride (UN) Handbook*. Tech. rep. LA-UR-24-29668. Los Alamos National Laboratory, 2024.
- [2] A. F. Q. W. Group. *Accelerated Fuel Qualification White Paper*. General Atomics, INL, LANL, ORNL, Framatome, Westinghouse, 2021.
- [3] J. D. Hales et al. *BISON theory manual the equations behind nuclear fuel analysis*. Tech. rep. Idaho National Lab.(INL), Idaho Falls, ID (United States), 2016.
- [4] T. Honda and T. Kikuchi. “Porosity Dependence of Elastic Modulus of Uranium Nitride”. In: *Journal of Nuclear Science and Technology* 6.4 (1969), pp. 221–222. DOI: 10.1080/18811248.1969.9732871.
- [5] A. Padel and C. H. deNovion. “Constantes elastiques des carbures, nitrures et oxydes d’uranium et de plutonium”. In: 33.1 (1969), pp. 40–51. DOI: 10.1016/0022-3115(69)90006-3.
- [6] H. Whaley, W. Fulkerson, and R. Potter. “Elastic Moduli and Debye Temperature of Polycrystalline Uranium Nitride by Ultrasonic Velocity Measurements”. In: 31.3 (1969), pp. 345–350. DOI: 10.1016/0022-3115(69)90234-7.
- [7] A. R. Hall. “Elastic moduli and internal friction of some uranium ceramics”. In: 37.3 (1970), pp. 314–323. DOI: 10.1016/0022-3115(70)90161-3.
- [8] H. Muta, K. Kurosaki, M. Uno, and S. Yamanaka. “Thermophysical properties of several nitrides prepared by spark plasma sintering”. In: 389.1 (2009), pp. 186–190. DOI: 10.1016/j.jnucmat.2009.01.027.
- [9] E. O. Speidel and D. L. Keller. *Fabrication and properties of hot-pressed uranium mononitride*. Report. Battelle Memorial Inst., Columbus, OH, 1963. DOI: 10.2172/4674236.
- [10] K. M. Taylor and C. H. McMurtry. *Synthesis and Fabrication of Refractory Uranium Compounds. Summary Report for May 1959 through December 1960*. Report. Carborundum Co. Research and Development Div., Niagara Falls, NY, 1961. DOI: 10.2172/4843747.
- [11] M. Guinan and C. F. Cline. “Elastic properties of uranium mononitride at 298 K”. In: 43.2 (1972), pp. 205–206. DOI: 10.1016/0022-3115(72)90153-5.

- [12] J. Adachi et al. “Mechanical properties at sub-microscale and macroscale of polycrystalline uranium mononitride”. In: *Journal of Nuclear Materials* 384.1 (2009), pp. 6–11. DOI: 10.1016/j.jnucmat.2008.09.019.
- [13] M. A. DeCrescente, M. S. Freed, and S. D. Caplow. *Uranium nitride fuel development, SNAP-50*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1965. DOI: 10.2172/4324037.
- [14] F. L. Carlsen and O. L. Harms. *Thermal Expansion of UN*. Report. Oak Ridge National Lab., Oak Ridge, TN, 1964.
- [15] T. G. Kollie and J. P. Moore. *Electrical Resistivity and Thermal Conductivity of Uranium Mononitride*. Report. Oak Ridge National Lab., Oak Ridge, TN, 1970.
- [16] J. P. Moore, W. Fulkerson, and D. L. McElroy. “Thermal Conductivity, Electrical Resistivity, and Seebeck Coefficient of Uranium Mononitride”. In: *Journal of the American Ceramic Society* 53.2 (1970), pp. 76–82. DOI: 10.1111/j.1151-2916.1970.tb12014.x.
- [17] Y. Takahashi, M. Murabayashi, Y. Akimoto, and T. Mukaibo. “Uranium mononitride: heat capacity and thermal conductivity from 298 to 1000 K”. In: *Journal of Nuclear Materials* 38.3 (1971), pp. 303–308. DOI: 10.1016/0022-3115(71)90059-6.
- [18] B. A. Hayes and M. A. DeCrescente. *Thermal Conductivity and Electrical Resistivity of Uranium Mononitride*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1965. DOI: OSTI:4587070.
- [19] R. W. Endebrock, E. Foster Jr, and D. L. Keller. *Preparation and Properties of Cast UN*. Report. Battelle Memorial Inst., Columbus, OH, 1964. DOI: OSTI:4651895.
- [20] E. F. Westrum Jr. and C. M. Barber. “Uranium mononitride: Heat capacity and thermodynamic properties from 5 to 350 K”. In: *The Journal of Chemical Physics* 45.2 (1966), pp. 635–639. DOI: 10.1063/1.1727621.
- [21] J. F. Counsell, R. M. Dell, and J. F. Martin. “Thermodynamic properties of uranium compounds. Part 2.—Low-temperature heat capacity and entropy of three uranium nitrides”. In: *Transactions of the Faraday Society* 62 (1966), pp. 1736–1747. DOI: 10.1039/TF9666201736.
- [22] C. Affortit. “Chaleur spécifique de UC et UN”. In: *Journal of Nuclear Materials* 34.1 (1970), pp. 105–107. DOI: 10.1016/0022-3115(70)90014-0.
- [23] L. C. Harrington. *Heat Capacity and Enthalpy of Uranium Mononitride from 0 to 1200C (32 to 2200F)*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1963. DOI: 10.2172/4509604.
- [24] F. L. Oetting and J. M. Leitnaker. “The chemical thermodynamic properties of nuclear materials I. Uranium mononitride”. In: *The Journal of Chemical Thermodynamics* 4.2 (1972), pp. 199–211. DOI: 10.1016/0021-9614(72)90057-2.

- [25] E. H. P. Cordfunke and R. P. Muis. “The heat capacity of uranium mononitride”. In: *Journal of Nuclear Materials* 42.2 (1972), pp. 233–234. DOI: 10.1016/0022-3115(72)90031-1.
- [26] J. B. Holt and M. Y. Almassy. “Nitrogen Diffusion in Uranium Nitride as Measured by Alpha Particle Activation of ^{15}N ”. In: *Journal of the American Ceramic Society* 52.12 (1969), pp. 631–635. DOI: 10.1111/j.1151-2916.1969.tb16064.x.
- [27] M. Bradbury and H. Matzke. “Self-diffusion of plutonium in uranium-plutonium mononitride”. In: *Journal of Nuclear Materials* 75.1 (1978), pp. 68–76. eprint: doi: 10.1016/0022-3115(78)90029-6.
- [28] D. Reimann, D. Kroegbr, and T. Lundy. “Cation self-diffusion in UN_{1+x} ”. In: *Journal of Nuclear Materials* 38.2 (1971), pp. 191–196.
- [29] V. Kocevski et al. “Finite temperature properties of uranium mononitride”. In: *Journal of Nuclear Materials* 576 (2023), p. 154241. eprint: doi:10.1016/j.jnucmat.2023.154241.
- [30] S. L. Hayes, J. K. Thomas, and K. L. Peddicord. “Material property correlations for uranium mononitride: II. Mechanical properties”. In: *Journal of Nuclear Materials* 171.2-3 (1990), pp. 271–288. DOI: 10.1016/0022-3115(90)90375-W.
- [31] M. H. Fassler, F. J. Huegel, and M. A. DeCrescente. *Compressive Creep of UC AND UN*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1965. DOI: OSTI:4594868.
- [32] R. Vandervoort, W. Barmore, and C. Cline. “Compressive Creep of Polycrystalline Uranium Mononitride in Nitrogen”. In: *Metallurgical Society of AIME* 242 (1968), pp. 1466–1467.
- [33] U. Masaaki and I. Michio. “Compressive Creep of Sintered UC-UN Solid Solutions”. In: *Journal of Nuclear Materials* 49.1 (1973), pp. 91–97. DOI: 10.1016/0022-3115(73)90065-2.
- [34] H. Matzke. *Science of advanced LMFBR fuels : solid state physics, chemistry, and technology of carbides, nitrides, and carbonitrides of uranium and plutonium*. North-Holland Physics Publishing, Amsterdam. 1986. DOI: OSTI:6332409.
- [35] S. L. Hayes, J. K. Thomas, and K. L. Peddicord. “Material property correlations for uranium mononitride: I. Physical properties”. In: *Journal of Nuclear Materials* 171.2-3 (1990), pp. 262–270. DOI: 10.1016/0022-3115(90)90374-V.
- [36] M. W. D. Cooper et al. “Simulations of self-and Xe diffusivity in uranium mononitride including chemistry and irradiation effects”. In: *Journal of Nuclear Materials* 587 (2023), p. 154685. DOI: 10.1016/j.jnucmat.2023.154685.
- [37] G. T. Craven et al. “Data-driven methods for diffusivity prediction in nuclear fuels”. In: *Computational Materials Science* 230 (2023), p. 112442. DOI: 10.1016/j.commatsci.2023.112442.

- [38] L. Johnson et al. “Machine learning method to determine concentrations of structural defects in irradiated materials”. In: *Comput. Mater. Sci.* 242 (2024), p. 113079.
- [39] C. B. Carter and M. G. Norton. *Ceramic Materials: Science and Engineering*. Second. Vol. 716. Springer, 2007. DOI: 10.1007/978-1-4614-3523-5.
- [40] Q. Yin, A. Kutepov, K. Haule, G. Kotliar, S. Y. Savrasov, and W. E. Pickett. “Electronic correlation and transport properties of nuclear fuel materials”. In: *Physical Review B* 84.19 (2011), p. 195111. DOI: 10.1103/PhysRevB.84.195111.
- [41] K. Kurosaki, K. Yano, K. Yamada, M. Uno, and S. Yamanaka. “A Molecular Dynamics Study of the Thermal Conductivity of Uranium Mononitride”. In: *Journal of Alloys and Compounds* 311.2 (2000), pp. 305–310. DOI: 10.1016/S0925-8388(00)01127-0.
- [42] B. Szpunar and J. A. Szpunar. “Thermal conductivity of uranium nitride and carbide”. In: *International Journal of Nuclear Energy* 2014 (2014). DOI: 10.1155/2014/178360.
- [43] B. Szpunar, J. I. Ranasinghe, L. Malakkal, and J. A. Szpunar. “First principles investigation of thermal transport of uranium mononitride”. In: *Journal of Physics and Chemistry of Solids* 146 (2020), p. 109636. DOI: 10.1016/j.jpcs.2020.109636.
- [44] T. Kikuchi, T. Takahashi, and S. Nasu. “Porosity dependence of thermal conductivity of uranium mononitride”. In: *Journal of Nuclear Materials* 45.4 (1973), pp. 284–292. DOI: 10.1016/0022-3115(73)90162-1.
- [45] S. L. Hayes, J. K. Thomas, and K. L. Peddicord. “Material property correlations for uranium mononitride: IV. Thermodynamic properties”. In: *Journal of Nuclear Materials* 171.2-3 (1990), pp. 300–318. DOI: 10.1016/0022-3115(90)90377-Y.
- [46] S. C. Weaver, J. L. Scott, R. L. Senn, and B. H. Montgomery. *Effects of irradiation on uranium nitride under space-reactor conditions*. Report. Oak Ridge National Lab., Oak Ridge, TN, 1969. DOI: 10.2172/4467048.
- [47] S. I. Porollo, S. N. Ivanov, E. E. Marinenko, and L. M. Zabudko. *Analysis of experimental data on gas release and swelling of UN fuel irradiated in BR-10 reactor*. Conference Paper. 2017.
- [48] S. I. Porollo and S. N. Ivanov. “Corrosion Damage and Mechanical Properties of BR-10 Reactor Fuel-Rod Cladding with Mononitride Uranium Fuel”. In: *Atomic Energy* 125 (2019), pp. 185–193. DOI: 10.1007/s10512-018-00464-3.
- [49] K. J. Bowles and R. E. Gluyas. *Evaluation of refractory-metal-clad uranium nitride and uranium dioxide fuel pins after irradiation for times up to 10 450 hours at 990 C*. Report. NASA Lewis Research Center, Cleveland, OH, 1975.
- [50] R. G. Rohal, T. N. Tambling, and R. L. Smith. *Testing of uranium nitride fuel in T-111 cladding at 1200 K cladding temperature*. Report. NASA Lewis Research Center Cleveland, OH, 1973.

- [51] D. L. Keller and W. Chubb. *Developments in the Technology of Nuclear Fuels During August 1967 through July 1968. Annual Report*. Report. Battelle Memorial Inst., Columbus, OH, 1968. DOI: 10.2172/4282667.
- [52] D. L. Keller. *Progress on Development of Fuels and Technology for Advanced Reactors During October - December 1970. Progress Report*. Tech. rep. Battelle Memorial Inst., Columbus, OH (United States), 1971.
- [53] R. Matthews, K. Chidester, C. Hoth, R. Mason, and R. Petty. “Fabrication and testing of uranium nitride fuel for space power reactors”. In: 151.3 (1988), p. 345. DOI: 10.1016/0022-3115(88)90029-3.
- [54] Y. Arai, A. Maeda, K.-i. Shiozawa, and T. Ohmichi. “Chemical forms of solid fission products in the irradiated uranium—plutonium mixed nitride fuel”. In: 210.1-2 (1994), pp. 161–166.
- [55] D. Brucklacher and W. Dienst. “Creep behavior of ceramic nuclear fuels under neutron irradiation”. In: 42.3 (1972), pp. 285–296.
- [56] R. E. Gluyas and G. K. Watson. *Materials technology for an advanced space power nuclear reactor concept: Program summary*. Report. NASA Lewis Research Center, Cleveland, OH, 1975.
- [57] S. B. Ross, M. S. El-Genk, and R. Matthews. “Uranium nitride fuel swelling correlation”. In: 170.2 (1990), pp. 169–177. eprint: doi:10.1016/0022-3115(90)90409-G.
- [58] H. Matzke. “Atomic mechanisms of mass transport in ceramic nuclear fuel materials”. In: *Journal of the Chemical Society, Faraday Transactions* 86.8 (1990), pp. 1243–1256. DOI: 10.1039/FT9908601243.
- [59] R. Chen et al. “Machine learning models for volumetric swelling in uranium nitride”. In: (2025), p. 155980.
- [60] J. T. Rizk et al. “Mechanistic nuclear fuel performance modeling of uranium nitride”. In: 606 (2025), p. 155604.
- [61] R. A. Lefebvre, B. R. Langley, L. P. Miller, M. Delchini, M. L. Baird, and J. P. Lefebvre. *Neams workbench status and capabilities*. Tech. rep. Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), 2019.
- [62] R. C. Martineau. *The MOOSE multiphysics computational framework for nuclear power applications: A special issue of nuclear technology*. 2021.
- [63] T. Barani, G. Pastore, A. Magni, D. Pizzocri, P. Van Uffelen, and L. Luzzi. “Modeling intra-granular fission gas bubble evolution and coarsening in uranium dioxide during in-pile transients”. In: 538 (2020), p. 152195.
- [64] C. Matthews, D. Schwen, and A. C. Klein. “Radiation re-solution of fission gas in non-oxide nuclear fuel”. In: 457 (2015), pp. 273–278. DOI: 10.1016/j.jnucmat.2014.11.108. URL: <https://www.sciencedirect.com/science/article/pii/S002231151400912X>.

- [65] R. J. White. “The development of grain-face porosity in irradiated oxide fuel”. In: 325.1 (2004), pp. 61–77.
- [66] V. Kocovski, M. W. Cooper, A. J. Claisse, and D. A. Andersson. “Development and application of a uranium mononitride (UN) potential: Thermomechanical properties and Xe diffusion”. In: 562 (2022), p. 153553.
- [67] A. Schneider, J. Rizk, M. Kosmidou, C. Matthews, D. A. Andersson, and M. W. Cooper. “Radiation induced athermal diffusivity in uranium mononitride”. In: 601 (2024), p. 155313.
- [68] L. Börgesson. “Abaqus”. In: *Developments in geotechnical engineering*. Vol. 79. Elsevier, 1996, pp. 565–570.
- [69] M. D. Morris. “Factorial sampling plans for preliminary computational experiments”. In: *Technometrics* 33.2 (1991), pp. 161–174. eprint: doi:10.2307/1269043.
- [70] F. Campolongo, J. Cariboni, and A. Saltelli. “An effective screening design for sensitivity analysis of large models”. In: *Environ. Model. Softw.* 22.10 (2007), pp. 1509–1518. eprint: doi:10.1016/j.envsoft.2006.10.004.
- [71] A. Saltelli et al. *Global sensitivity analysis: the primer*. John Wiley & Sons, 2008.
- [72] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of markov chain monte carlo*. CRC press, 2011.
- [73] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. “Equation of state calculations by fast computing machines”. In: *The journal of chemical physics* 21.6 (1953), pp. 1087–1092.
- [74] M. D. Hoffman, A. Gelman, et al. “The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.” In: *J. Mach. Learn. Res.* 15.1 (2014), pp. 1593–1623.
- [75] M. Kosmidou et al. “Temperature dependency of dislocation evolution in Kr-irradiated Uranium mononitride (UN) through in-situ TEM observation and modeling”. In: *Acta Materialia* 288 (2025), p. 120824.
- [76] A. E. Slaughter et al. “MOOSE Stochastic Tools: A module for performing parallel, memory-efficient in situ stochastic simulations”. In: *SoftwareX* 22 (2023), p. 101345.
- [77] S. L. Dhulipala et al. *Massively parallel bayesian model calibration and uncertainty quantification with applications to nuclear fuels and materials*. Tech. rep. Idaho National Laboratory (INL), Idaho Falls, ID (United States), 2023.
- [78] C. J. T. Braak. “A Markov Chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces”. In: *Statistics and Computing* 16.3 (2006), pp. 239–249.

Bibliography

- [79] M. AbdulHameed, B. Beeler, C. O. Galvin, and M. W. Cooper. “Assessment of uranium nitride interatomic potentials”. In: *Journal of Nuclear Materials* 600 (2024), p. 155247.
- [80] K. S. Chaudri et al. “Coupled analysis for new fuel design using UN and UC for SCWR”. In: *Prog. Nucl. Energy* 63 (2013), pp. 57–65. eprint: doi:10.1016/j.pnucene.2012.11.001.
- [81] M. W. Cooper et al. “Fission gas diffusion and release for Cr_2O_3 -doped UO_2 : From the atomic to the engineering scale”. In: *Journal of Nuclear Materials* 545 (2021), p. 152590. eprint: doi:10.1016/j.jnucmat.2020.152590.
- [82] K. Forsberg and A. R. Massih. “Kinetics of fission product gas release during grain growth”. In: *Model. Simul. Mater. Sci. Eng.* 15.3 (2007), pp. 335–353. eprint: doi: 10.1088/0965-0393/15/3/011.
- [83] J. Herman and W. Usher. “SALib: An open-source Python library for sensitivity analysis”. In: *Journal of Open Source Software* 2.9 (2017), p. 97.
- [90] C. P. Kempter and R. O. Elliott. “Thermal Expansion of UN, UO_2 , $\text{UO}_2\cdot\text{ThO}_2$, and ThO_2 ”. In: *The Journal of Chemical Physics* 30.6 (1959), pp. 1524–1526. DOI: 10.1063/1.1730230.
- [91] U. C. Nunez, D. Prieur, R. Bohler, and D. Manara. “Melting point determination of uranium nitride and uranium plutonium nitride: A laser heating study”. In: *Journal of Nuclear Materials* 449.1-3 (2014), pp. 1–8. DOI: 10.1016/j.jnucmat.2014.02.021.
- [92] J. H. Harding and D. G. Martin. “A recommendation for the thermal conductivity of UO_2 ”. In: *Journal of Nuclear Materials* 166.3 (1989), pp. 223–226. DOI: 10.1016/0022-3115(89)90218-3.
- [93] K. L. Peddicord, M. E. Cunningham, and A. Tripathi. *Porosity correction to thermal conductivity based on analytical temperature solutions*. Conference Paper. 1978. DOI: IAEA:9418589.
- [94] S. B. Ross, M. S. El-Genk, and R. B. Matthews. “Thermal conductivity correlation for uranium nitride fuel between 10 and 1923 K”. In: *Journal of Nuclear Materials* 151.3 (1988), pp. 318–326. DOI: 10.1016/0022-3115(88)90026-8.

- [95] S. L. Hayes, J. K. Thomas, and K. L. Peddicord. “Material property correlations for uranium mononitride: III. Transport properties”. In: *Journal of Nuclear Materials* 171.2-3 (1990), pp. 289–299. DOI: 10.1016/0022-3115(90)90376-X.
- [96] G. Rossiter. “Understanding and modelling fuel behaviour under irradiation”. In: *Nuclear Fuel Cycle Science and Engineering*. Ed. by W. P. Limited. Woodhead Publishing Limited, 2012, pp. 396–426.
- [97] P. Van Uffelen and M. Suzuki. “3.19: Oxide Fuel Performance Modeling and Simulations”. In: *Comprehensive Nuclear Materials* 3 (2012), pp. 535–577. DOI: 10.1016/B978-0-08-056033-5.00071-9.
- [98] F. J. Huegel. *Swelling of Irradiated UC AND UN*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1965. DOI: 10.2172/4678747.
- [99] A. J. E. Foreman. *Calculations on the Rate of Swelling of Gas Bubbles in Uranium*. Report. Atomic Energy Research Establishment, Harwell, Berks, England, 1956. DOI: OSTI:4369123.
- [100] W. D. Callister. *Fundamentals of Materials Science and Engineering*. Fifth. Vol. 471660817. John Wiley & Sons, Inc., 2000.
- [101] L. Zabudko and V. Poplavsky. “Investigations of Nitride Fuels for Fast Reactors in Russia”. In: *Symposium on Nitride Fuel Cycle Technology*. Vol. 015. Japan Atomic Energy Research Inst., pp. 1–14.
- [102] S. I. Porollo and S. N. Ivanov. “Corrosion Damage and Mechanical Properties of BR-10 Reactor Fuel-Rod Cladding with Mononitride Uranium Fuel”. In: *Atomic Energy* 125 (2019), pp. 185–193. DOI: 10.1007/s10512-018-00464-3.
- [103] S. S. Voss. *SNAP (Space Nuclear Auxiliary Power) Reactor Overview*. Report. Air Force Weapons Lab, Kirtland AFB, NM, 1984.
- [104] S. C. Weaver and J. L. Scott. *Comparison of reactor fuels for high temperature applications*. Report. Oak Ridge National Lab., Oak Ridge, TN, 1965.
- [105] R. E. Gluyas and A. F. Lietzke. *Materials technology program for a compact fast reactor for space power*. Report. NASA Lewis Research Center, Cleveland, OH, 1971.
- [106] J. G. Slaby, B. L. Siegel, L. Gedeon, and R. J. Galbo. *Irradiation of three T-111 clad uranium nitride fuel pins for 8070 hours at 990 C (1815 F)*. Report. NASA Lewis Research Center Cleveland, OH, 1973.
- [107] L. Van Brutzel, R. Dingreville, and T. Bartel. *Nuclear fuel deformation phenomena*. Report. 2015.
- [108] J. G. Slaby and B. L. Siegel. *Examination of T-111 clad uranium nitride fuel pins irradiated up to 13,000 hours at a clad temperature of 990 C*. Report. NASA Lewis Research Center Cleveland, OH, 1973.

- [109] K. R. Thoms. *Design, fabrication, and operation of capsules for the irradiation testing of candidate advanced space reactor fuel pins*. Report. Oak Ridge National Lab., Oak Ridge, TN, 1975.
- [110] R. G. Rohal and T. N. Tambling. *Unrestrained swelling of uranium-nitride fuel irradiated at temperatures ranging from 1100 to 1400 K (1980 to 2520 R)*. Report. NASA Lewis Research Center Cleveland, OH, 1973.
- [111] D. S. Dutt, C. M. Cox, R. A. Karnesky, and M. K. Millhollen. *Performance testing of refractory alloy-clad fuel elements for space reactors*. Report. Hanford Engineering Development Lab., Richland, WA, 1985. DOI: OSTI:5052097.
- [112] R. A. Karnesky and R. E. Mason. *Postirradiation examination results from SP-1*. Report. Hanford Engineering Development Lab., Richland, WA, 1986. DOI: OSTI: 6332376.
- [113] B. J. Makenas, D. M. Paxton, S. Vaidyanathan, M. Marietta, and C. W. Hoth. “SP-100 Fuel Pin Performance: Results from Irradiation Testing”. In: *AIP Conference Proceedings*. Vol. 301. American Institute of Physics, pp. 403–412. DOI: 10.1063/1.2950209.
- [114] D. L. Keller. *Progress on Development of Fuels and Technology for Advanced Reactors During July 1969 - June 1970. Progress Report*. Report. Battelle Memorial Inst., Columbus, OH, 1971. DOI: 10.2172/4083188.
- [115] E. D. Albrecht, A. J. Rothman, J. D. Lee, J. M. Johnson, and W. N. Hayes Jr. *High-Temperature Irradiation and Post-Irradiation Analysis of Uranium Nitride Fuel*. Report. California Univ., Livermore, CA, Lawrence Radiation Lab., 1969. DOI: 10.2172/4733870.
- [116] E. D. Albrecht, A. J. Rothman, J. D. Lee, J. M. Johnson, and W. N. Hayes Jr. *Irradiation studies of uranium nitride*. Report. California Univ., Livermore, CA. Lawrence Radiation Lab., 1969. DOI: OSTI:4768037.
- [117] R. Brennan, M. Golt, and M. Ivill. *Comparison of Dynamic Methods for Determining Elastic Property Measurements of Solid Materials*. Report. DEVCOM Army Research Lab., Aberdeen Proving Ground, MD, 2020.
- [118] S. F. Demuth. “SP100 space reactor design”. In: *Progress in Nuclear Energy* 42.3 (2003), pp. 323–359. DOI: 10.1016/S0149-1970(03)90003-5.
- [119] W. F. Lyon III. “Thermal analysis of the FSP-1 fuel pin irradiation test”. In: *AIP Conference Proceedings*. Vol. 217. American Institute of Physics, pp. 892–901. DOI: 10.1063/1.40104.
- [120] C. P. Kempter, J. C. McGuire, and M. R. Nadler. “Crystallographic data. 177. Uranium mononitride”. In: *Analytical Chemistry* 31.1 (1959), pp. 156–157. DOI: 10.1021/ac60145a601.

- [121] B. J. Makenas, J. W. Hales, and A. L. Ward. “Fuels irradiation testing for the SP-100 program”. In: *AIP Conference Proceedings*. Vol. 217. American Institute of Physics, pp. 886–891. DOI: 10.1063/1.40103.
- [122] M. H. Krasner, H. W. Davison, and A. J. Diaguila. *Conceptual design of a compact fast reactor for space power*. Report. NASA Lewis Research Center Cleveland, OH, 1971. DOI: OSTI:4021520.
- [123] W. Heneveld, R. Paschall, T. Springer, V. Swanson, A. Thiele, and R. Tuttle. *Experimental physics characteristics of a heavy-metal-reflected fast-spectrum critical assembly*. Report. Atomics International Div. North American Rockwell Corp., Canoga Park, CA, 1971. DOI: OSTI:4487619.
- [124] W. Mayo, P. G. Klann, and C. L. Whitmarsh Jr. *Nuclear design and experiments for a space power reactor*. Report. NASA Lewis Research Center Cleveland, OH, 1971.
- [125] R. R. Metroka. *Fabrication of uranium mononitride compacts*. Report. NASA Lewis Research Center Cleveland, OH, 1970. DOI: 10.2172/4674236.
- [126] L. Raring. *Fuels and materials highlights: SNAP-50 and advanced programs*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1965. DOI: 10.2172/4284895.
- [127] W. Doll. *SNAP-50/SPUR program annual technical report, 1964*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1964. DOI: 10.2172/4315637.
- [128] H. Banach and H. Richings. *Status of the SNAP-50 reactor fuels parametric studies. Part I*. Report. Pratt and Whitney Aircraft, Middletown, CT, Connecticut Advanced Nuclear Engineering Lab., 1964. DOI: 10.2172/4330618.
- [129] M. Uno, T. Nishi, and M. Takano. “2.03 - Thermodynamic and Thermophysical Properties of the Actinide Nitrides”. In: *Comprehensive Nuclear Materials* 2 (2012), pp. 61–85. DOI: 10.1016/B978-0-08-056033-5.00010-0.
- [130] R. S. Sen, M. A. Pope, A. M. Ougouag, and K. O. Pasamehmetoglu. “Assessment of possible cycle lengths for fully encapsulated microstructure fueled light water reactor concepts”. In: *Nuclear Engineering and Design* 255 (2013), pp. 310–320. DOI: 10.1016/j.nucengdes.2012.11.007.
- [131] H. J. Matzke. “Point Defects and Transport Properties in Carbides”. In: *Solid State Ionics* 12 (1984), pp. 24–45. DOI: OSTI:6216063.
- [132] D. K. Reimann, D. M. Kroeger, and T. S. Lundy. “Citation self-diffusion in UN(1+x)”. In: *Journal of Nuclear Materials* 38.2 (1971), pp. 191–196. DOI: 10.1016/0022-3115(71)90042-0.

Appendix A

Direct MCMC Implementation with Multi-Variable, Multi-Variance Framework

1.1 Overview and Significance

Initial investigations into direct Bayesian calibration using the INL MOOSE Stochastic Tools framework [76] were undertaken to establish a baseline understanding of computational requirements for the 16-dimensional parameter estimation problem. The Stochastic Tools module provides an ensemble MCMC implementation based on the Affine Invariant Ensemble sampler with Differential Evolution proposals, enabling efficient parallelization across multiple compute nodes through walker-based sampling strategies [77].

The critical algorithmic innovation developed here—and the primary contribution of this appendix—is the multi-variable, multi-variance framework that enables simultaneous calibration against multiple observable quantities (swelling and fission gas release) with independent noise parameters for each observable or observable group. This capability proves essential for realistic inverse problems where different observables carry different measurement uncertainties and represent different physical mechanisms. Rather than forcing all observables to share a single noise parameter or implementing separate single-observable samplers, the unified framework treats multi-observable calibration as a fundamental architectural requirement.

While the direct BISON-coupled approach ultimately proved computationally impractical, the multi-variable architecture developed here directly informed and enabled the surrogate-accelerated calibration methodology adopted in Chapter 7. The insights regarding multi-variable likelihood structure, multi-variance state management, and convergence-aware decision logic proved invaluable for building the surrogate-based system.

1.2 Affine Invariant Ensemble Sampler with Multi-Variance Support

1.2.1 Foundational Sampler

The affine invariant ensemble (AIE) sampler extends the Metropolis-Hastings algorithm with ensemble-based proposal generation [78]. Rather than proposing parameter updates independently for each chain, AIE maintains an ensemble of N_w walkers that collectively explore the posterior. New proposals are generated by computing differential vectors between random walker pairs, producing proposals that maintain affine invariance—a property ensuring that proposals scale appropriately with the local curvature of the posterior. The proposal for walker i at iteration t takes the form

$$\boldsymbol{\theta}_i^{\text{prop}} = \boldsymbol{\theta}_i^t + \gamma (\boldsymbol{\theta}_a^t - \boldsymbol{\theta}_b^t) + \epsilon, \quad (\text{A.1})$$

where $\boldsymbol{\theta}_a^t$ and $\boldsymbol{\theta}_b^t$ are states from randomly selected walkers ($a \neq b \neq i$), $\gamma = 2.38/\sqrt{2d_\theta}$ represents a tuning parameter depending on parameter dimensionality d_θ [78], and $\epsilon \sim \mathcal{N}(0, b^2)$ adds calibrated noise with $b = 10^{-6} \times \text{scale}$ to ensure walkers remain distinct. The walker-dependent proposal structure exploits multi-walker parallelism naturally, enabling efficient use of ensemble-based sampling on high-performance computing systems.

1.2.2 Multi-Variance Architecture

The critical algorithmic contribution extending the standard AIE framework is the introduction of independent variance parameter sampling. Standard ensemble MCMC assumes a single noise parameter σ shared across all observations. Multi-observable calibration problems demand independent noise parameters $\boldsymbol{\sigma} = \{\sigma_1, \sigma_2, \dots, \sigma_K\}$ corresponding to K distinct observable groups (e.g., swelling and fission gas release measurements).

The multi-variance extension maintains parallel proposal chains for both model parameters $\boldsymbol{\theta}$ and variance parameters $\boldsymbol{\sigma}$. For each walker i , the sampler maintains

$$\mathbf{X}_i^t = [\boldsymbol{\theta}_i^t, \boldsymbol{\sigma}_i^t] \in \mathbb{R}^{d_\theta + d_\sigma}, \quad (\text{A.2})$$

where $d_\theta = 16$ parameters govern fuel behavior and $d_\sigma = 2$ variance parameters correspond to swelling and fission gas release noise, respectively. The extended state space doubles the dimensionality but maintains the affine invariance property by applying differential evolution to both parameter and variance components through unified proposal mechanisms.

Variance proposals follow the same affine-invariant structure as parameter proposals. For variance component j of walker i ,

$$\sigma_{i,j}^{\text{prop}} = \sigma_{i,j}^t + \gamma (\sigma_{a,j}^t - \sigma_{b,j}^t) + \epsilon_\sigma, \quad (\text{A.3})$$

where the tuning parameter γ is recomputed based on total state dimensionality ($d_\theta + d_\sigma$), and ϵ_σ represents noise scaled to variance parameter ranges. This unified treatment ensures

that variance parameters explore the posterior using the same statistical machinery as model parameters rather than through ad-hoc perturbation schemes.

1.3 Multi-Observable Extension and Variable Grouping

1.3.1 Multi-Variable Gaussian Likelihood with Grouped Noise

The central architectural innovation enabling simultaneous calibration of multiple observables is a flexible multi-variable Gaussian likelihood function that supports three distinct noise assignment strategies:

1. **Shared Noise Mode:** All K observables share a single noise parameter σ , reducing parameterization but assuming equal measurement uncertainty across distinct physical quantities
2. **Per-Variable Noise Mode:** Each observable k receives independent noise σ_k , enabling differential uncertainty quantification for swelling (σ_{swell}) and fission gas release (σ_{fgr})
3. **Per-Group Noise Mode:** Variables are clustered into groups by physical mechanism or measurement origin. For example, all swelling observations from different tests could share a common σ_{swell} , while maintaining separate σ_{fgr} for gas release

For the uranium mononitride fuel model, per-variable noise mode is most appropriate because swelling and fission gas release are fundamentally distinct physical processes (void nucleation/growth versus bubble migration/release) measured with different instrumental precisions and model approximation errors.

1.3.1.1 Sample-Major Data Organization

The likelihood function organizes multi-observable model predictions in sample-major layout, where all observables for a single experimental configuration are grouped together:

$$\mathbf{y}_{\text{model}} = [y_0^{\text{swell}}, y_0^{\text{fgr}}, y_1^{\text{swell}}, y_1^{\text{fgr}}, \dots, y_{N_{\text{obs}}-1}^{\text{swell}}, y_{N_{\text{obs}}-1}^{\text{fgr}}], \quad (\text{A.4})$$

where subscript denotes observation index and superscript denotes observable type. This organization enables vectorized likelihood computation while maintaining cache efficiency on modern CPUs. Experimental observations $\mathbf{y}_{\text{exp}}$ follow identical organization, with corresponding noise parameters organized as

$$\boldsymbol{\sigma}_{\text{model}} = [\sigma_{\text{swell}}, \sigma_{\text{fgr}}, \sigma_{\text{swell}}, \sigma_{\text{fgr}}, \dots, \sigma_{\text{swell}}, \sigma_{\text{fgr}}]. \quad (\text{A.5})$$

For each experimental test c and sample (model evaluation) s , the multi-variable Gaussian likelihood aggregates across both observables:

$$\log p(\mathbf{d}_s^c \mid \boldsymbol{\theta}, \boldsymbol{\sigma}) = \sum_{k=1}^K \log \mathcal{N}(d_{s,k}^c; \mu_k(\boldsymbol{\theta}), \sigma_k^2), \quad (\text{A.6})$$

where $d_{s,k}^c$ represents experimental observation k for sample s in test configuration c , $\mu_k(\boldsymbol{\theta})$ denotes the model prediction via surrogate or BISON evaluation, and σ_k represents the independent noise parameter for observable k .

1.3.1.2 Multi-Variable Reporter and Data Aggregation Pipeline

A critical infrastructure component supporting multi-variable calibration is the Multi-VariableReporter, which aggregates disparate BISON model outputs into the sample-major layout required by the likelihood function. BISON evaluations produce swelling and fission gas release predictions separately through distinct output interfaces. The reporter combines these outputs with precise ordering:

1. For each walker (28 walkers total), BISON produces output predictions for each observable at each experimental configuration
2. MultiVariableReporter extracts current iteration's outputs (avoiding duplication from previous iterations)
3. Outputs are organized in sample-major order: $[\text{sample}_0^{\text{obs}_0}, \text{sample}_0^{\text{obs}_1}, \text{sample}_1^{\text{obs}_0}, \text{sample}_1^{\text{obs}_1}, \dots]$
4. This aggregated vector is passed directly to the multi-variable likelihood function

The reporter interface tracks previous iteration size to compute only incremental new data, avoiding memory overhead from storing entire history. This design proves essential for efficient communication between BISON's output reporters and the MCMC decision module across large-scale parallel systems.

1.3.2 Multi-Variance State Management and Prior Specification

The extended state space (A.2) requires rigorous state persistence mechanisms across MCMC iterations. The sampler maintains separate vector structures:

1. **Parameter state:** `_new_samples[walker][parameter]` containing proposed model parameters
2. **Variance state (single-variance mode):** `_new_var_samples[walker]` for scalar variance parameter

3. **Variance state (multi-variance mode):**

`_new_var_samples_vec`[walker][variance_group] for independent noise per observable

4. **Previous parameter state:** `_previous_state`[walker][parameter] retained for rejection handling

5. **Previous variance state:** `_previous_state_var_multi`[variance_group][walker] or `_previous_state_var_single`[walker] for variance reversion

Prior distributions follow the extended state dimensionality. Model parameters receive uniform priors on normalized interval $p(\theta_i) = \text{Uniform}(0, 1)$ for $i = 1, \dots, 16$, reflecting weak prior information. Variance parameters receive per-group half-normal priors

$$p(\sigma_k) = \text{HalfNormal}(\mu = 0, \sigma = 0.05), \quad k = 1, 2, \quad (\text{A.7})$$

with $\mu_k \approx 0.05$ to 0.08 depending on observable. This prior admits large posterior variability while concentrating prior mass near realistic measurement uncertainty scales (5-10% relative noise in log-space).

The multi-variance state management preserves complete backward compatibility with legacy single-variance samplers through conditional logic. When `variance_names` parameter is empty, the sampler operates in legacy mode using `_new_var_samples`[walker] and `_var_prior`. When `variance_names` is populated with observable group names, multi-variance mode activates, using `_new_var_samples_vec` and `_var_priors` vector structures.

1.3.3 Multi-Variable Decision Logic with Convergence Masking

The acceptance-rejection decision stage represents the computational bottleneck where model evaluations occur and MCMC convergence assessment happens. For each walker i , the log-space Metropolis-Hastings ratio incorporates extended state components:

$$\log \alpha_i = \min \left(0, \log \frac{p(\mathbf{d} \mid \boldsymbol{\theta}^{\text{prop}}, \boldsymbol{\sigma}^{\text{prop}}) p(\boldsymbol{\theta}^{\text{prop}}) p(\boldsymbol{\sigma}^{\text{prop}})}{p(\mathbf{d} \mid \boldsymbol{\theta}^{\text{prev}}, \boldsymbol{\sigma}^{\text{prev}}) p(\boldsymbol{\theta}^{\text{prev}}) p(\boldsymbol{\sigma}^{\text{prev}})} \right), \quad (\text{A.8})$$

where the posterior likelihood spans all observables through Equation (A.6), and prior contributions separate into model parameter priors and variance parameter priors. The evidence calculation decomposes into three additive components in log-space:

$$\Delta \log p = \underbrace{\sum_{j=1}^{d_\theta} [\log p(\theta_j^{\text{prop}}) - \log p(\theta_j^{\text{prev}})]}_{\text{Parameter priors}} + \underbrace{\sum_{k=1}^K [\log p(\sigma_k^{\text{prop}}) - \log p(\sigma_k^{\text{prev}})]}_{\text{Variance priors}} + \underbrace{\Delta \log L}_{\text{Likelihood}}, \quad (\text{A.9})$$

enabling targeted debugging and diagnostics. If any component (e.g., likelihood) dominates due to numerical issues, decomposition isolates the problem.

1.3.3.1 Convergence-Based Automatic Rejection

A critical capability developed for realistic multi-observable problems is convergence-aware sample rejection. BISON’s nonlinear solver occasionally fails to converge at extreme burnup ($>4\%$ FIMA) or unusual parameter combinations. Rather than attempting inference on unconverged simulation data—which contains undefined or extrapolated numerical values—the framework automatically masks unconverged samples through convergence-status tracking.

For each sample (c, s, w) (configuration c , sample s , walker w), a convergence indicator $\text{conv}_{c,s,w} \in \{0, 1\}$ is computed by BISON and communicated through the convergence reporter. In the decision module, if any configuration for walker w reports non-convergence, the transition probability is automatically set to zero:

$$\alpha_i = \begin{cases} \text{Metropolis-Hastings ratio from Eq. (A.8)} & \text{if } \text{conv}_{c,s,i} = 1 \text{ for all } c \\ 0 & \text{if } \text{conv}_{c,s,i} = 0 \text{ for any } c \end{cases} \quad (\text{A.10})$$

This mechanism ensures that even if model evaluation produces numerical output, the unconverged proposal is automatically rejected and the walker reverts to its previous accepted state. This conservative approach maintains sample quality at the cost of occasionally rejecting otherwise plausible proposals.

1.3.4 State Reversion and Output Reorganization

Upon proposal rejection (Metropolis-Hastings acceptance probability below random threshold), the sampler must carefully revert all state components. The reversion operation restores both parameter and variance vectors to their previous accepted values, and crucially, restores all model output predictions to prevent downstream likelihood evaluations from using mismatched parameter-output pairs.

For multi-variable systems, output reversion organizes data by configuration and observable:

$$\mathbf{y}_{\text{revert}}[c \times N_w \times K + w \times K + k] = \mathbf{y}_{\text{prev}}[c \times N_w \times K + w \times K + k], \quad (\text{A.11})$$

where configuration index c ranges over test cases, walker index w ranges over ensemble members, observable index k ranges over output types (swelling, gas release), and $N_w = 28$ walkers. This organization ensures that rejected proposals cannot contaminate subsequent likelihood evaluations through output misalignment.

1.4 Computational Performance and Scaling Analysis

1.4.1 Deployment Configuration

The extended multi-variable, multi-variance framework was deployed on a high-performance computing cluster with 28 walkers operating in parallel, 12 dedicated cores per walker, for a total of 336 cores allocated to a single MCMC ensemble. Each BISON simulation required 10 to 30 minutes depending on burnup and temperature conditions, with significant variability arising from nonlinear solver convergence behavior at extreme burnup values. The extended state space (16 parameters + 2 variance parameters totaling 18 total degrees of freedom) increased proposal dimensionality by 12.5% compared to parameter-only sampling, but the affine invariance property maintained tuning parameter relationships.

Over a 7-day maximum-allocation continuous run, the framework completed approximately 400 iterations of the AIE sampler. This computational investment corresponds to approximately 95 core-days of computing resources per MCMC iteration. The multi-variance overhead—including separate proposal generation for variance parameters, independent state tracking, and per-observable likelihood evaluation—added approximately 5-10% computational cost compared to single-variance implementations, a modest penalty for the flexibility gained.

1.4.2 Computational Cost Analysis and Dimensionality Scaling

The 95 core-day per iteration cost breaks down as follows across system components. For a typical iteration involving 28 walkers making 1 to 3 proposals each (depending on rejection rate), approximately 50 to 60 model evaluations occur. Each evaluation requires 10 to 30 minutes depending on operating conditions, averaging 18 minutes in the experimental burnup range. This yields $50 \text{ evaluations} \times 18 \text{ minutes} = 900 \text{ core-minutes}$, or approximately 15 core-hours per iteration when accounting for solver convergence behavior and I/O overhead.

The remaining computational overhead includes MCMC bookkeeping (state management for both parameters and variances, multi-variable likelihood computation, convergence diagnostics), which consumes approximately 10% of total runtime due to vectorized operations and in-core caching. Parallel communication overhead (MPI synchronization between walkers) consumes approximately 5% when using efficiently parallelized MOOSE communicators, but can spike to 20% when filesystem I/O limits throughput. The measured 95 core-days per iteration reflects average performance across multiple 7-day allocations with typical HPC contention and I/O bottleneck patterns.

Scaling to the 16-dimensional extended state space (16 parameters + 2 variances) increases the tuning parameter $\gamma = 2.38/\sqrt{2 \times 18} = 0.397$, compared to $\gamma = 0.410$ for pure 16-dimensional parameter space. This modest change preserves acceptance rates in the target 25-30% range for ensemble MCMC methods [72].

1.4.3 Fundamental Barriers to Direct 16+2 Dimensional Inference

While the multi-variance framework proves elegant and statistically sound, fundamental computational barriers render direct BISON-coupled calibration impractical even with parallel optimization. Achieving 95% confidence coverage of the 16-dimensional posterior support requires approximately 2 to 3 million samples per effective dimension, or roughly 36 to 54 million total samples accounting for correlation structure. With 28-walker parallelism and typical 25% acceptance rate for differential evolution proposals, this corresponds to 130 to 200 million total proposals, demanding 6 to 10 billion BISON evaluations.

This vastly exceeds both practical dissertation timescales and standard HPC allocation limits (typically 30 to 100 core-days for individual research problems). The dimensional barrier compounds with model complexity: each dimension increases difficulty exponentially through the curse of dimensionality in Bayesian inference.

1.5 Model Complexity and Nonlinear Solver Challenges

1.5.1 Physical Nature of Computational Barriers

The computational barrier stems not from algorithmic deficiency (the multi-variable, multi-variance framework is statistically rigorous and the ensemble MCMC implementation is sound) but rather from the inherent expense of mechanistic modeling. Each BISON model evaluation involves solving coupled nonlinear PDEs across 2D fuel pin geometry with temperature-dependent properties and state-dependent source terms. The coupled physics includes heat conduction (temperature), gas diffusion (species concentration), and mechanical deformation (stress-strain evolution) across the full pin cross-section. Temperature-dependent material properties require root-finding and Newton iterations at each quadrature point.

Nonlinear solver convergence becomes particularly challenging at extreme burnup (3.5%+ FIMA) where high-burnup acceleration mechanisms introduce bifurcations and discontinuities in material response. The framework's convergence masking capability (Section 1.3) handles these failures gracefully, but their frequency increases solver time variability and reduces effective sampling efficiency.

While advanced implicit solvers (PETSC, Trilinos) and HPC parallelization (up to 336 cores deployed here) have reduced individual evaluation time to 10 to 30 minutes, this fundamental limitation cannot be overcome through pure algorithmic or computational improvements without abandoning mechanistic fidelity.

1.6 Transition to Surrogate Acceleration and Knowledge Transfer

The direct MCMC framework, while scientifically sound and representing sophisticated algorithmic development, conclusively demonstrated that direct BISON-coupled inference remains impractical for routine fuel model calibration. The 28,500 to 47,500 core-day computational requirement renders the approach infeasible within dissertation timescales and standard HPC resource availability.

This computational barrier motivated the transition to surrogate-based acceleration documented in Chapter 7. By replacing expensive BISON evaluations with trained Gaussian Process emulators during MCMC sampling, computational demands reduce from tens of thousands of core-days to 8 to 12 hours, enabling practical Bayesian calibration while maintaining statistical rigor through explicit treatment of surrogate approximation uncertainty in the likelihood function.

The algorithmic innovations developed for direct MCMC proved invaluable for surrogate-based design:

1. **Multi-variable likelihood architecture:** The sample-major data organization (Equation (A.4)) and multi-observable likelihood formulation (Equation (A.6)) directly translated to the surrogate-accelerated framework without modification. Surrogates replaced BISON calls but retained identical likelihood structure.
2. **Per-variable noise parameterization:** The multi-variance framework enabling independent σ_{swell} and σ_{fgr} proved essential for surrogate-based calibration, allowing realistic treatment of differential measurement uncertainties between distinct physical observables.
3. **Convergence awareness:** The convergence masking strategy (Equation (A.10)) informed data quality assessment in Chapter 7. While surrogate models don't directly call BISON, understanding which parameter combinations cause solver failure influenced outlier detection and data refinement strategies.
4. **Parallel architecture understanding:** The multi-walker ensemble structure with 28 walkers informed the multi-chain NUTS strategy employed in surrogate-accelerated calibration. Ensemble parallelism remains central to both direct and surrogate implementations.
5. **State management insights:** The careful management of parameter and variance state vectors, previous states, and output organization directly informed the state-handling logic in surrogate-accelerated decision modules.
6. **Regime-specific calibration strategy:** Understanding the computational cost of high-dimensional inference motivated the regime-specific calibration approach,

partitioning the 16-dimensional problem into lower-dimensional regime-specific sub-problems with improved MCMC mixing and convergence properties.

The direct MCMC implementation thus serves not merely as a cautionary computational tale, but as a rigorous architectural foundation enabling the sophisticated multi-objective Bayesian inference methodology ultimately adopted for this dissertation.

1.7 Summary and Scientific Significance

The direct MCMC implementation with multi-variable, multi-variance framework represents a significant algorithmic contribution to ensemble-based Bayesian inference for expensive computational models. The framework enables simultaneous calibration of multiple observables with independent noise parameters, convergence-aware sample rejection, and rigorous state management across extended parameter and variance dimensions. While computational barriers ultimately necessitated surrogate acceleration, the scientific and engineering insights developed through this implementation directly enabled the successful surrogate-accelerated calibration methodology that forms the technical core of this dissertation’s uncertainty quantification framework.