

Improving Fire Resistance Prediction of Glulam Timber Columns Using Gaussian Process Surrogate Modelling

A Masters thesis by
J.D. van der Wulp

Delft University of Technology

Improving Fire Resistance Prediction of Glulam Timber Columns Using Gaussian Process Surrogate Modelling

by

J.D. van der Wulp

to obtain the degree of Master of Science
at the Delft University of Technology.

Student number: 5009790
University: Delft University of Technology
Faculty: Faculty of Civil Engineering and Geosciences
Department: Materials, Mechanics, Management & Design (3MD)
Project duration: August 11, 2025 – February 27, 2026
Thesis committee: Prof. dr. ir. L.J. Sluys,
Dr. I.B.C.M. Rocha,
Dr. Z. Nan,
Prof. dr. ir. E.O.L Lantsoght,
Ir. M.M.J. Spanenburg,

TU Delft, Supervisor
TU Delft, Supervisor
TU Delft, Supervisor
TU Delft, USFQ, Supervisor
BAM Advies & Engineering,
Supervisor

In collaboration with BAM Advies & Engineering.

Cover: T3 West Midtown in Atlanta. Photo by Rion Rizzo.
(Image source: DLR Group, 2021)

Preface

This thesis marks the end of my journey in the Civil Engineering Master's programme at Delft University of Technology. In this project, I investigated the delayed failure behaviour and burnout resistance of glulam timber columns exposed to fire scenarios that include both the heating and cooling phases. In my view, there is a need for the construction industry to become more sustainable, and I am convinced that timber buildings have an important role to play in that transition. At the same time, fire safety challenges in mid-rise and high-rise timber buildings can still constrain the decision to adopt timber solutions, which is exactly why I found it exciting to contribute to a topic that may help make these applications more feasible in the future.

Over the course of this thesis project, I particularly enjoyed combining the different fields from my Master's degree. Coming from a structural engineering track, supported by modules in building engineering, mechanics and dynamics, and complemented by a cross-over module in data science and artificial intelligence, I had the opportunity to apply this knowledge within a single project. I found it both challenging and rewarding to integrate these foundations with the disciplines of structural fire engineering, structural mechanics, and machine learning within one coherent study. The topic also confronted me with the inherent complexity of delayed failure, where non-linearities across thermal exposure, material behaviour, and structural response are central rather than incidental, making the work demanding but also intellectually engaging.

I would like to express my sincere gratitude to the members of my assessment committee for their support and guidance throughout this project. First, I would like to thank Dr. Zhuojun Nan, my daily supervisor, for the continuous support and constructive feedback. I greatly appreciate our weekly meetings on the technical details and the broader direction of the research, which made the thesis process both clearer and more enjoyable. Next, I would like to express my great appreciation to Dr. Iuri Rocha for the guidance and the in-depth discussions on machine learning and surrogate modelling, which proved to be of indispensable value throughout the project. Furthermore, I am sincerely grateful to Prof. Dr. Ir. Eva Lantsoght for joining the project from abroad and for taking the time to provide sharp and comprehensive feedback that helped to refine this thesis further. Finally, I would like to thank Prof. Dr. Ir. Bert Sluys, as chair of the assessment committee, for overseeing the thesis process. The guidance from the start of the project and the freedom given to follow my own research path are very much appreciated.

I would also like to express my gratitude to my company supervisor, Ir. Mark Spanenburg, for the guidance and encouragement throughout the project. I really appreciate the support during the more challenging and stressful phases of the thesis. It was a pleasure to work together. I am also grateful to BAM Advies & Engineering for their support throughout the project and for the trust and freedom I was given to shape the direction of the work.

In addition, I am grateful to the researchers with whom I had the opportunity to be in contact during this project, in particular Dr. Jingxian Zhao, Dr. Jeremy Chang, Dr. Cheng Chen, and Dr. Thomas Gernay, for their time and valuable input.

Lastly, I am grateful to the SFPE Foundation for awarding me a Student Research Grant in support of this project (Grant No. 26-SRG-001-01; DOI: 10.64167/r1ey-1ngc). I greatly appreciate their support and the recognition of the work, which provided an important opportunity and motivation to further develop and carry out the research.

*J.D. van der Wulp
Delft, February 2026*

Summary

The growing global interest in mass timber construction has increased the need for reliable assessment of structural fire performance beyond regulated fire resistance ratings. Current Eurocode fire resistance provisions are primarily based on standard fire tests (e.g. ISO 834 [1]), which consider the heating phase only. However, experimental studies have shown that glued laminated timber (glulam) columns may fail during the cooling phase of natural fires, a phenomenon known as *delayed failure* [2]. This phenomenon is particularly relevant for glulam timber columns due to the combustible nature of the material, slow heat penetration towards the core caused by low thermal conductivity, and rapid degradation of strength and stiffness at relatively low temperatures. To capture structural performance over the complete fire and post-fire duration, the metric *burnout resistance* is therefore introduced, which evaluates whether a member maintains its load-bearing capacity throughout the full fire event, leading to a final state of either failure or no failure.

Since current design provisions are largely established for standardised fires, it remains uncertain whether commonly used thermo-mechanical models can reliably reproduce delayed failure under realistic natural fire exposures that include cooling phases. Moreover, the introduction of a metric such as burnout resistance raises the practical question of how it can be assessed efficiently across a broad parameter space to characterise its behaviour and to identify the conditions under which it becomes governing.

Motivated by these challenges, this thesis aims to improve the understanding and prediction of delayed failure behaviour and burnout resistance of glulam timber columns exposed to natural fire scenarios, including both the heating and cooling phases. Specifically, it (i) evaluates the applicability of non-linear thermo-mechanical finite element modelling in SAFIR for describing delayed failure of glulam timber columns under realistic fire exposures, and (ii) investigates Gaussian Process Classification (GPC) as a surrogate modelling framework to enable efficient parametric assessment of burnout resistance.

The thesis adopts a combined numerical–surrogate framework. A modelling framework is developed in the non-linear finite element software SAFIR for isolated spruce glulam (GL24h) columns with square cross-sections, four-sided fire exposure, and eccentric axial loading. Numerical predictions are compared with available full-scale experimental data from natural fire tests. The modelling framework is defined parametrically to support subsequent surrogate modelling using Gaussian Process Classification (GPC) and an active learning strategy. The GPC is trained on binary failure outcomes from the parametric simulations, while active learning is used to target additional simulations near the estimated decision boundary. The surrogate models are constructed progressively from low-dimensional to six-dimensional formulations, considering structural, fire, and material variables.

The finite element simulations are found to reproduce the observed response during the heating phase with reasonable agreement. However, predictions of delayed failure become less reliable when collapse occurs late in the cooling phase. This shortcoming is primarily linked to the thermal modelling, which does not sufficiently capture sustained elevated internal temperatures during the cooling phase of the fire. When the internal temperature histories are matched more closely to measurements, the predicted failure times show better alignment with the experimental observations. The surrogate modelling indicates that GPC is able to approximate complex failure boundaries obtained from deterministic simulations and can enable efficient exploration of interacting variables in higher-dimensional parameter spaces, provided that the underlying numerical model is sufficiently accurate.

These findings imply that timber fire performance cannot be assessed solely using standard fire resistance characteristics or charring depth assessments when cooling phases are included. They further indicate a need for improved thermal modelling of timber during decay and cooling. Finally, the results illustrate that a validated numerical–surrogate workflow is able to support design-oriented evaluation of burnout resistance across realistic fire scenarios.

Contents

Preface	i
Summary	ii
List of Figures	vii
List of Tables	ix
Nomenclature	x
1 Introduction	1
1.1 Background and Motivation	1
1.2 Problem Description	2
1.3 Research Gaps	3
1.4 Research Method	4
1.5 Research Objectives and Questions	4
1.6 Scope and Limitations	5
1.7 Thesis Outline	5
2 Theory	7
2.1 Fire Behaviour of Structural Timber	7
2.1.1 Fire-Induced Degradation Processes	7
2.1.2 Heat Transfer Mechanisms	9
2.1.3 Thermal and Mechanical Properties	10
2.2 Fire Exposure and Standard Fire Curves	12
2.2.1 Standard ISO 834 Fire Curve	12
2.2.2 Parametric Fires	13
2.2.3 Mass Timber Compartment Fires	14
2.3 SAFIR Finite Element Method	14
2.3.1 Thermal Analysis	15
2.3.2 Mechanical Analysis	17
2.4 Gaussian Process Classification	19
2.4.1 Gaussian Processes	20
2.4.2 Formulation of Gaussian Process Classification	21
2.4.3 Active Learning for Gaussian Process Classification	23
3 Literature Review	27
3.1 Delayed Failure of Glulam Timber Columns	27
3.2 Burnout Resistance and the DHP Concept	30
3.3 Delayed Failure Modelling of Glulam Timber Columns	30
4 Methodology	32
4.1 Research Approach	32
4.1.1 General Research Workflow	33
4.1.2 General Modelling Framework	34
4.1.3 Justification of the Modelling Approach	35
4.2 SAFIR Numerical Modelling Framework	36
4.2.1 Fire Curve Definition	37
4.2.2 Thermal Analysis	39
4.2.3 Mechanical Analysis	44
4.3 GPC Surrogate Modelling Framework	47
4.3.1 GPC Surrogate Model Setup	48

4.3.2	Setup of the Active Learning Algorithm	49
4.3.3	GPC Surrogate Model Evaluation	53
5	Results	55
5.1	Numerical Modelling Results	55
5.1.1	Standard ISO 834 Fire Exposure	55
5.1.2	Standardised Parametric Fire Exposure	57
5.1.3	Natural Fire Exposure	58
5.2	GPC Surrogate Modelling Results	66
5.2.1	6D Structural Design and Fire Variables	66
5.2.2	6D Material Variables	72
6	Discussions	78
6.1	Numerical Modelling Discussions	78
6.1.1	Interpretation of the Numerical Modelling Results	78
6.1.2	Limitations of the Numerical Modelling	80
6.2	GPC Surrogate Modelling Discussions	81
6.2.1	Interpretation of the GPC Surrogate Modelling Results	82
6.2.2	Limitations of the GPC Surrogate Modelling	83
6.3	Implications for Fire Engineering	84
7	Conclusions and Recommendations	86
7.1	Numerical Modelling Conclusions	87
7.2	GPC Surrogate Modelling Conclusions	87
7.3	Recommendations for Further Research	88
7.3.1	Numerical Modelling Recommendations	88
7.3.2	GPC Surrogate Modelling Recommendations	89
	References	90
A	Alternative Research Approaches	93
B	Symmetry Mesh Modelling	97
C	Thermal Mesh Sensitivity Analysis	99
D	Numerical Restart Algorithm	101
E	Simplified Analytical Approach	107
F	Charring Depth Validation	113
G	Moment Initialisation	118
H	Mechanical Mesh Sensitivity Analysis	120
I	Plastic Compression Strain	122
J	GPC AL: 2D Trial Functions	125
K	GPC AL: 3D Trial Functions	134
L	GPC AL: 6D Trial Function	139
M	GPC AL: 12D Trial Function	143
N	Bounds Determination	149
O	Results: Monitoring Metrics	153

List of Figures

1.1	Ascent MKE building in Milwaukee, Wisconsin	1
2.1	Schematic representation of in-depth fire-induced degradation processes and internal zones in timber during fire exposure from Bartlett et al. [19].	8
2.2	Temperature-dependent thermal properties of timber according to EN 1995-1-2 [1].	11
2.3	Temperature-dependent reduction factors for the mechanical properties of timber according to EN 1995-1-2 [1].	12
2.4	ISO 834 standard fire curve as defined in EN 1991-1-2 [20].	13
2.5	Illustrative Eurocode parametric fire curves based on Annex A of EN 1991-1-2 [20].	14
2.6	Comparison of experimentally measured gas temperature histories in mass timber compartments with standard and parametric fire curves [13].	15
2.7	SAFIR thermal analysis iteration process [22].	17
2.8	SAFIR mechanical analysis iteration process [25].	19
2.9	Graphical representation of a Gaussian Process with inputs x , latent function values $f(x)$, and observed variables y [26].	21
2.10	Logit and Probit sigmoid-shaped link functions	22
2.11	Laplace and EP approximation to a non-Gaussian posterior [27].	22
2.12	EP approximation of a sigmoid likelihood term in Gaussian Process Classification [26].	23
3.1	Cross-sectional temperature field for ISO 834 fire failure and delayed failure during the cooling phase from Gernay [12].	28
3.2	Isotherm depths for different time instances from Gernay [12].	28
3.3	Overview of experimental furnace tests on glulam timber columns by Gernay et al. [31]	29
3.4	Overview of natural fire tests on glulam timber columns as illustrated in Renard et al. [33]	29
3.5	Illustration of the DHP principle by Robert et al. [30].	30
3.6	Iterative procedure for determining the burnout resistance of a structural member as mentioned in Gernay [34].	30
4.1	Methodological phases	33
4.2	General modelling framework	35
4.3	Numerical simulation framework	37
4.4	Definition of the fire curve parameters.	38
4.5	Effect of the Γ -factor on the fire curve heating phase.	39
4.6	Mesh definition of the two-dimensional cross-section with refined corner regions.	41
4.7	Approximated fire curve used as thermal boundary condition in the simplified analytical approach.	42
4.8	Weibull distribution reduction factors for compression and tensile strength [46].	45
4.9	Constitutive stress–strain relationship for timber used in the mechanical analysis.	47
5.1	Comparison of thermal numerical simulations against experimental data for Test 10 by Renard et al. [32].	59
5.2	Comparison of thermal numerical simulation against experimental data for Test 10 of Renard et al. [32].	60
5.3	Comparison of thermal numerical simulation against experimental data for Test 11 of Renard et al. [32].	61
5.4	Comparison of thermal numerical simulation against experimental data for Test 15 of Renard et al. [32].	62
5.5	Comparison of thermal numerical simulation against experimental data for Test 17 of Renard et al. [32].	63

5.6	Comparison of thermal numerical simulation against experimental data for Test 9 of Renard et al. [32].	64
5.7	Comparison of thermal numerical simulation against experimental data for Test 12 of Renard et al. [32].	65
5.8	Two-dimensional GPC surrogate model for (t_h, B) , illustrating the decision boundary between failure (red) and no failure (blue).	67
5.9	Three-dimensional GPC surrogate model for (t_h, B, F) , illustrating the two-dimensional decision boundary surface separating failure (red) and non-failure (blue).	68
5.10	Two-dimensional slices of the three-dimensional GPC surrogate model for (t_h, B, F) , illustrating the projected decision boundaries between failure (red) and non-failure (blue).	68
5.11	Comparison of the decision boundary in the (t_h, B) plane obtained from the six-dimensional surrogate model for structural design and fire variables and the originally trained two-dimensional model.	69
5.12	Comparison of the decision boundary in the (t_h, F) plane obtained from the six-dimensional surrogate model for structural design and fire variables and the trained three-dimensional model.	70
5.13	Comparison of the decision boundary in the (F, B) plane obtained from the six-dimensional surrogate model for structural design and fire variables and the trained three-dimensional model.	70
5.14	Comparison of three-dimensional decision boundary surfaces in the (t_h, B, F) space obtained from the six-dimensional surrogate model for structural design and fire variables and the trained three-dimensional model.	71
5.15	2D-slices of the (t_h, B) plane with quarter fractiles for fixed variables obtained from the six-dimensional surrogate model for structural design and fire variables.	72
5.16	Two-dimensional GPC surrogate model for (t_h, ρ) , illustrating the decision boundary between failure (red) and no failure (blue).	73
5.17	Three-dimensional GPC surrogate model for (t_h, ρ, E) , illustrating the decision boundary surface separating failure (red) and non-failure (blue).	74
5.18	Two-dimensional slices extracted from the three-dimensional GPC surrogate model for (t_h, ρ, E) , illustrating projected decision boundaries between failure (red) and non-failure (blue).	74
5.19	Two-dimensional slices extracted from the six-dimensional GPC surrogate model incorporating material variables, illustrating projected decision boundaries between failure (red) and non-failure (blue).	75
5.20	Three-dimensional slice extracted from the six-dimensional GPC surrogate model incorporating material variables for (t_h, ρ, E)	76
5.21	Stochastic evaluation of the six-dimensional GPC surrogate model incorporating material variability, showing the mean decision boundary and 5% fractile bounds.	76
6.1	Workflow for implications into fire engineering	85
A.1	Parametric PINN approach	94
A.2	Surrogate Neural Network approach	94
A.3	Inverse PINN approach	95
D.1	Position of nodes in the cross-section for temperature development measurements.	102
D.2	Restart algorithm for central node	102
D.3	Errors of restart algorithm for the central node	103
D.4	Behaviour of restart algorithm for cooling for node 316	104
D.5	Restart algorithm for ISO834 heating	105
E.1	Approximated fire curve used as thermal boundary condition in the simplified analytical approach.	108
F.1	Fire curve and calculated versus measured charring depths for Test 9.	113
F.2	Fire curve and calculated versus measured charring depths for Test 10.	114
F.3	Fire curve and calculated versus measured charring depths for Test 12.	114

F.4	Fire curve and calculated versus measured charring depths for Test 14.	115
F.5	Fire curve and calculated versus measured charring depths for Test 15.	115
F.6	Fire curve and calculated versus measured charring depths for Test 17.	116
F.7	Fire curves and calculated versus measured charring depths for Test 18.	116
J.1	GPC prediction field for the two-dimensional nonlinear boundary trial function.	126
J.2	Monitoring metrics for the two-dimensional nonlinear boundary trial function.	127
J.3	GPC prediction field for the two-dimensional parabola boundary trial function.	128
J.4	Monitoring metrics for the two-dimensional parabola boundary trial function.	129
J.5	GPC prediction field for the two-dimensional square-root boundary trial function.	130
J.6	Monitoring metrics for the two-dimensional square-root boundary trial function.	131
J.7	GPC prediction field for the two-dimensional sinus boundary trial function.	132
J.8	Monitoring metrics for the two-dimensional sinus boundary trial function.	133
K.1	Estimated decision boundary for the three-dimensional nonlinear boundary trial function.	135
K.2	Monitoring metrics for the three-dimensional nonlinear boundary trial function.	136
K.3	Estimated decision boundary for the three-dimensional sinus boundary trial function.	137
K.4	Monitoring metrics for the three-dimensional sinus boundary trial function.	137
L.1	Median slices of the six-dimensional considered trial function	140
L.2	Random slices of the six-dimensional considered trial function	141
L.3	Monitoring metrics for the six-dimensional considered trial function	142
M.1	Median slices of the twelve-dimensional considered trial function	145
M.2	Random slices of the twelve-dimensional considered trial function	146
M.3	Monitoring metrics for the twelve-dimensional considered trial function	147
O.1	Monitoring metrics 2D (t_h, B) structural design and fire variables	153
O.2	Monitoring metrics 3D (t_h, B, F) structural design and fire variables	154
O.3	Monitoring metrics 6D ($t_h, B, F, l, T_{max}, r_c$) structural design and fire variables	154
O.4	Monitoring metrics 2D (t_h, ρ) material variables	155
O.5	Monitoring metrics 3D (t_h, ρ, E) material variables	155
O.6	Monitoring metrics 6D ($t_h, \rho, E, B, f_c, e_0$) material variables	156

List of Tables

2.1	NR-(S)MOCU-RO algorithm: n -th iteration, by Zhao et al. [29]	26
2.2	NR-SMOCU-SGD algorithm: n -th iteration, by Zhao et al. [28]	26
4.1	Temperature-dependent thermal conductivity from Hopkin et al. [44].	43
4.2	Values and Bounds of input variables for simulation data generation	49
4.3	Probabilistic descriptions of GL24h material properties adopted from Schilling et al. [40] and Blaß [47].	53
5.1	Numerical simulation results for Test 1 of Gernay et al. [31].	56
5.2	Numerical simulation results for Tests 2 & 5 of Gernay et al. [31].	56
5.3	Numerical simulation results for Tests 3 & 6 of Gernay et al. [31].	57
5.4	Numerical simulation results for Tests 4 & 7 of Gernay et al. [31].	58
5.5	Numerical mechanical simulation results for Test 10 of Renard et al. [32].	59
5.6	Numerical mechanical simulation results for Test 10 of Renard et al. [32].	60
5.7	Numerical mechanical simulation results for Test 11 of Renard et al. [32].	61
5.8	Numerical mechanical simulation results for Test 15 of Renard et al. [32].	62
5.9	Numerical mechanical simulation results for Test 17 of Renard et al. [32].	63
5.10	Numerical mechanical simulation results for Test 9 of Renard et al. [32].	64
5.11	Numerical mechanical simulation results for Test 12 of Renard et al. [32].	65
B.1	Input parameters used for the symmetry modelling.	97
B.2	Comparison of total runtimes for full and symmetry-based models.	98
C.1	Input parameters used for the thermal mesh sensitivity analyses.	99
C.2	Thermal mesh sensitivity results	100
D.1	Input parameters used for the restart thermal simulations	101
D.2	Fire curve variables corresponding to tests 3 and 6 of Gernay et al. [31]	102
E.1	Input parameters used for the comparison between the simplified analytical model and the non-linear SAFIR thermal analysis (Tests 1–4).	111
E.2	Input parameters used for the comparison between the simplified analytical model and the non-linear SAFIR thermal analysis (Tests 5–8).	111
E.3	Comparison of SAFIR and analytical model results (Tests 1–4).	111
E.4	Comparison of SAFIR and analytical model results (Tests 5–8).	111
G.1	Input parameters used for the moment initialisation analyses.	118
G.2	Comparison of moment initialisation approaches	119
H.1	Input parameters used for the mechanical mesh sensitivity analyses.	120
H.2	Mechanical mesh sensitivity results.	121
I.1	Test-specific input parameters for test H27A of Stanke et al. [35].	123
I.2	Simulation results for the H27A test of Stanke et al. [35].	123
I.3	Test-specific input parameters for test R20A of Stanke et al. [35].	123
I.4	Simulation results for the R20A test of Stanke et al. [35].	123
I.5	Test-specific input parameters for test R15A of Stanke et al. [35].	124
I.6	Simulation results for the R15A test of Stanke et al. [35].	124
I.7	Test-specific input parameters for test H40 of Stanke et al. [35].	124
I.8	Simulation results for the H40 test of Stanke et al. [35].	124

N.1	Corresponding fire resistance time R for the limiting scenarios defined by the structural variable bounds.	150
N.2	Probabilistic descriptions of the applied material variables.	152

Nomenclature

Abbreviations

Abbreviation	Definition
AL	Active Learning
ARD	Automatic Relevance Determination
BALD	Bayesian Active Learning by Disagreement
CLT	Cross-Laminated Timber
DHP	Duration of Heating Phase
EER	Estimated Error Reduction
EP	Expectation Propagation
FEM	Finite Element Method
FRR	Fire Resistance Rating
GLT	Glued Laminated Timber
GP	Gaussian Process
GPC	Gaussian Process Classification
MES	Maximum Entropy Sampling
MOCU	Mean Objective Cost of Uncertainty
OBC	Optimal Bayesian Classifier
PDF	Probability Density Function
PDE	Partial Differential Equation
PINN	Physics-Informed Neural Network
RBF	Radial Basis Function
RO	Random Optimisation
SGD	Stochastic Gradient Descent
SMOCU	Smooth Mean Objective Cost of Uncertainty
<i>SAFIR Modules</i>	
WOODEC5	SAFIR thermal module using Eurocode 5 effective thermal properties
HOPKINS	SAFIR thermal module based on the conductivity model by Hopkin et al.
HOPKINS+FLUX	SAFIR thermal module using HOPKINS module extended with an additional moving boundary heat-flux term
WOODEC5	SAFIR mechanical module using Eurocode 5 strength and stiffness reduction factors
WOODPRBWE	SAFIR mechanical module with Weibull-based strength reduction factors

1

Introduction

This chapter introduces the thesis by outlining the background and motivation of the research and defining the problem under investigation. It identifies the existing research gaps and presents the adopted research method. Subsequently, the research objectives and questions are formulated, followed by a clarification of the scope and limitations of the study. The chapter concludes with an overview of the thesis structure.

1.1. Background and Motivation

In recent years, there has been a growing global interest in the use of mass timber construction. Engineered timber products such as cross-laminated timber (CLT) panels and glued laminated (glulam) timber elements are increasingly applied in building projects. This development is driven by several factors, including the environmental benefits of timber, its architectural and aesthetic qualities, and its potential for cost- and time-efficient construction [3, 4]. Moreover, mass timber is recognised as an effective means of storing large amounts of carbon within the built environment [5], which has stimulated the design and construction of larger and taller timber buildings worldwide [6]. An example of such a development is shown in Figure 1.1.





Figure 1.1: Ascent MKE building in Milwaukee, Wisconsin
(Image source: Thornton Tomasetti, 2024)

At the same time, fire safety remains a critical consideration in structural design. The use of timber as a primary load-bearing material introduces specific challenges for fire engineering, with its combustible nature being one important factor among others. Traditionally, the fire performance of structural elements is assessed using standardised fire curves, such as the ISO 834 curve, which represents a continuously growing heating phase of a fire. Within the Eurocode framework, this performance is expressed by the fire resistance rating (FRR), defined as the time, in minutes, that an element is able to maintain its load-bearing capacity under ISO 834 heating conditions [1].

However, real fires typically consist of both a heating phase and a subsequent cooling phase, during which temperatures decrease after reaching a peak. While the limitations of the current fire resistance assessment are recognised, relatively little attention is given to the fact that structural members, irrespective of material, may fail after the peak fire temperature has been reached. Contributing mechanisms include delayed heat transfer within the cross-section, ongoing material degradation, and stress redistributions caused by incompatible deformations or interactions between structural components. Consequently, fire engineering research has recently started to shift more focus on structural behaviour during the cooling phase of fires exposures, in which simulation and experimental investigations have demonstrated that structural members, including timber elements, may experience failure during or even after the decay phase of a fire [2].

This severe scenario, commonly referred to as *delayed failure*, poses significant risks to building occupants, firefighters, and first responders. In addition, it compromises property protection and reduces the resilience of the built environment. Ensuring structural integrity throughout the full fire duration is therefore of central importance. The issue is particularly relevant for tall buildings, where prolonged evacuation procedures and extended firefighting operations require that structural integrity is maintained not only throughout the full fire duration, but also after complete burnout.

Practical fire events further highlight the importance of maintaining structural integrity throughout the full fire duration and the potential severity of structural failure. In 2014, a timber-based laboratory building at the University of Nottingham was reported as destroyed by fire, demonstrating that complete loss of a timber structure may occur in practice [7, 8]. In tall buildings, late-stage fire-induced failures also have severe outcomes. The Plasco Building in Tehran collapsed after several hours of fire exposure in 2017, and the collapse of the 7 World Trade Center after a prolonged fire further underlines the seriousness of this risk in high-rise structures [9, 10]. Together, these cases show that structural integrity cannot be evaluated on the heating phase alone, and that late-stage failure mechanisms may govern structural performance.

In response to the concern of delayed failure, the assessment of structural performance under fire has gradually shifted towards considering the complete fire duration, encompassing both heating and cooling phases. Within this context, a new performance metric, referred to as *burnout resistance*, has been introduced. Burnout resistance describes the ability of a structural member to maintain its intended load-bearing function throughout the entire fire exposure and post-fire phase. The concept was initiated by Gernay and Franssen [11] and is quantified using the Duration of Heating Phase (DHP), which represents the minimum heating duration that leads to failure of a structural element. In this way, burnout resistance provides a more realistic representation of fire performance by directly relating structural survival to the complete fire duration.

1.2. Problem Description

Despite the increasing application of mass timber in modern buildings, the fire performance of timber structures remains a subject of concern. A key issue is the occurrence of delayed failure during the cooling phase of natural fires, which has been shown to compromise structural integrity after peak temperatures have been reached. Such failures challenge conventional assumptions regarding fire resistance and raise questions about the adequacy of current design approaches for ensuring safety and resilience.

Within the current Eurocode framework, fire resistance is defined with respect to the ISO 834 standard fire and is limited to the heating phase of a fire. Hence, the fire resistance rating does not explicitly account for the cooling phase or for the possibility of failure occurring after the peak fire temperature. For timber structures, this limitation is particularly pronounced. Timber exhibits a relatively low thermal conductivity, which delays heat penetration into the cross-section, while its strength and stiffness degrade rapidly at elevated temperatures. These characteristics make timber especially susceptible to delayed heating effects and, consequently, to delayed failure [12]. Furthermore, as a combustible material, timber may contribute to compartment fire dynamics, potentially increasing both the severity and duration of a fire [13–15]. If auto-extinction does not occur, complete burnout of the structure may follow and structural survival cannot be ensured.

At the level of the structural system, these material- and member-level vulnerabilities translate into

broader challenges for mass timber buildings and several studies have identified resistance to full burnout as a critical challenge for mass timber structures [16, 17].

Taken together, these material- and system-level vulnerabilities indicate that achieving burnout resistance in timber structures remains a significant challenge. While the concept of burnout resistance offers a framework for assessing performance over the full fire duration, its implications for timber structures are not yet fully reflected in current design practice. Existing design approaches primarily focus on the charring behaviour of timber when determining its residual load-bearing capacity. In addition, more advanced calculation methods rely on thermal and mechanical material properties that are calibrated under standard fire (e.g. ISO 834) heating conditions. As such, the applicability of these assessment tools to fire scenarios that include a cooling phase is uncertain. More broadly, the conditions under which delayed failure governs structural performance, and the extent to which timber structures can maintain load-bearing capacity during and after a fire, remain insufficiently characterised.

Concerns regarding delayed failure have already begun to influence design recommendations and regulatory discussions. For example, the draft version of the Dutch NTA 6125 proposes extensions to fire resistance requirements for timber buildings in order to address the risk of delayed failure [18]. However, the current Eurocode provisions do not explicitly account for this phenomenon. This inconsistency highlights the growing tension between emerging design recommendations and the existing regulatory framework, and raises questions regarding the extent to which such additional measures are justified.

The investigation of delayed failure is further complicated by the practical limitations of available assessment methods. Experimental testing at full scale is time-consuming and particularly costly, making it impractical to investigate a wide range of influencing factors. The understanding of such phenomena in the literature relies therefore predominantly on physical modelling as a means of describing real-world structural behaviour in a controlled and systematic manner.

Because the coupled thermal and mechanical response of structural members under fire is highly non-linear, especially when the cooling phase is included, it cannot be determined in advance whether simplified calculation models are capable of adequately capturing the structural response. Advanced numerical techniques, such as finite element modelling, are commonly adopted as the most accurate available representation of the underlying physical processes. However, higher modelling fidelity comes at the expense of substantially increased computational cost, which naturally restricts the feasibility of efficiently identifying the influences of individual variables and systematic assessments of different scenarios.

1.3. Research Gaps

The growing recognition of delayed failure and burnout resistance has highlighted important limitations in the current assessment of timber structures under fire. While multiple studies have demonstrated the relevance of structural behaviour during the cooling phase, several key gaps remain in the ability to reliably analyse and predict this behaviour within an engineering design context. A more detailed examination of the identification of these research gaps is presented in the literature review in Chapter 3.

The first gap concerns the applicability of existing numerical modelling approaches for capturing delayed failure in timber structures. Current design codes and most established modelling procedures have been developed and calibrated primarily for standardised fire scenarios that consider only the heating phase. Hence, it remains unclear whether commonly used numerical models and material formulations can accurately capture the coupled thermal-mechanical behaviour of timber when the cooling phase is included.

The second gap relates to the practical applicability of numerical modelling for the assessment of burnout resistance across a wide range of scenarios. While detailed thermo-mechanical simulations have the potential of providing high-fidelity insight into structural behaviour, their computational cost makes systematic assessments impractical. This limits the ability to draw general conclusions regarding delayed failure behaviour based on numerical modelling alone.

As a result, there is a clear need for computationally efficient assessment techniques that are able to support systematic exploration of structural response under fire. This increased efficiency may be achieved through computationally less demanding approaches, such as

reduced-order modelling or data-driven techniques. In this context, one promising approach is surrogate modelling, which has been applied within the traditional structural fire resistance framework to approximate computationally expensive numerical models, thereby enabling efficient systematic assessments. However, the application of surrogate models to the assessment of delayed failure and burnout resistance in timber structures remains unexplored. In particular, it is not yet established how surrogate models can be formulated to reliably capture the transition between structural survival and failure.

Together, these gaps indicate a lack of validated and computationally efficient methodologies for assessing burnout resistance in timber structures under realistic fire conditions. Addressing these gaps is essential to support the development of reliable assessment tools and to support discussions on whether current design provisions for mass timber buildings are adequate.

1.4. Research Method

To address the identified research gaps, this thesis adopts a modelling-based research method in which numerical simulation and surrogate modelling techniques are integrated. The research is conducted within a computational framework, focusing on the assessment of structural performance and failure under fire exposure.

The physical response of the structure is investigated using non-linear finite element modelling. This approach enables the representation of complex material behaviour, geometric nonlinearity, and progressive failure mechanisms that govern structural performance under elevated temperatures. In this way the numerical modelling provides a framework in describing the delayed failure phenomena. For this purpose, the specialised finite element software SAFIR is employed, which is specifically developed for fire-related structural analyses and is widely applied in fire engineering design.

In addition to the physics-based simulations, a surrogate modelling framework is introduced to evaluate the structural burnout resistance. Gaussian Process Classification (GPC) is applied to predict whether a given set of governing parameters leads to failure or no failure. By learning from the finite element simulation data, the classification model provides an efficient manner for burnout resistance prediction across a wide parameter space. The parameter space is physically motivated by a defined set of input variables describing structural, fire, and material characteristics, selected to represent the dominant physical mechanisms and their interactions governing the structural response under fire.

By integrating non-linear finite element modelling with Gaussian Process Classification, the adopted research method combines detailed physical simulation with data-driven inference. This integrated approach allows the underlying physical processes to be captured while also enabling an efficient prediction and assessment framework. More detailed formulations and justifications of the applied methods are presented in Chapter 4.

1.5. Research Objectives and Questions

This research is structured around two main objectives, each addressing a key aspect of the assessment of delayed failure and burnout resistance of glulam timber columns subjected to realistic fire exposures. Together, these objectives support the broader aim of improving the understanding of delayed failure behaviour and burnout resistance under fire scenarios including both the heating and cooling phases.

1. The first objective is to investigate the applicability of non-linear thermo-mechanical finite element modelling in SAFIR for describing the delayed failure behaviour of glulam timber columns under natural fire exposure.
2. The second objective is to investigate the applicability of Gaussian Process Classification as a surrogate modelling framework for assessing burnout resistance.

Based on these objectives, the following main research question is formulated:

How are the delayed failure behaviour and burnout resistance of glulam timber columns under natural fire exposure, including both heating and cooling phases, modelled and predicted in a reliable and computationally efficient manner?

To address this main question, the following sub-questions are formulated to guide the analysis and structure of the thesis.

- Which physical processes and governing variables influence the thermo-mechanical response of glulam timber columns subjected to realistic fire exposures, including both the heating and cooling phases?
- To what extent is non-linear thermo-mechanical finite element modelling in SAFIR capable of reproducing experimentally observed delayed failure behaviour of glulam timber columns under natural fire exposure?
- To what extent is Gaussian Process Classification a reliable and useful surrogate modelling approach for approximating burnout resistance outcomes predicted by numerical simulations of glulam timber columns?
- How can Gaussian Process Classification be applied to support efficient parametric and probabilistic assessment of burnout resistance for glulam timber columns?

1.6. Scope and Limitations

This study focuses on the fire-induced behaviour of single, isolated glulam timber columns and does not consider complete structural frames, multi-element assemblies, or system-level interactions. Glued laminated timber (glulam) made of spruce is considered, with GL24h adopted as the reference strength class due to its widespread structural application, the availability of validated material data, and its use in the experimental studies employed for model validation. The columns are taken to have a square cross-section and to be fully exposed to fire on all four sides. Structural boundary conditions are initially modelled as pinned–roller supports, unless stated otherwise. This implies both lateral displacements restrained at the bottom support and only the horizontal displacement restrained at the top support. A vertical applied force is considered at the top of the column in which an initial eccentricity is applied in accordance with the experimental configurations used for model validation.

Fire exposure is modelled using prescribed fire curves, derived either from experimental data or from schematic fire modelling approaches. The structural response is assumed to depend directly on the applied gas temperature history, in which the gas temperature is taken as uniform over the height of the compartment. Schematic fire scenarios are defined using the relevant characteristics of Eurocode parametric fire curves, with the ISO 834 standard fire curve considered for comparison where appropriate. Experimental fire curves are defined based on the mean compartment gas temperature. Both the heating and cooling phases of the fire exposure are included in the analysis.

The assessment is conducted exclusively through modelling methods, and no additional experimental testing is performed. One single non-linear thermo-mechanical finite element software package, SAFIR, is used throughout the study to ensure methodological consistency.

Several simplifying assumptions are adopted, which define limitations of the study. Fire dynamics within the compartment, including ventilation effects and spatial temperature variations, are not explicitly modelled. The phenomenon of auto-extinction is not considered. No active or passive fire protection measures are included, and the analysis is limited to unprotected timber members. Furthermore, the glulam members are assumed to be manufactured using heat-resistant adhesives, in which the possibility of delamination is not considered.

Experimental data used in this study are limited to the research programme described in Section 3.1, conducted by the CERIB research institute in France. For the natural fire tests, only thermal measurement data are included, as no mechanical deflection data are available.

1.7. Thesis Outline

In order to provide a clear structure, the thesis is organised in multiple chapters.

Chapter 2 sets out the theory underlying the study. It describes the fire behaviour of timber, the governing heat transfer equations, and the thermal and mechanical properties of the material. In addition,

it presents the characteristics of the fire curves used in this research and outlines the core principles of the SAFIR finite element software. Finally, the chapter introduces the fundamentals of Gaussian Process Classification as the surrogate modelling technique.

Chapter 3 presents a review of the relevant literature. It introduces the delayed failure phenomena of glulam timber columns as investigated in existing studies. Subsequently, the concepts of burnout resistance and the Duration of Heating Phase (DHP) are presented. Finally, current delayed failure modelling approaches for glulam timber columns are discussed.

Chapter 4 sets out the methodology. It explains the general research approach and its justification, and provides a detailed setup of the numerical and surrogate models.

Chapter 5 presents the results of the analyses, including outcomes from the numerical simulations and the surrogate modelling based on Gaussian process Classification.

Chapter 6 discusses the findings. The results are interpreted and the limitations of the study are examined. In addition a discussion is provided in which the implications for fire engineering practice are considered.

Finally, Chapter 7 concludes the thesis by summarising the main findings and providing recommendations for further research.

2

Theory

This chapter presents the theory underlying the assessment of the fire behaviour and burnout resistance of glue-laminated structural timber columns. The chapter first describes the fire behaviour of structural timber, addressing the governing physical processes, fundamental heat transfer mechanisms, and the temperature-dependent thermal and mechanical material properties. Subsequently, fire exposure models and standard fire curves adopted in design codes are discussed.

The chapter then outlines the theoretical principles of the finite element method as implemented in SAFIR, including both the thermal and mechanical analyses. Finally, Gaussian Process Classification is introduced, with emphasis on its fundamental characteristics and active learning strategies.

2.1. Fire Behaviour of Structural Timber

Structural timber exposed to fire undergoes a range of interacting thermal, chemical, and mechanical processes that govern its temperature evolution and load-bearing capacity. These processes ultimately determine the fire resistance and burnout behaviour of timber members. This section introduces the fundamental mechanisms underlying the fire behaviour of structural timber, with emphasis on the degradation processes induced by elevated temperatures, the governing heat transfer mechanisms, and the temperature-dependent thermal and mechanical material properties. Together, these aspects provide the physical basis for the broader understanding of the fire behaviour of structural timber.

2.1.1. Fire-Induced Degradation Processes

Structural timber exposed to fire undergoes a sequence of temperature-driven degradation mechanisms that progressively alter its material composition and structural integrity, resulting in a change of material state. These fire-induced processes are governed by coupled thermal and chemical interactions and evolve as a function of temperature, time, and oxygen availability. Understanding these mechanisms is essential for interpreting the development of damage in timber members and for understanding their residual load-bearing capacity under fire exposure.

The qualitative description of the fire-induced degradation processes presented in this subsection is primarily based on the comprehensive review by Bartlett et al. [19].

An overview of the coupled in-depth processes and the resulting internal zones that develop during fire exposure is illustrated in Figure 2.1.

Moisture Transport and Evaporation

Upon exposure to fire, the temperature of a timber member initially rises as heat is transferred from the exposed surface into the material. At this early stage, before the onset of chemical decomposition, the thermal response is strongly influenced by the presence of moisture within the timber. As temperatures approach approximately 100 °C, free water contained in the timber begins to evaporate, marking the first significant fire-induced degradation process. The energy supplied during this phase is largely consumed by the phase change of water, thereby delaying the rate of temperature increase within the timber member.

The evaporation of moisture is accompanied by the transport of water vapour within the timber cross-section. Part of the generated vapour migrates away from the heated surface towards cooler regions,

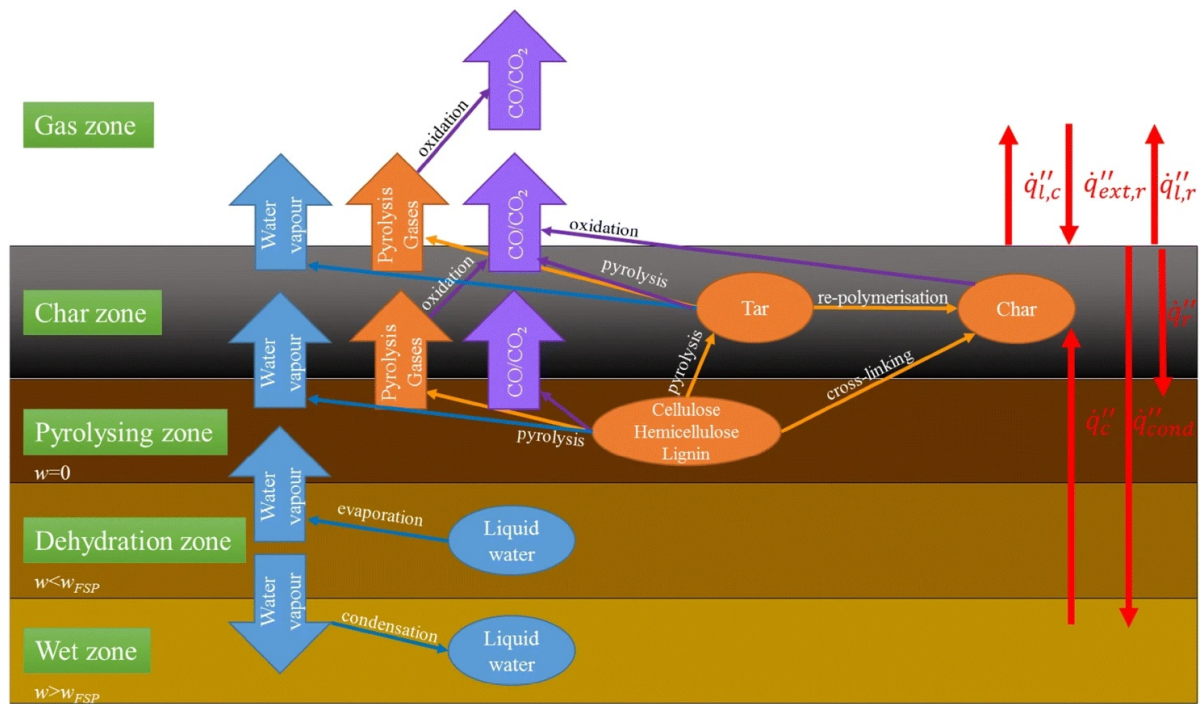


Figure 2.1: Schematic representation of in-depth fire-induced degradation processes and internal zones in timber during fire exposure from Bartlett et al. [19].

where it may re-condense, locally increasing the moisture content. As a result, distinct internal zones develop within the timber member, consisting of a dry zone near the exposed surface, a dehydrating zone, and a wetter inner zone. This moisture redistribution modifies the thermal gradients within the timber cross-section and delays the onset of subsequent degradation processes in the interior regions of the member.

At higher heat fluxes, dehydration may occur simultaneously with the onset of thermal decomposition, whereas at lower heat fluxes these processes tend to remain more clearly separated. In both cases, moisture evaporation acts as an effective heat sink, slowing the temperature rise until evaporation is largely complete, typically up to temperatures of approximately 115 °C. Bound water is released at higher temperatures, contributing further to the delayed thermal response during the early stages of fire exposure.

Pyrolysis

After most of the free moisture has been removed, continued heating leads to the onset of pyrolysis. Pyrolysis refers to the thermal decomposition of a substance. It is an endothermic process that can occur in the absence of oxygen and should be distinguished from combustion, which is an exothermic oxidation process. Pyrolysis therefore precedes flaming combustion and takes place within the solid material. During this process, the main timber polymers, cellulose, hemicellulose, and lignin, decompose into solid char, liquid tars, and gaseous products. The composition of these products depends on temperature, heating rate, and material properties.

In timber members, pyrolysis occurs within a narrow zone beneath the exposed surface. This zone progresses inward as heating continues and develops in stages. Slow decomposition occurs below approximately 200 °C, more active pyrolysis develops up to around 300 °C, and rapid decomposition occurs at higher temperatures.

At lower temperatures, mass loss is limited and the released products are mainly non-combustible. At higher temperatures, pyrolysis produces larger quantities of flammable gases, which can ignite after leaving the timber and mixing with oxygen. Where moisture evaporation is still ongoing, the temperature rise and pyrolysis rate are reduced. In dried surface layers, pyrolysis proceeds more rapidly and forms a distinct pyrolysing zone that governs subsequent char development.

Oxidation and Combustion

The gas products released during pyrolysis provide the fuel for oxidation and combustion once sufficient oxygen is available. Combustion occurs in the gas phase outside the timber surface and involves the exothermic oxidation of these released gases. The heat released during combustion increases the heat flux to the timber surface and accelerates further pyrolysis in the underlying material.

Flaming combustion is sustained when the rate of gas production is sufficient to maintain ignition. During this stage, most of the available oxygen is consumed by the flame and combustion remains largely localised to the region outside the solid timber. The combustion rate is therefore primarily controlled by the rate of pyrolysis.

Under conditions of limited oxygen supply, combustion transitions to smouldering. Smouldering is a slow form of oxidation that occurs at lower temperatures and without visible flames. It can persist for extended periods and contributes to continued material degradation and heat release.

Charring

Charring is the process in which timber is converted into a carbon-rich solid during fire exposure. This process represents a permanent change in material state and results from sustained pyrolysis and oxidation at elevated temperatures. The char layer forms at the exposed surface and progresses inward as heating continues. Char has a much lower density and thermal conductivity than virgin wood and therefore acts as a protective layer.

Although the char layer provides thermal insulation, it has limited mechanical strength. Char is brittle and may crack or detach from the surface under sustained heating or mechanical disturbance. When char falls off, new uncharred timber is exposed to fire. This exposure increases the local heat flux and accelerates further pyrolysis and combustion.

The formation of char leads to a reduction of the effective load-bearing cross-section of the timber member. As the char depth increases, the remaining section available to resist mechanical loads decreases.

2.1.2. Heat Transfer Mechanisms

Fire-induced degradation of structural timber involves the simultaneous action of several heat transfer mechanisms. Heat is transferred within the material primarily by conduction, while moisture migration within the porous structure can contribute to additional energy transport through phase change effects. At the exposed surface, thermal loading is governed by radiation and convection from the surrounding fire environment. These coupled processes can be described using general energy conservation equations, which form the basis for physically based thermal models of fire-exposed timber members.

Transient Heat Conduction

The thermal response of structural timber exposed to fire can be described using a transient energy conservation framework. This formulation balances the rate of change of stored thermal energy with conductive heat transfer and internal energy exchange. It allows material properties to vary with temperature and material state, which is essential for representing the evolving behaviour of fire-exposed timber.

Heat transfer within the timber is governed primarily by conduction and is commonly described by Fourier's law. The thermal conductivity, density, and specific heat capacity may change as degradation progresses, reflecting moisture loss, chemical decomposition, and char formation. The associated internal energy effects can be represented through a volumetric source or sink term.

The most general statement of energy conservation for a solid domain can be written as

$$\frac{\partial}{\partial t}(\rho c T) = \nabla \cdot (k \nabla T) + Q, \quad (2.1)$$

where ρ is the material density (kg m^{-3}), c the specific heat capacity ($\text{J kg}^{-1} \text{K}^{-1}$), k the thermal conductivity ($\text{W m}^{-1} \text{K}^{-1}$), and Q a general volumetric heat source or sink (W m^{-3}). The left-hand side represents the energy storage term, meaning the rate of change of thermal energy stored in the material per unit volume. The term Q provides a compact representation of internal physical or chemical processes, such as moisture evaporation or material decomposition.

In practical heat transfer modelling, the temperature field is treated as the time-dependent state variable, while material properties such as ρ , c , and k are prescribed as functions of temperature and may vary

spatially across the domain. In this way, any time variation of the material properties is introduced through the evolving temperature field rather than through an explicit dependence on time. The energy storage term is therefore commonly written in terms of the temperature rate, leading to the transient heat conduction equation

$$\rho(\mathbf{x}, T) c(\mathbf{x}, T) \frac{\partial T(\mathbf{x}, t)}{\partial t} = \nabla \cdot (k(\mathbf{x}, T) \nabla T(\mathbf{x}, t)) + Q(\mathbf{x}, T). \quad (2.2)$$

More advanced models could additionally consider moisture movement and other energy effects. Such formulations provide a more complete description of the coupled heat and mass transfer processes but are often simplified in structural fire analyses, where heat conduction with effective thermal properties is preferred.

Radiation

At the exposed surface of structural timber, thermal loading is partly imposed through radiative heat transfer. Radiation refers to the transfer of thermal energy by electromagnetic waves and does not require a material medium. In fire exposure, radiative heat transfer originates from flames, hot gases, and surrounding heated surfaces and acts directly at the boundary of the timber member.

The net radiative heat flux q_{rad} (W m^{-2}) to the surface is expressed using the Stefan–Boltzmann law as

$$q_{\text{rad}} = \varepsilon \sigma (T_{\text{g}}^4 - T_{\text{s}}^4), \quad (2.3)$$

where ε is the surface emissivity (–), σ the Stefan–Boltzmann constant ($5.67 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$), T_{g} the compartment gas temperature (K), and T_{s} is the surface temperature of the timber (K). This formulation accounts for both incoming thermal radiation from the fire and outgoing radiation emitted by the heated surface.

In thermal analyses, radiative heat transfer is introduced as a boundary condition at the exposed surface of the timber member. The resulting radiative heat flux contributes to the surface energy balance and serves as an input for the subsequent conduction of heat into the material.

Convection

At the exposed surface of structural timber, thermal loading also occurs through convective heat transfer. Convection refers to the transfer of thermal energy between a solid surface and an adjacent fluid due to temperature differences and fluid motion. In fire exposure, convective heat transfer arises from the movement of hot gases in contact with the timber surface.

The convective heat flux q_{conv} (W m^{-2}) at the surface is described as

$$q_{\text{conv}} = h_{\text{c}}(T_{\text{g}} - T_{\text{s}}), \quad (2.4)$$

where h_{c} is the convective heat transfer coefficient ($\text{W m}^{-2} \text{ K}^{-1}$), T_{g} is the compartment gas temperature (K), and T_{s} is the surface temperature of the timber (K). The convective heat transfer coefficient depends on the properties and flow conditions of the surrounding gases and may vary with the fire scenario.

In thermal analyses, convective heat transfer is applied as a boundary condition at the exposed surface of the timber member. Together with radiative heat transfer, convection defines the thermal boundary loading imposed by the fire environment and contributes to the surface energy balance that governs heat conduction into the timber cross-section.

2.1.3. Thermal and Mechanical Properties

The fire-induced degradation processes and heat transfer mechanisms described in the preceding sections are reflected in the temperature-dependent thermal and mechanical properties of structural timber. These properties provide the link between the evolving thermal state of the material and its mechanical response under fire exposure. For modelling purposes these properties define heat propagation through the timber cross-section and degradation in stiffness and strength.

Thermal Properties

The thermal response of structural timber under fire exposure is governed by three primary material properties, which include density, specific heat capacity, and thermal conductivity. Together, these properties control the rate of heat storage and heat transfer within the timber cross-section and therefore directly influence the temperature development.

In structural fire design, temperature-dependent thermal properties are commonly adopted from Eurocode provisions [1]. These relationships are calibrated for transient heat conduction analyses without explicit internal heat source or sink terms (e.g. $Q(x, T) = 0$) and are effectively calibrated based on standard fire exposure scenarios, such as the ISO 834 fire curve (see Subsection 2.2.1). As a result, the prescribed thermal properties implicitly incorporate the combined effects of multiple fire-induced physical processes.

The specific heat capacity exhibits a pronounced increase around 100 °C, reflecting the additional energy required for moisture evaporation. This peak effectively represents a heat sink associated with phase change processes and delays the temperature rise within the timber during the early stages of fire exposure.

The thermal conductivity increases with temperature, affecting the magnitude of conductive heat transfer within the timber. This increase reflects the effective calibration of the conductivity for standard fire exposure conditions and implicitly accounts for processes such as cracking and gas movement.

The density of timber decreases with increasing temperature, representing mass loss due to chemical decomposition and charring. This reduction in density influences both the thermal inertia and the mechanical response of the remaining cross-section.

Figure 2.2 illustrates the temperature dependency of the Eurocode thermal properties used for fire analysis.




Figure 2.2: Temperature-dependent thermal properties of timber according to EN 1995-1-2 [1].

Mechanical Properties

The mechanical response of structural timber under fire exposure is governed by the temperature-dependent reduction of strength and stiffness. As the temperature of the material increases, the available load-bearing capacity and deformation resistance are progressively reduced.

In structural fire design, this temperature dependency is commonly represented using reduction factors for strength and stiffness. These factors are provided in EN 1995-1-2 [1] and are applied to the ambient-temperature material properties to obtain the remaining strength and stiffness as a function of temperature. Separate reduction relationships are specified for tensile and compressive strength, accounting for the different failure mechanisms governing timber behaviour under specified loading conditions.

In addition, the formation of char is commonly represented by the 300 °C isotherm in structural fire engineering. This is reflected in the applied reduction factors, for which strength and stiffness are taken as zero beyond this temperature.

The temperature-dependent reduction factors for strength and stiffness applied in this study are illustrated in Figure 2.3.

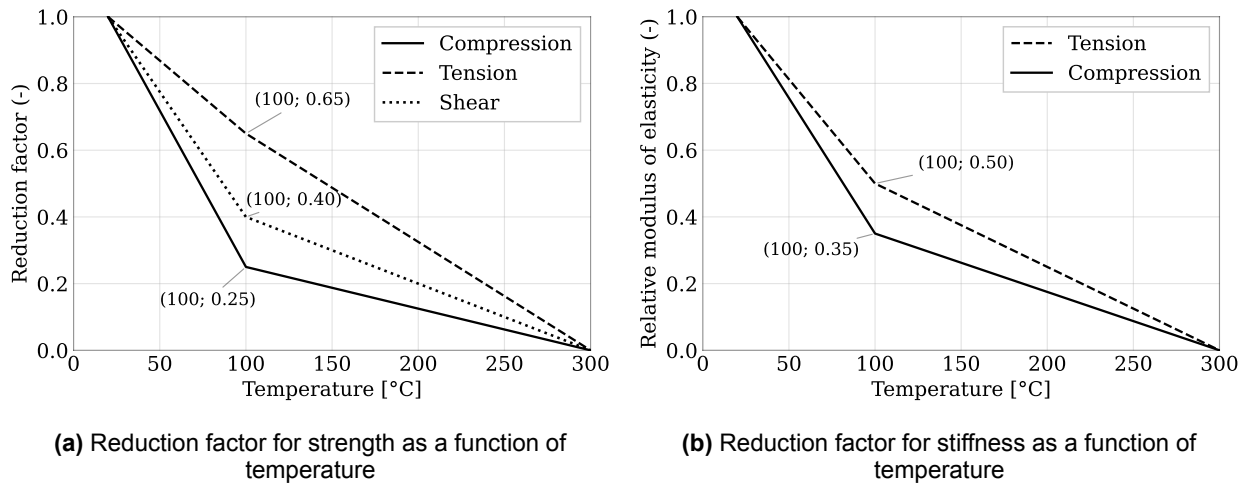


Figure 2.3: Temperature-dependent reduction factors for the mechanical properties of timber according to EN 1995-1-2 [1].

2.2. Fire Exposure and Standard Fire Curves

Fire exposure models describe the temporal evolution of gas temperatures to which structural members are subjected during a fire. These models prescribe the thermal boundary conditions in structural fire analysis and therefore represent a key consideration in the assessment of fire resistance.

To enable consistent assessment and comparison, idealised fire curves are commonly employed in structural fire engineering. This section introduces the fire exposure models considered in this thesis. First, the standard ISO 834 fire curve is presented, which forms the basis for conventional fire resistance design. Subsequently, parametric fire models are discussed, which account for compartment characteristics and fire load. Finally, fire exposure specific to mass timber compartments is addressed, highlighting the influence and complexity of timber contribution to the fire dynamics.

2.2.1. Standard ISO 834 Fire Curve

The ISO 834 standard fire curve is the most widely used fire exposure model in structural fire design and is described in NEN-EN 1991-1-2 [20]. It represents an idealised, fully developed fire and is intended to describe the heating phase of a fire through a prescribed time–temperature relationship. Fire resistance ratings in current design practice are defined with respect to this standardised fire exposure.

The ISO 834 fire curve defines the compartment gas temperature as a function of time according to

$$T_g(t) = T_0 + 345 \log_{10}(8t + 1), \quad (2.5)$$

where T_g is the gas temperature (°C), T_0 is the initial ambient temperature (°C), typically taken as 20 °C, and t is the time (min). This formulation prescribes a monotonically increasing temperature and does not include a cooling phase.

As a result, the ISO 834 fire curve represents a conservative thermal loading scenario, as the gas temperature continues to increase indefinitely with time. Despite this simplification, the curve provides a consistent and reproducible basis for assessing fire resistance and forms the reference exposure for standard fire testing and Eurocode-based design.

An illustration of the ISO 834 fire curve is shown in Figure 2.4.




Figure 2.4: ISO 834 standard fire curve as defined in EN 1991-1-2 [20].

2.2.2. Parametric Fires

In addition to nominal fire curves, Eurocode 1 provides parametric fire models that describe the gas temperature development in enclosed compartments based on simplified fire physics [20]. These models are intended to represent post-flashover fires and account explicitly for the influence of ventilation conditions, thermal properties of the compartment boundaries, and the available fire load. In contrast to the ISO 834 fire curve, parametric fires include both a heating phase and a decay phase.

The heating phase of the Eurocode parametric fire is defined by a temperature–time relationship expressed as a function of a dimensionless time parameter. The compartment gas temperature T_g (°C) during the heating phase is given by

$$T_g(t^*) = 20 + 1325 \left(1 - 0.324 e^{-0.2 t^*} - 0.204 e^{-1.7 t^*} - 0.472 e^{-19 t^*} \right), \quad (2.6)$$

where t^* is the dimensionless time defined as

$$t^* = \Gamma t, \quad (2.7)$$

with t the physical time (h) and Γ a scaling factor that governs the rate of fire development.

The factor Γ incorporates the combined influence of ventilation conditions and the thermal inertia of the compartment boundaries. It is defined as

$$\Gamma = \frac{(O/b)^2}{(0.04/1160)^2}, \quad (2.8)$$

where O is the opening factor and b is the thermal inertia of the enclosure surfaces.

The opening factor represents the degree of ventilation of the compartment, while the thermal inertia characterises the heat-absorbing capacity of the boundary materials.

The duration of the heating phase and the transition to the decay phase are governed by the available design fire load density $q_{f,d}$. Together, the opening factor O , the thermal inertia b , and the fire load density $q_{f,d}$ define the shape, intensity, and duration of the parametric fire curve. Figure 2.5 illustrates a set of parametric fire curves, presenting different shapes and behaviours.

Annex A of EN 1991-1-2 provides detailed expressions, parameter limits, and decay formulations for the calculation of these quantities as a function of compartment geometry, boundary conditions, and fire load characteristics. For a complete overview reference is made to the Eurocode EN 1991-1-2 [20].

It is noted that different combinations of opening factor, thermal inertia, and fire load density are able to produce similar temperature–time histories, as these parameters influence both the time scale and the achievable temperature level in coupled ways. As a result, distinct compartment configurations can lead to comparable thermal fire exposure when represented using the parametric fire model.




Figure 2.5: Illustrative Eurocode parametric fire curves based on Annex A of EN 1991-1-2 [20].

2.2.3. Mass Timber Compartment Fires

Natural compartment fires are characterised by a wide variety of temperature–time histories resulting from combined influence of ventilation conditions, compartment geometry and fire load.

In particular to mass timber buildings, fire development is further influenced by the presence of the load bearing structure. Unlike compartments with a non-combustible nature, mass timber compartments contain additional material contributing to the fire dynamics.

Experimental investigations into mass timber compartment fires provide important insight into these differences. Large-scale tests with exposed cross-laminated timber (CLT) surfaces have shown that timber provides an additional fuel source, resulting in higher compartment temperatures and increased fire severity compared to otherwise identical compartments with non-combustible boundaries [13, 15]. Furthermore, mass timber compartment fires have been observed to exhibit prolonged decay phases and, in some cases, fire regrowth associated with continued pyrolysis and delamination of charred timber layers [13, 14]. This results in insights that the additional contribution of the timber load bearing structure leads to behaviour that is not captured by standard fire curves and only partially represented by parametric fire models.

Representative temperature–time histories from McNamee et al. [13] are presented in Figure 2.6 to visualise the characteristic features of mass timber compartment fires and to highlight their differences from conventional fire exposure models.

2.3. SAFIR Finite Element Method

SAFIR is a finite element software package developed for analysing structures exposed to fire, combining thermal and mechanical modelling within a single framework. It is intended to represent heat transfer and structural response under elevated temperatures, and it supports nonlinear material behaviour, large displacements, and time-dependent effects. The software is able to be applied to ambient-temperature and fire scenarios using both two-dimensional and three-dimensional finite element formulations.

The analysis procedure in SAFIR is based on a weakly coupled sequential approach. A transient thermal analysis is performed first to determine the temperature evolution within structural members. These temperature histories are subsequently used as input for the mechanical analysis, where they influence temperature-dependent material properties, thermal strains and structural equilibrium. This one-way coupling reflects the assumption that the structural response does not significantly affect the heat transfer process.

The formulation and capabilities of the SAFIR finite element method described in this section are based on published SAFIR reference papers and the associated technical documentation [21–25].

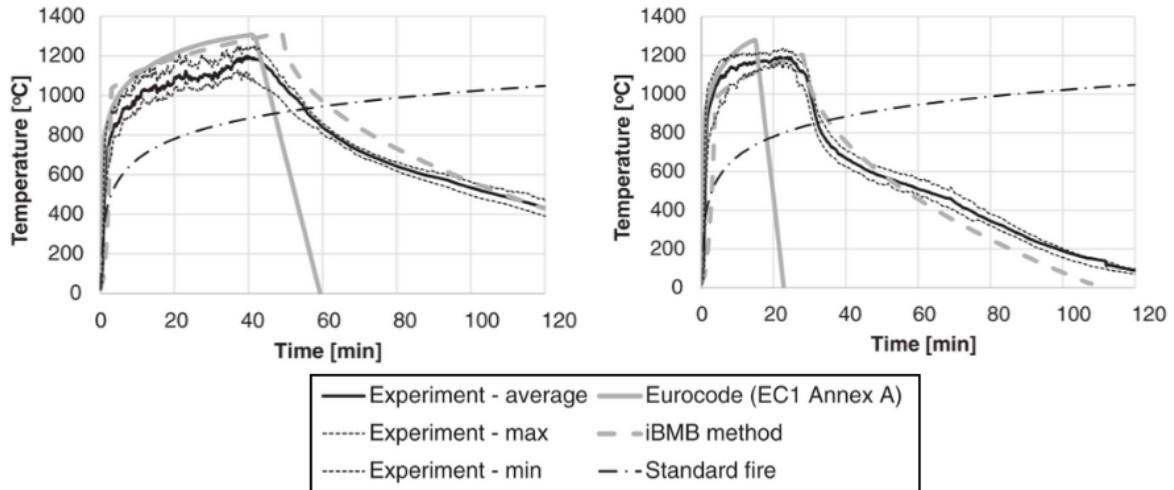


Figure 2.6: Comparison of experimentally measured gas temperature histories in mass timber compartments with standard and parametric fire curves [13].

2.3.1. Thermal Analysis

The thermal analysis in SAFIR is used to determine the transient temperature distribution within structural members and surrounding domains when exposed to fire. Its primary purpose is to compute the evolution of temperature fields as a function of time, which subsequently serve as input for the mechanical analysis. The thermal module is capable of modelling heat transfer in solids, heat exchange at exposed surfaces, and radiative and convective heat transfer within cavities and enclosed spaces. The analysis is able to be performed in two or three dimensions, depending on the adopted geometric idealisation and the type of structural element considered.

The thermal behaviour is governed by the transient heat conduction equation, as described in Subsection 2.1.2 with Equation 2.2. In SAFIR, internal heat generation is not considered, and the transient thermal field is solved within the solid domain using the finite element method. Material properties such as density, thermal conductivity and heat capacity may depend on temperature, allowing the nonlinear thermal response of materials exposed to fire to be represented.

At the boundaries of the solid domain, heat exchange with the surrounding environment is modelled through convection and radiation, resulting in a superposition of the defined Equations 2.4 and 2.3. The resulting normal heat flux q_n (W m^{-2}) at the surface is expressed as

$$q_n = h_c(T_g - T_s) + \varepsilon \sigma(T_g^4 - T_s^4), \quad (2.9)$$

where h_c is the convective heat transfer coefficient ($\text{W m}^{-2} \text{K}^{-1}$), T_g the gas temperature ($^{\circ}\text{C}$), T_s the surface temperature ($^{\circ}\text{C}$), σ the Stefan–Boltzmann constant ($5.67 \times 10^{-8} \text{ W m}^{-2} \text{K}^{-4}$), and ε the surface emissivity (–).

Finite Element Discretisation

Within each finite element, the temperature field is approximated using classical shape functions applicable for triangular and rectangular finite elements. The temperature at any point inside an element is expressed as

$$T(\mathbf{x}, t) = \mathbf{N}(\mathbf{x}) \mathbf{T}(t), \quad (2.10)$$

where \mathbf{N} is the vector of shape functions and \mathbf{T} the vector of nodal temperatures. The same shape functions are used to interpolate both the geometry and the temperature field, ensuring continuity of temperature at element boundaries.

By substituting the finite element approximation of the temperature field into the governing heat conduction equation and applying the Galerkin weighted residual method, the weak form of the thermal equilibrium equation is obtained. Green's theorem is used to rewrite the heat conduction term, allowing

heat fluxes at the element surfaces to be included directly in the finite element formulation. After assembly over all elements, the global system of equations governing heat transfer in the structure is written as

$$\mathbf{K} \mathbf{T} + \mathbf{C} \dot{\mathbf{T}} = \mathbf{g}, \quad (2.11)$$

where \mathbf{K} is the conductivity matrix, \mathbf{C} the heat capacity matrix, \mathbf{T} the vector of nodal temperatures, and \mathbf{g} the vector accounting for boundary heat fluxes. Numerical Gauss integration is used to evaluate the element matrices, and temperature-dependent material properties are accounted for at the integration points.

Time Integration and Iterative Convergence

The transient thermal equilibrium equation is solved incrementally in time using an implicit single time step integration scheme. SAFIR adopts a generalised midpoint formulation in which thermal equilibrium is enforced at an intermediate time level within each time increment. It is assumed that the nodal temperatures vary linearly between two successive time steps. Denoting the converged temperature vector at time t_n by \mathbf{T}_n and the unknown temperature vector at t_{n+1} by \mathbf{T}_{n+1} , an intermediate temperature state \mathbf{T}_θ is defined as

$$\mathbf{T}_\theta = \mathbf{T}_n + \theta (\mathbf{T}_{n+1} - \mathbf{T}_n), \quad (2.12)$$

where $\theta \in [0, 1]$ is a user-defined parameter. The choice of θ governs the numerical characteristics of the scheme, with $\theta = 1$ corresponding to a fully implicit backward Euler method, $\theta = 0.5$ to the trapezoidal rule, and smaller values of θ introducing increasingly explicit behaviour.

Under the assumption of linear temperature variation in time, the temperature rate at the intermediate time level is constant and given by

$$\dot{\mathbf{T}}_\theta = \frac{\mathbf{T}_\theta - \mathbf{T}_n}{\theta \Delta t}, \quad (2.13)$$

where $\Delta t = t_{n+1} - t_n$ is the time increment, also specified by the user.

Substituting this expression into the semi-discrete thermal equilibrium equation 2.11, and writing it in a residual form at time $t_\theta = t_n + \theta \Delta t$ gives,

$$\mathbf{r}(\mathbf{T}_\theta) = \mathbf{K}(\mathbf{T}_\theta) \mathbf{T}_\theta + \mathbf{C}(\mathbf{T}_\theta) \dot{\mathbf{T}}_\theta - \mathbf{g}(\mathbf{T}_\theta), \quad (2.14)$$

which represents the out-of-balance thermal forces at the intermediate time level. The matrices \mathbf{K} and \mathbf{C} and the load vector \mathbf{g} may depend on temperature through temperature-dependent material properties and boundary conditions.

At the beginning of each time increment, an initial estimate of \mathbf{T}_θ is obtained by linear extrapolation from previously converged solutions. As this estimate does not generally satisfy the equilibrium condition $\mathbf{r}(\mathbf{T}_\theta) = \mathbf{0}$, the solution is corrected iteratively using a Newton-type procedure. Linearisation of the residual about the current iterate $\mathbf{T}_\theta^{(i)}$ leads to

$$\mathbf{r}(\mathbf{T}_\theta^{(i)} + \Delta \mathbf{T}_\theta) \approx \mathbf{r}(\mathbf{T}_\theta^{(i)}) + \frac{\partial \mathbf{r}}{\partial \mathbf{T}} \Delta \mathbf{T}_\theta = \mathbf{0}. \quad (2.15)$$

The temperature correction is therefore obtained as

$$\Delta \mathbf{T}_\theta = - \left(\frac{\partial \mathbf{r}}{\partial \mathbf{T}} \right)^{-1} \mathbf{r}(\mathbf{T}_\theta^{(i)}). \quad (2.16)$$

The first derivative of the residual with respect to the temperature vector is obtained from Equation 2.14 and is approximated as

$$\frac{\partial \mathbf{r}}{\partial \mathbf{T}} \approx \mathbf{K} + \frac{\mathbf{C}}{\theta \Delta t} - \frac{\partial \mathbf{g}}{\partial \mathbf{T}}. \quad (2.17)$$

To obtain a symmetric iteration matrix, SAFIR neglects the terms $\partial \mathbf{K} / \partial \mathbf{T}$ and $\partial \mathbf{C} / \partial \mathbf{T}$. In addition, the contribution associated with $\partial \mathbf{g} / \partial \mathbf{T}$ is treated such that the resulting iteration matrix remains symmetric. Convergence is assessed using the residual vector $\mathbf{r}(\mathbf{T}_\theta)$ defined in Equation 2.14. As a result, the introduced approximations influence the number of iterations and the computational efficiency, but do not affect the equilibrium condition that defines the converged solution.

The Newton iterations are repeated at the fixed intermediate time level t_θ until the residual and/or the temperature corrections fall below prescribed tolerances. Once convergence is achieved, the temperature at the end of the time step is recovered algebraically from an inverse calculation of Equation 2.12, and the solution advances to the next time increment.

SAFIR allows the time increment Δt to vary during the analysis, enabling adaptive or dynamic time-stepping strategies in which smaller time steps are adopted during periods of rapid temperature change. Temperature-dependent material properties may be treated as reversible or irreversible, depending on whether they are evaluated based on the current temperature or on the maximum temperature previously reached at the integration points. The overall time integration and convergence procedure is illustrated schematically in Figure 2.7.




Figure 2.7: SAFIR thermal analysis iteration process [22].

2.3.2. Mechanical Analysis

The mechanical analysis in SAFIR is used to determine the structural response of a system subjected to mechanical loading under elevated temperatures. Using the temperature histories obtained from the thermal analysis, the mechanical module evaluates the evolution of displacements, internal forces, stresses, and strains over time. The formulation accounts for geometric nonlinearity and material nonlinearity arising from temperature-dependent material behaviour. As a result, the analysis is able to capture deformations due to applied loading, stiffness degradation, and strength reduction during fire exposure.

Finite Element Discretisation

The mechanical problem is formulated using a displacement-based finite element method. The displacement field within each element is approximated using appropriate shape functions, such that the vector of displacements can be written as

$$\mathbf{u} = \mathbf{N} \mathbf{p}, \quad (2.18)$$

where \mathbf{N} is the matrix of shape functions and \mathbf{p} the vector of nodal displacements. This approximation allows the displacement field to be expressed in terms of a finite number of degrees of freedom associated with the element nodes.

The incremental strain field is derived from the incremental nodal displacements through the strain–displacement relationship,

$$d\boldsymbol{\varepsilon} = \mathbf{B} d\mathbf{p}, \quad (2.19)$$

where \mathbf{B} is the strain–displacement matrix. In SAFIR, this matrix not only contains the spatial derivatives of the shape functions, as in small-deformation formulations, but also accounts for nodal displacements in the co-rotational configuration, enabling the treatment of large displacements and rotations.

Strain Decomposition and Constitutive Behaviour

A fundamental assumption in the mechanical formulation of SAFIR is the concept of strain decomposition. In structures exposed to fire, stresses are assumed to be generated exclusively by the stress-related strain, which is defined as the difference between the total strain and the initial and thermal strain components. This relationship is expressed as

$$\varepsilon_{\sigma} = \varepsilon_{\text{total}} - \varepsilon_i - \varepsilon_{\text{th}}, \quad (2.20)$$

where $\varepsilon_{\text{total}}$ is obtained from the nodal displacements, ε_i represents initial strains, and ε_{th} denotes thermal strains induced by temperature changes.

The constitutive law is evaluated at the integration points using the current temperature, such that material strength and the tangent modulus evolve during the analysis. Plastic strains are updated only after convergence has been achieved within a time step, ensuring consistency between the stress state and the equilibrium configuration.

Incremental Equilibrium

The transient mechanical behaviour of the structure in SAFIR is computed incrementally in time using a step-by-step procedure. At each time step, the temperature field corresponding to the new time level is read from the thermal analysis and introduced into the mechanical calculation.

At the beginning of a time step, the nodal displacements are kept constant and the total strain therefore remains unchanged. The temperature increment induces an increment of thermal strain at the integration points. Through the strain decomposition hypothesis, the corresponding stress-related strain is determined and used to evaluate the temperature-dependent constitutive law. As a result, each integration point is governed by its own material stress–strain relationship reflecting the local temperature history. The resulting stresses are computed at the integration points and subsequently integrated to obtain the internal nodal forces and the element stiffness matrices.

Due to the change in stresses and material properties, the internal nodal forces are generally no longer in equilibrium with the applied external forces. Mechanical equilibrium is therefore enforced iteratively in incremental form. The linearised equilibrium equation governing the displacement correction is written as

$$(\mathbf{K}_u + \mathbf{K}_s) d\mathbf{p} = \mathbf{f}^{\text{ext}} - \mathbf{f}^{\text{int}}, \quad (2.21)$$

where $d\mathbf{p}$ is the incremental nodal displacement vector, \mathbf{f}^{ext} the vector of external nodal forces, and \mathbf{f}^{int} the vector of internal nodal forces. \mathbf{K}_u represents the linear elastic and geometric stiffness matrices, while \mathbf{K}_s denotes the stress generated stiffness matrix.

The out-of-balance force vector drives the incremental displacement correction, which updates the strain field and, through the constitutive law, the stress state. This procedure is repeated through successive Newton iterations at constant temperature until equilibrium is satisfied within a prescribed tolerance. Convergence completes the time step, after which the solution advances to the next time increment.

The principle of this incremental–iterative equilibrium procedure is illustrated schematically in Figure 2.8.




Figure 2.8: SAFIR mechanical analysis iteration process [25].

Beam Finite Element Formulation

SAFIR includes several finite element formulations for mechanical analysis, such as truss, beam, shell, spring, and solid elements, allowing the modelling of a wide range of structural systems. Detailed descriptions of these elements are provided in the SAFIR documentation and the original reference paper.

The beam finite element is particularly suited for modelling frame-like structures and is therefore highlighted. The element represents a straight Euler-Bernoulli beam element with a constant cross-section along its length. Bending behaviour is described according to the Bernoulli hypothesis, whereby plane sections remain plane and normal to the deformed axis, while shear deformation is neglected.

The beam element is formulated with three nodes, in which two end nodes and one intermediate node. Each end node possesses seven degrees of freedom, accounting for translations, rotations, and warping, while the intermediate node includes one degree of freedom introduced to capture the nonlinear component of the axial displacement. This formulation enables an accurate representation of axial-bending interaction in geometrically nonlinear analyses.

The formulation assumes small strains within the element, while allowing for large displacements and rotations of the structure as a whole. Large rigid body motions are accounted for using a co-rotational formulation, which separates rigid body movements from deformational components. Axial forces, bending moments, and torsional effects are included.

The longitudinal stresses and stiffness are obtained by integrating the stress–strain response over the cross-section, which is discretised into fibres. Each fibre is assigned a temperature and material state, enabling a consistent representation of non-uniform temperature distributions across the section. Numerical integration along the beam axis is performed using Gauss integration with a limited number of integration points.

This fibre-based integration strategy ensures compatibility between the thermal and mechanical analyses, as the same sectional discretisation is used to evaluate temperature fields and mechanical response.

2.4. Gaussian Process Classification

Gaussian Process Classification (GPC) is a supervised, Bayesian, non-parametric framework for solving binary and multi-class classification problems. In GPC, a Gaussian Process is used to define a probabilistic prior over latent functions, which is subsequently combined with observed data through Bayes' theorem to obtain a posterior predictive distribution. This formulation allows classification decisions to be expressed in probabilistic terms, providing not only class predictions but also associated measures of uncertainty. Such uncertainty quantification is particularly valuable in engineering applications, where data are often limited and model predictions support risk-informed decision-making.

In this section, the theoretical foundations of Gaussian Processes are first introduced, followed by the formulation of Gaussian Process Classification and its extension to active learning strategies.

2.4.1. Gaussian Processes

The probabilistic foundation of Gaussian Processes is the multivariate Gaussian (normal) distribution. A D -dimensional random vector \mathbf{x} is said to follow a multivariate Gaussian distribution with mean vector $\mathbf{m} \in \mathbb{R}^D$ and covariance matrix $\Sigma \in \mathbb{R}^{D \times D}$ if its joint probability density function is given by [26]

$$p(\mathbf{x} \mid \mathbf{m}, \Sigma) = (2\pi)^{-D/2} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{m})^\top \Sigma^{-1}(\mathbf{x} - \mathbf{m})\right). \quad (2.22)$$

A key property of the multivariate Gaussian distribution is closure under marginalisation and conditioning, which enables analytical Bayesian inference.

A Gaussian Process (GP) extends this concept to functions. Formally, a Gaussian Process is defined as a collection of random variables, any finite subset of which follows a joint Gaussian distribution [26]. In supervised learning, these random variables correspond to the values of an underlying latent function evaluated at different input locations.

A Gaussian Process is fully specified by a mean function $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}, \mathbf{x}')$, defined as

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})], \quad (2.23)$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]. \quad (2.24)$$

The prior over the latent function $f(\mathbf{x})$ is written as

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (2.25)$$

For any finite set of input locations $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, the corresponding latent function values $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ follow a multivariate Gaussian distribution,

$$\mathbf{f} \sim \mathcal{N}(\mathbf{m}, \mathbf{K}(\mathbf{X}, \mathbf{X})), \quad (2.26)$$

where the mean vector $\mathbf{m} = [m(\mathbf{x}_1), \dots, m(\mathbf{x}_n)]^\top$ is obtained by evaluating the mean function at the input locations, and $\mathbf{K}(\mathbf{X}, \mathbf{X})$ is the covariance matrix obtained by evaluating the covariance function pairwise over the input set. The statistical dependence between latent function values is fully determined by the covariance matrix \mathbf{K} , which encodes prior assumptions such as smoothness and characteristic length scales.

Observed data are linked to the latent function through an observation model. In the graphical representation of a Gaussian Process, illustrated in Figure 2.9, observations are conditionally independent given the corresponding latent function values, while statistical dependencies between observations arise indirectly through the joint Gaussian prior over the latent function. This conditional independence structure allows the same Gaussian Process prior to be combined with different likelihood functions, forming the basis for both regression and classification models.

A defining property of Gaussian Processes is consistency under marginalisation, meaning that introducing additional input locations does not affect the distributions of existing variables. When observational information is available, Bayesian updating is performed by conditioning the joint Gaussian prior on the observed data.

In representation without the introduction of noise, the joint prior distribution of latent function values at the training inputs \mathbf{X} and at a set of test inputs \mathbf{X}^* is given by

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}^* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{m} \\ \mathbf{m}^* \end{bmatrix}, \begin{bmatrix} \mathbf{K}(\mathbf{X}, \mathbf{X}) & \mathbf{K}(\mathbf{X}, \mathbf{X}^*) \\ \mathbf{K}(\mathbf{X}^*, \mathbf{X}) & \mathbf{K}(\mathbf{X}^*, \mathbf{X}^*) \end{bmatrix}\right), \quad (2.27)$$

where \mathbf{m}^* denotes the mean function evaluated at the test inputs.




Figure 2.9: Graphical representation of a Gaussian Process with inputs \mathbf{x} , latent function values $f(\mathbf{x})$, and observed variables y [26].

Conditioning this joint prior distribution on the observed latent function values at the training inputs yields the posterior Gaussian Process. The resulting posterior mean and covariance at the test inputs are given by

$$\boldsymbol{\mu}^* = \mathbf{m}^* + \mathbf{K}(\mathbf{X}^*, \mathbf{X})\mathbf{K}(\mathbf{X}, \mathbf{X})^{-1}(\mathbf{f} - \mathbf{m}), \quad (2.28)$$

$$\boldsymbol{\Sigma}^* = \mathbf{K}(\mathbf{X}^*, \mathbf{X}^*) - \mathbf{K}(\mathbf{X}^*, \mathbf{X})\mathbf{K}(\mathbf{X}, \mathbf{X})^{-1}\mathbf{K}(\mathbf{X}, \mathbf{X}^*), \quad (2.29)$$

which together define the predictive distribution of the latent function under the Gaussian Process posterior [26].

As discussed above, the covariance function governs the dependencies between latent function values and encodes prior assumptions about the behaviour of the underlying function. In practice, covariance functions are parameterised by a set of hyperparameters, for example controlling the characteristic length scales and overall variance of the function. These hyperparameters are not fixed a priori, but are typically inferred from the data. Within the Gaussian Process framework, this is commonly achieved by maximising the log marginal likelihood, which balances data fit against model complexity through automatic Bayesian regularisation [26].

2.4.2. Formulation of Gaussian Process Classification

Gaussian Process Classification (GPC) extends the Gaussian Process framework to supervised classification problems by linking a latent Gaussian Process to discrete class labels through a probabilistic likelihood function. While the framework naturally generalises to multi-class classification, the following formulation is presented for the binary case, which is adopted throughout this work.

Consider a training dataset consisting of input–output pairs $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where $\mathbf{x}_i \in \mathbb{R}^d$ denotes the input variables and $y_i \in \{0, 1\}$ is a binary class label. As in the general Gaussian Process framework, an underlying latent function $f(\mathbf{x})$ is introduced and assigned a Gaussian Process prior according to Equation 2.25. The corresponding latent function values $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ therefore follow a multivariate Gaussian distribution with mean vector \mathbf{m} and covariance matrix $\mathbf{K}(\mathbf{X}, \mathbf{X})$, as described in the previous subsection.

In classification problems, the observed labels are not modelled as noisy real-valued observations of the latent function. Instead, class labels are linked to the latent function through a probabilistic likelihood that maps real-valued latent variables to class probabilities. For binary classification, this is achieved using a sigmoid-shaped link function, with common choices including the logistic (logit) and probit links [26]. An illustration of the logit and probit link functions is provided in Figure 2.10.

Adopting the probit formulation, the conditional probability of class membership is given by

$$p(y_i = 1 \mid f_i) = \Phi(f_i), \quad (2.30)$$




Figure 2.10: Logit and Probit sigmoid-shaped link functions

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. Accordingly, the likelihood for a single observation is expressed as a Bernoulli distribution,

$$p(y_i | f_i) = \Phi(f_i)^{y_i} [1 - \Phi(f_i)]^{1-y_i}. \quad (2.31)$$

Taking conditional independence of the class labels given the latent function values, the likelihood for the full dataset factorises as

$$p(\mathbf{y} | \mathbf{f}) = \prod_{i=1}^n p(y_i | f_i), \quad (2.32)$$

where $\mathbf{y} = [y_1, \dots, y_n]^\top$. Together with the Gaussian Process prior over \mathbf{f} , this defines a Bayesian classification model.

The posterior distribution over the latent function values is obtained by combining the prior and likelihood,

$$p(\mathbf{f} | \mathbf{y}, \mathbf{X}) \propto p(\mathbf{y} | \mathbf{f}) p(\mathbf{f} | \mathbf{X}). \quad (2.33)$$

Due to the Bernoulli likelihood introduced by a sigmoid-shaped link function, the resulting posterior distribution is no longer Gaussian and therefore not analytically tractable. Consequently, approximate inference methods are required to characterise the posterior. Common approaches include the Laplace approximation and Expectation Propagation (EP), which yield a Gaussian approximation to the posterior distribution over the latent function values [26].

Expectation Propagation (EP) approximates the intractable posterior distribution by a Gaussian through iterative moment matching, aiming to capture the posterior density globally. In contrast, the Laplace approximation constructs a Gaussian by performing a local second-order expansion of the log posterior around its mode, yielding a Gaussian centred at the maximum a posteriori estimate. A qualitative comparison of the Gaussian approximations obtained using EP and the Laplace approximation is illustrated in Figure 2.11 [27].




Figure 2.11: Laplace and EP approximation to a non-Gaussian posterior [27].

Specific to Gaussian Process Classification, EP operates on the non-Gaussian likelihood terms induced by the sigmoid-shaped link function. For each data point, EP constructs a so-called *cavity distribution*, obtained by temporarily removing the Gaussian approximation of the corresponding likelihood term from the global posterior approximation. The cavity represents the Gaussian belief about the latent function value based on all other observations. This cavity distribution is then combined with the exact sigmoid likelihood, and a new Gaussian approximation is obtained through moment matching. The interaction between the sigmoid likelihood, the cavity distribution, and the resulting Gaussian EP approximation is illustrated in Figure 2.12.

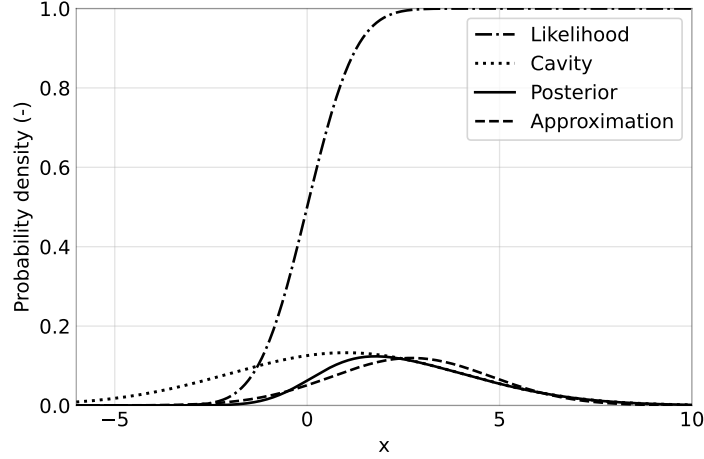


Figure 2.12: EP approximation of a sigmoid likelihood term in Gaussian Process Classification [26].

Under the EP Gaussian approximation, the marginal posterior distribution of the latent function value at a new input location \mathbf{x}^* , $f^* = f(\mathbf{x}^*)$, is formulated by,

$$q(f^* | \mathbf{X}, \mathbf{y}, \mathbf{x}^*) = \mathcal{N}(f^* | \mu_*, \sigma_*^2), \quad (2.34)$$

with posterior mean μ_* and variance σ_*^2 .

The predictive probability for the corresponding class label is then obtained by marginalising over the latent function value using this approximate posterior,

$$p(y^* = 1 | \mathbf{X}, \mathbf{y}, \mathbf{x}^*) = \int \Phi(f^*) q(f^* | \mathbf{X}, \mathbf{y}, \mathbf{x}^*) df^*. \quad (2.35)$$

For the probit likelihood, this integral admits a closed-form solution,

$$p(y^* = 1 | \mathbf{X}, \mathbf{y}, \mathbf{x}^*) = \Phi\left(\frac{\mu_*}{\sqrt{1 + \sigma_*^2}}\right), \quad (2.36)$$

which provides an explicit expression for the predictive class probability in Gaussian Process Classification [28].

2.4.3. Active Learning for Gaussian Process Classification

Active learning aims to reduce the number of labelled samples required to train an accurate classifier by actively selecting, at each iteration, the data point that is expected to be most informative. This selection is guided by an acquisition function, which assigns an utility value to candidate input locations. Rather than sampling training data uniformly at random, the acquisition function is optimised to identify the input whose label is expected to provide the greatest improvement to the model. Within the GPC framework, active learning naturally exploits the probabilistic predictive distributions and uncertainty estimates provided by the model.

Several acquisition functions have been proposed for GPC. Common uncertainty-based approaches include Maximum Entropy Sampling (MES) and Bayesian Active Learning by Disagreement (BALD).

MES selects query points with high predictive entropy, while BALD aims to maximise the mutual information between the predicted label and the model posterior. Although effective, these criteria do not directly target the reduction of classification error. In contrast, the approach covered in this section introduced by Zhao et al. [28], is based on Estimated Error Reduction (EER), which explicitly selects query points by estimating their expected impact on future classification performance.

In the EER framework, the quality of a classifier is measured in terms of its expected misclassification error under the *Optimal Bayesian Classifier* (OBC). The OBC refers to the theoretical Bayes-optimal decision rule that assigns to each input \mathbf{x}_s the class with the highest posterior probability,

$$\hat{y}_s = \arg \max_{y_s} p(y_s | \mathbf{x}_s). \quad (2.37)$$

For a given input \mathbf{x}_s , the corresponding misclassification probability is therefore

$$1 - \max_{y_s} p(y_s | \mathbf{x}_s), \quad (2.38)$$

since $\max_{y_s} p(y_s | \mathbf{x}_s)$ represents the probability assigned to the predicted class. In the binary case, this corresponds to selecting the larger of $p(y_s = 0 | \mathbf{x}_s)$ and $p(y_s = 1 | \mathbf{x}_s)$. If the predictive distribution is fully certain, the misclassification probability becomes zero. Otherwise, it reflects the remaining probability mass associated with an incorrect decision.

Let \mathbf{x}_s denote a test input drawn from the instance distribution $p(\mathbf{x}_s)$. The expected OBC classification error is then given by

$$\mathbb{E}_{\mathbf{x}_s} \left[1 - \max_{y_s} p(y_s | \mathbf{x}_s) \right]. \quad (2.39)$$

which represents the minimal achievable classification error under the current predictive distribution.

The Mean Objective Cost of Uncertainty (MOCU) is defined as the increase in this expected error due to uncertainty in the latent function. Active learning then aims to select a query \mathbf{x}^* that maximally reduces this uncertainty after observing the corresponding label y^* .

Formally, the expected reduction in MOCU induced by querying \mathbf{x}^* can be expressed as

$$U_M(\mathbf{x}^*) = \mathbb{E}_{\mathbf{x}_s} \left[\mathbb{E}_{y^* | \mathbf{x}^*} \left[\max_{y_s} p(y_s | \mathbf{x}_s, \mathbf{x}^*, y^*) \right] - \max_{y_s} p(y_s | \mathbf{x}_s) \right], \quad (2.40)$$

which serves as an EER-based acquisition function. This expression quantifies how the predictive distribution at future test points \mathbf{x}_s is expected to change after incorporating the new observation (\mathbf{x}^*, y^*) .

Direct optimisation of $U_M(\mathbf{x}^*)$ is challenging, as the presence of the maximisation operator yields a non-smooth objective. To enable gradient-based optimisation, a smooth approximation Soft-MOCU (SMOCU) is introduced by replacing the maximisation with a LogSumExp approximation. The resulting smooth acquisition function is given by

$$U_S(\mathbf{x}^*) = \mathbb{E}_{\mathbf{x}_s} \left[\mathbb{E}_{y^* | \mathbf{x}^*} \left[\frac{1}{k} \text{LogSumExp}(k \cdot p(y_s | \mathbf{x}_s, \mathbf{x}^*, y^*)) \right] - \frac{1}{k} \text{LogSumExp}(k \cdot p(y_s | \mathbf{x}_s)) \right], \quad (2.41)$$

where the parameter k controls the tightness of the approximation.

Applying EER to GPC is computationally demanding. Evaluating either $U_M(\mathbf{x}^*)$ or $U_S(\mathbf{x}^*)$ requires computing the predictive distribution $p(y_s | \mathbf{x}_s, \mathbf{x}^*, y^*)$ for both possible labels $y^* \in \{0, 1\}$ and for many samples \mathbf{x}_s . A naive implementation would retrain the GPC model for each candidate query and label, making it computationally infeasible as EP-based GPC inference scales cubically with the number of observations.

To address this computational limitation, two complementary strategies are introduced that reduce the cost of evaluating the acquisition function while preserving its theoretical formulation. The first strategy concerns efficient estimation of the expectation over \mathbf{x}_s , while the second avoids repeated retraining of the GPC model.

First, the expectation over \mathbf{x}_s is estimated using importance sampling, which can be written as

$$U(\mathbf{x}^*) = \mathbb{E}_{\mathbf{x}_s \sim p(\mathbf{x}_s)} [g(\mathbf{x}_s; \mathbf{x}^*)] = \mathbb{E}_{\mathbf{x}_s \sim \tilde{p}(\mathbf{x}_s; \mathbf{x}^*)} \left[\frac{p(\mathbf{x}_s)}{\tilde{p}(\mathbf{x}_s; \mathbf{x}^*)} g(\mathbf{x}_s; \mathbf{x}^*) \right]. \quad (2.42)$$

The variance of the estimator is reduced when the sampling distribution places more weight on regions that contribute most to the integral. In practice, a convenient and effective choice is $\tilde{p}(\mathbf{x}_s; \mathbf{x}^*) \propto k(\mathbf{x}_s, \mathbf{x}^*) p(\mathbf{x}_s)$, which exploits the locality of the covariance kernel. Under this choice, sample points that are weakly correlated with the candidate query \mathbf{x}^* receive little weight, reflecting the fact that they contribute less to the expected error reduction.

When the instance distribution $p(\mathbf{x}_s)$ is uniform over a bounded domain and a squared exponential (RBF) kernel is employed, the resulting importance sampling distribution $\tilde{p}(\mathbf{x}_s; \mathbf{x}^*)$ is proportional to a Gaussian density centred at \mathbf{x}^* . In this case, $\tilde{p}(\mathbf{x}_s; \mathbf{x}^*)$ can be implemented as a truncated Gaussian distribution over the domain of interest, allowing samples to be drawn efficiently from regions close to \mathbf{x}^* .

Second, retraining of the GPC model is avoided by expressing the updated predictive distribution through a joint predictive formulation. Rather than refitting the model after observing a hypothetical label y^* at a candidate query location \mathbf{x}^* , the updated predictive distribution at a test point \mathbf{x}_s is written as

$$p(y_s | \mathbf{x}_s, \mathbf{x}^*, y^*) = \frac{p(y_s, y^* | \mathbf{x}_s, \mathbf{x}^*)}{p(y^* | \mathbf{x}^*)}, \quad (2.43)$$

where all probabilities are conditioned on the currently observed dataset, which is omitted from the notation for brevity. This expression shows that the effect of a potential observation (\mathbf{x}^*, y^*) on future predictions can be captured through the joint distribution of the corresponding labels.

Under the EP Gaussian approximation, the joint distribution of the latent function values (f_s, f^*) remains Gaussian. As a result, the joint label probability $p(y_s, y^* | \mathbf{x}_s, \mathbf{x}^*)$ can be evaluated analytically without retraining the model. For the probit likelihood, this joint probability can be written as

$$p(y_s = 1, y^* = 1 | \mathbf{x}_s, \mathbf{x}^*) = \int \Phi\left(\frac{\tilde{\mu}_*(f_s)}{\sqrt{1 + \tilde{\sigma}_{**}}}\right) \Phi(f_s) \phi(f_s | \mu_s, \sigma_{ss}) df_s, \quad (2.44)$$

where $\phi(\cdot | \mu_s, \sigma_{ss})$ denotes the Gaussian density of the latent function value at \mathbf{x}_s . The quantities $\tilde{\mu}_*(f_s)$ and $\tilde{\sigma}_{**}$ are obtained from the joint Gaussian distribution of (f_s, f^*) evaluated at the input locations $(\mathbf{x}_s, \mathbf{x}^*)$, and follow from standard Gaussian conditioning rules.

By conditioning on one latent variable, the original two-dimensional integral over (f_s, f^*) is reduced to a one-dimensional integral over f_s . This reduction enables efficient, constant-time evaluation of the acquisition function for each candidate query point, without repeated model retraining.

Together, the importance sampling strategy and the joint predictive formulation enable practical implementation of the EER-based acquisition function. Based on these concepts, two active learning algorithms are formulated.

The first, NR-(S)MOCU-RO, employs random optimisation by sampling candidate query locations \mathbf{x}^* , estimating their acquisition values using importance sampling and the joint predictive distribution, and selecting the query that maximises the resulting utility.

The second algorithm, NR-SMOCU-SGD, uses the smoothness of the SMOCU acquisition function $U_S(\mathbf{x}^*)$ to enable gradient-based optimisation in continuous input spaces. Gradients of the acquisition function are obtained via the chain rule, using derivatives of the predictive mean and covariance with respect to the query location, which in turn depend on derivatives of the kernel function. A detailed derivation of the gradient expressions is provided by Zhao et al. [28].

The corresponding procedures are summarised in Table 2.1 and Table 2.2, respectively.

Table 2.1: NR-(S)MOCU-RO algorithm: n -th iteration, by Zhao et al. [29]

Step	Operation
1	function RandomOptimization($p(\mathbf{x}), q(f X, \mathbf{y})$)
2	Sample M_1 samples of $\mathbf{x}^* \sim p(\mathbf{x})$
3	for each \mathbf{x}^* do
4	Calculate $p(y \mathbf{x}^*)$ by (2.36)
5	Sample M_2 samples of $\mathbf{x}_s \sim \tilde{p}(\mathbf{x}_s; \mathbf{x}^*)$
6	for $y^* \in \{0, 1\}$ do
7	for each \mathbf{x}_s do
8	Calculate $p(y_s, y^* \mathbf{x}_s, \mathbf{x}^*)$ by (2.44)
9	Calculate $p(y_s \mathbf{x}_s)$ by (2.36)
10	Calculate posterior $p(y_s \mathbf{x}_s, \mathbf{x}^*, y^*) = p(y_s, y^* \mathbf{x}_s, \mathbf{x}^*) / p(y_s \mathbf{x}_s)$
11	Calculate $g(\mathbf{x}_s; \mathbf{x}^*)$, the integrand of (2.42)
12	end for
13	end for
14	$U(\mathbf{x}^*) = \frac{1}{M_2} \sum_{\mathbf{x}_s} p(\mathbf{x}_s) g(\mathbf{x}_s; \mathbf{x}^*) / \tilde{p}(\mathbf{x}_s; \mathbf{x}^*)$
15	end for
16	return $\tilde{\mathbf{x}} = \arg \max_{\mathbf{x}^*} U(\mathbf{x}^*)$
17	end function

Table 2.2: NR-SMOCU-SGD algorithm: n -th iteration, by Zhao et al. [28]

Step	Operation
1	function GradientOpt($p(\mathbf{x}), q(f X, \mathbf{y})$)
2	Obtain initial point \mathbf{x}^* from RandomOptimization($p(\mathbf{x}), q(f X, \mathbf{y})$)
3	while not converge do
4	Sample M_2 samples of $\mathbf{x}_s \sim \tilde{p}(\mathbf{x}_s; \mathbf{x}^*)$
5	Calculate $p(y^* \mathbf{x}^*)$ and $\nabla p(y^* \mathbf{x}^*)$ (see Zhao et al. [28])
6	for each \mathbf{x}_s do
7	for $y \in \{0, 1\}$ do
8	Calculate $p(y_s, y^* \mathbf{x}_s, \mathbf{x}^*)$, $\nabla p(y_s, y^* \mathbf{x}_s, \mathbf{x}^*)$ and $p(y_s \mathbf{x}_s)$ (see Zhao et al. [28])
9	end for
10	Calculate $\nabla g^S(\mathbf{x}_s, \mathbf{x}^*)$ (see Zhao et al. [28])
11	end for
12	$\nabla U^S(\mathbf{x}^*) = \frac{1}{M_2} \sum_{\mathbf{x}_s} p(\mathbf{x}_s) \nabla g^S(\mathbf{x}_s; \mathbf{x}^*) / \tilde{p}(\mathbf{x}_s; \mathbf{x}^*)$
13	Update \mathbf{x}^* with $\nabla U^S(\mathbf{x}^*)$
14	end while
15	return \mathbf{x}^*
16	end function

3

Literature Review

This chapter provides an overview of the existing literature relevant to the delayed failure of glulam timber columns under fire exposure. First, delayed failure is discussed in a general context, with emphasis on experimental investigations and their key findings. Subsequently, the concept of burnout resistance, together with the Duration of Heating Phase, is covered. Finally, the chapter reviews currently applied modelling studies for predicting the observed delayed failure behaviour of glulam timber columns. Together, these topics provide the basis for the identification of the research gaps addressed in this thesis.

3.1. Delayed Failure of Glulam Timber Columns

The collapse of load-bearing structural elements does not necessarily occur during the heating phase of a fire. Experimental observations and numerical studies have shown that failure may take place during the decay or cooling phase, after the gas temperatures in the compartment have started to decrease. This phenomenon is referred to as delayed failure and poses a significant risk for occupants and fire-fighting personnel, as structural instability may develop when the fire appears to be under control [30]. The risk of delayed failure is particularly pronounced for timber columns. Timber is characterised by a low thermal conductivity, which leads to a delayed penetration of heat into the cross-section. During the early stages of the decreasing fire exposure, while the outer surface layers begin to cool, heat continues to propagate towards the core of the member, resulting in a continued temperature increase in the central zones of the section for a considerable period of time [31]. At the same time, timber exhibits a fast reduction in mechanical properties with increasing temperatures, see Figure 2.3. Consequently, structural capacity continues to deteriorate even after the fire temperature has peaked.

This behaviour was first investigated through numerical simulations performed by Gernay using the finite element software SAFIR. Simulations on glulam timber columns subjected to parametric fire curves showed that columns designed according to traditional fire resistance ratings may fail during the decay phase, even when the heating duration is shorter than the rated fire resistance time. Although collapse occurs later than the fire resistance rating, fire exposure for a shorter heating duration does not guarantee structural survival. Delayed failure therefore constitutes a distinct, physics-based failure mode, with direct implications for the assessment of structural integrity in fire design.

The research clearly showed the difference between failure under standard fire exposure and delayed failure, as shown in Figure 3.1. For the column failing under ISO 834 fire conditions (a), high temperatures are concentrated near the outer layers, while the core remains relatively cool at the moment of failure. In contrast, for the column exposed to a short heating phase followed by cooling (b–d), failure occurs under a markedly different temperature distribution. At the time of collapse, the outer layers have cooled, whereas a substantially larger portion of the inner cross-section has reached elevated temperatures, ultimately leading to failure.

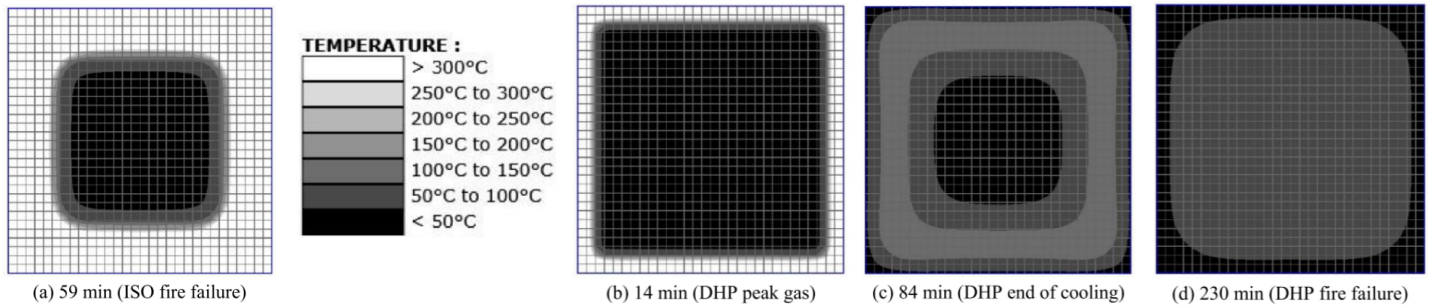


Figure 3.1: Cross-sectional temperature field for ISO 834 fire failure and delayed failure during the cooling phase from Gernay [12].

These results indicate that Eurocode-based fire assessment methods should extend beyond the traditional reliance on charring depth and the zero-strength layer when evaluating column failure. This is illustrated in Figure 3.2, which shows that isotherms in the range of 50–100 °C penetrate significantly deeper into the cross-section. As shown in Figure 2.3, temperatures of approximately 75 °C already result in a reduction of stiffness and strength of about 50 %.

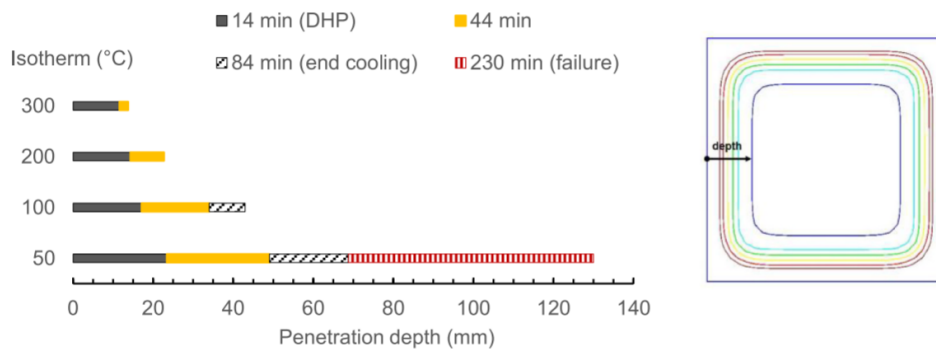


Figure 3.2: Isotherm depths for different time instances from Gernay [12].

To verify these numerical findings, a series of full-scale furnace tests was carried out on glulam timber columns exposed to parametric-shaped fire exposures, comprising an ISO 834 heating phase with varying heating durations and a subsequent controlled linear cooling phase. These experiments confirmed that columns are able to fail during the cooling phase when subjected to heating durations significantly shorter than their measured fire resistance time. Temperature measurements within the cross-sections further confirmed the sustained temperature increase in the core of the columns long after the end of the heating phase, in agreement with the previous numerical predictions [31]. An overview of the furnace tests performed is illustrated in Figure 3.3.

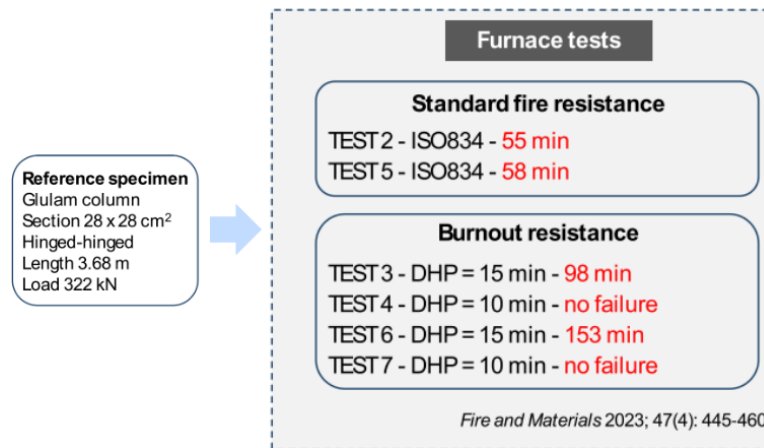


Figure 3.3: Overview of experimental furnace tests on glulam timber columns by Gernay et al. [31]

Subsequently, the investigation was extended to compartment fire tests with wood crib fuel loads, representing more realistic natural fire conditions. As these fires follow physically based temperature–time histories rather than the ISO 834 curve, direct comparison with standard fire resistance ratings is not applicable. Nevertheless, failures were observed during the decay and cooling phases of the fire, demonstrating that delayed failure is also a relevant risk in real fire scenarios [30, 32, 33]. In particular, delayed failure was observed in Tests 9, 10, 11, 12, 14, 18, and 19. Increasing the column cross-section resulted in longer survival times. As, for very large cross-sections, delayed failure was still observed long after fire extinguishment, driven by localised charring and continued smouldering, particularly near the column base. An overview of the natural fire tests is provided in Figure 3.4.

Overall, the combined numerical and experimental investigations demonstrate that delayed failure is a realistic and critical failure mode for glulam timber columns exposed to fire. The results clearly indicate that assessing fire performance solely on the basis of charring depth or standard fire resistance ratings is insufficient, as the progressive loss of strength and stiffness in the uncharred timber core plays a decisive role in post-peak fire stability.

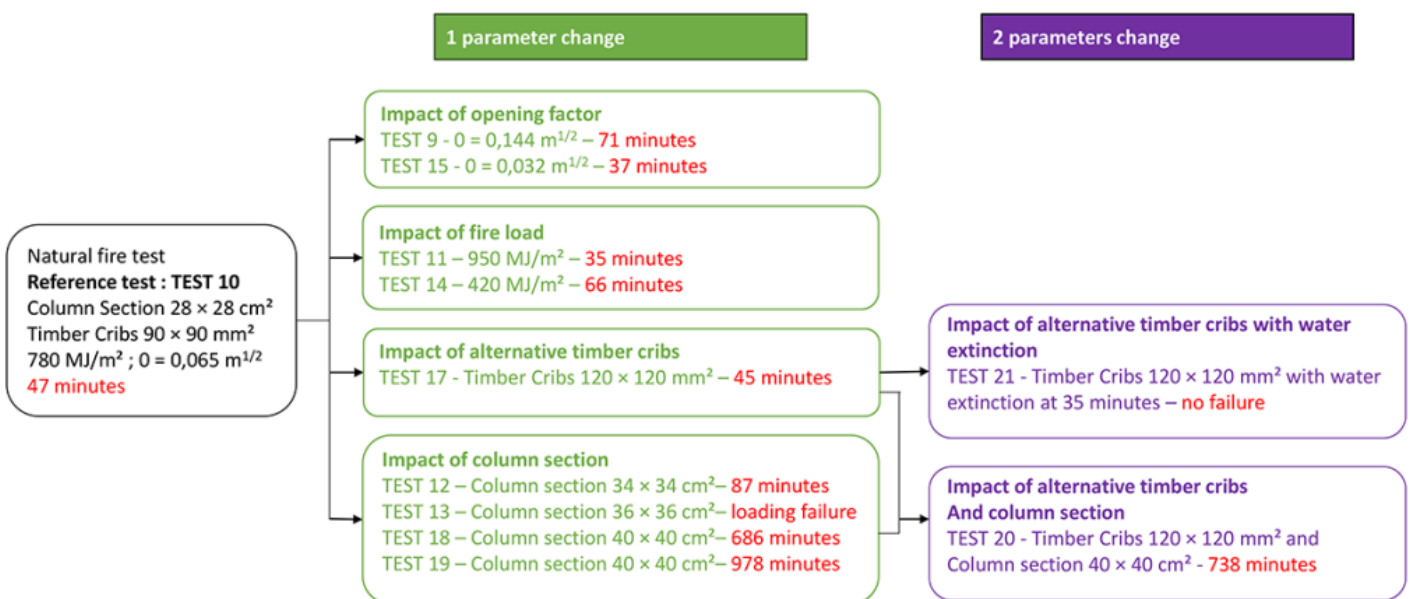


Figure 3.4: Overview of natural fire tests on glulam timber columns as illustrated in Renard et al. [33]

3.2. Burnout Resistance and the DHP Concept

The Duration of Heating Phase (DHP), introduced by Gernay and Franssen [11], is a standardised performance indicator to characterise the sensitivity of structural members to delayed failure. The indicator is defined using a set of standardised fire exposures based on the Eurocode parametric fire model with $\Gamma = 1$, yielding a heating branch that closely follows ISO 834, after which a prescribed cooling branch is applied. Within this framework, the DHP is the shortest heating duration (followed by the prescribed cooling) that will eventually lead to failure of the member, whether during heating, during cooling, or after the compartment temperature has returned to ambient. By definition, DHP is always smaller than or equal to the standard fire resistance rating, and it does not indicate the time of collapse, which may occur much later. Figure 3.5 provides an illustration of the DHP principle within the proposed standardised natural fire curves.




Figure 3.5: Illustration of the DHP principle by Robert et al. [30].

The DHP concept is used to define a complementary structural integrity metric, stated as the burnout resistance. In this terminology, the burnout resistance represents the binary outcome of either failure or survival of the structural member. Gernay and Franssen [11] propose reporting traditional fire resistance alongside burnout resistance to characterise performance under both continuous heating and complete fire events including a cooling phase.

Determining the burnout resistance (or DHP) requires an iterative procedure in which the member is analysed or tested under fires with increasing heating durations until failure is observed. Each analysis or test must be continued for a sufficiently long duration to ensure that any delayed collapse occurring during or after the cooling phase is properly captured. This procedure is typically implemented numerically, through thermo-mechanical FE simulations, but is also possible to be conducted experimentally. The iterative procedure is illustrated in Figure 3.6.

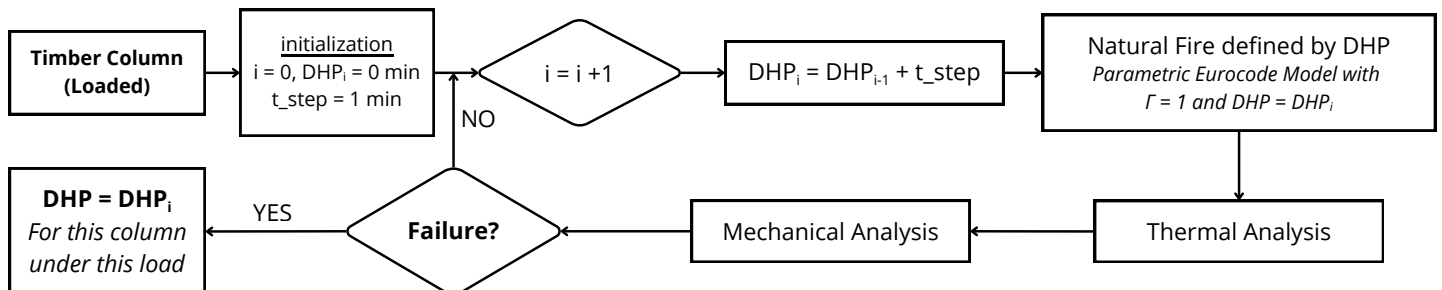


Figure 3.6: Iterative procedure for determining the burnout resistance of a structural member as mentioned in Gernay [34].

3.3. Delayed Failure Modelling of Glulam Timber Columns

As full-scale fire tests are costly, and the thermo-mechanical response of delayed failure is highly non-linear, numerical modelling is often considered as the approach to accurately describe such physical

problems. Numerical simulations enable detailed assessment of heat transfer, temperature-dependent material degradation, and structural stability.

The earlier mentioned numerical investigations in Section 3.1 by Gernay employed the finite element software SAFIR to model glulam timber columns exposed to both standard and parametrically shaped compartment fire curves. Using sequential thermal and mechanical analyses with temperature-dependent material properties, the study showed that SAFIR is able to reproduce the response of glulam timber columns tested under standard fire conditions by Stanke et al. [35], with conservative but reasonable accuracy. The same modelling framework was subsequently applied to fire scenarios including a cooling phase, where the simulations successfully captured delayed failure occurring well after the end of the heating phase [12].

With the availability of the furnace experiments shown in Figure 3.3, the finite element software SAFIR was applied to reproduce the experimental results and was found to provide reasonably accurate predictions. The thermal properties and temperature-dependent mechanical properties adopted from Eurocode 5 were sufficient to describe the response of glulam columns under parametric compartment fires, provided that material degradation was assumed to be non-reversible during cooling [36].

Despite their demonstrated reasonable accuracy for standard and standardised parametric fires with cooling phases, the applicability of numerical models to natural fire scenarios remains uncertain [30]. The thermal and mechanical material properties used in current models are largely calibrated against standard fire conditions, while experimental data describing timber behaviour during realistic natural fire decay phases remain limited [31]. Consequently, the validity of these modelling approaches for natural fires with varying heating and cooling characteristics has not yet been fully established.

In addition, an alternative numerical modelling approach was proposed by Han and Tesfamariam [37]. In this work, the Equivalent Section Temperature (EST) method implemented in Abaqus was introduced to enable reliability-based analysis of timber columns under fire exposure [37]. The study showed that this simplified thermo-mechanical approach, in which the cross-section is reduced to an equivalent uniform temperature, was able to reproduce the furnace tests reported by Gernay. Furthermore, within the EST framework, a surrogate modelling technique based on Polynomial Chaos Kriging was applied to quantify the influence of material uncertainty on the fire resistance of timber columns. The authors explicitly noted that extending surrogate modelling to fire scenarios including a decay phase introduces additional complexity, as structural members may either fail or survive.

4

Methodology

This chapter presents the methodology adopted in this study. First, it introduces the overall research approach, by describing the research workflow, the general modelling framework, and the justification for the chosen modelling approach. Second, it details the technical setup of the numerical modelling framework implemented in SAFIR, followed by the surrogate modelling framework based on Gaussian Process Classification.

4.1. Research Approach

This thesis adopts a model-based research approach that combines numerical simulations with a surrogate modelling technique in order to address the stated objectives of the research. The overall approach is to first evaluate the ability of numerical modelling to capture delayed failure behaviour in timber structures exposed to realistic fire scenarios followed by the examination whether this behaviour can be efficiently represented using surrogate models for assessment purposes.

The first stage of the research develops a non-linear thermo-mechanical finite element model to simulate the structural response of timber columns subjected to fire exposure, including both the heating and cooling phases. The numerical model is validated against the available experimental tests from Gernay et al. [31] and Renard et al. [32], as described in Section 3.1, and is assessed in whether it is capable to capture the governing structural response.

For this purpose, the non-linear finite element software SAFIR is employed. This software is specifically developed for the analysis of structures exposed to fire and enables thermal and mechanical analyses with temperature-dependent material behaviour. The numerical modelling framework is formulated in a parametric manner, allowing for systematic variation of input variables and efficient investigation of their influence on the structural response. Moreover, the parametric setup provides a consistent and suitable basis for the subsequent surrogate modelling stage.

The second stage of the research develops a surrogate modelling framework based on Gaussian Process Classification, using the parametrically defined numerical model as its foundation. Given that burnout resistance is characterised by a clear distinction between structural survival and failure, GPC is explored as a surrogate modelling technique to represent this binary response. By learning the relationship between governing input parameters and the resulting failure state, the surrogate model approximates the outcome of computationally expensive numerical simulations. An active learning strategy is integrated within the GPC framework to iteratively select the most informative simulation cases, thereby improving sampling efficiency and model accuracy.

Subsequently, the capabilities of the surrogate modelling framework are investigated to enable efficient assessments across a wider range of scenarios than would be feasible using numerical modelling alone. This includes investigating accuracy of parametric influences, supporting probabilistic interpretation of delayed failure behaviour, and evaluating the potential of surrogate models as tools for design-oriented assessment.

4.1.1. General Research Workflow

To ensure a structured and tractable research approach, an overall research workflow is established through which the study is organised into a sequence of defined phases. Each phase addresses a specific task and is completed before proceeding to the subsequent phase. This stepwise workflow facilitates systematic model development, validation, and analysis, while maintaining transparency in the adopted methodology.

An overview of the concerned phases is provided in Figure 4.1. Each individual phase and their underlying composition is described in detail in the subsequent paragraphs of this section.

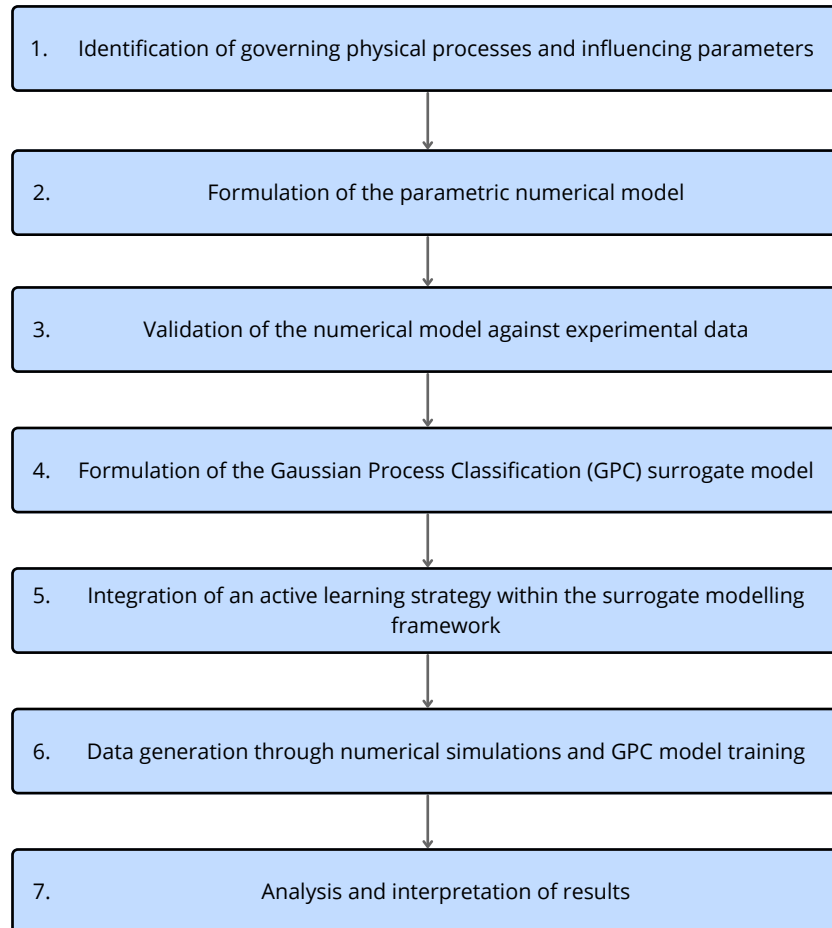


Figure 4.1: Methodological phases

1. Identification of governing physical processes and influencing parameters

The first phase focuses on identifying the governing physical processes associated with the fire behaviour of glued laminated timber columns and the mechanisms leading to delayed failure. Based on this understanding, the most influential design and modelling parameters are selected. These parameters define the parametric formulation of the numerical model and form the basis for the subsequent application of Gaussian Process Classification to relate variations in input parameters to the binary failure outcome.

2. Formulation of the parametric numerical model

The second phase focuses on the development of a numerical model capable of representing the structural response of glued laminated timber columns under fire exposure. In this phase, physical settings, modelling assumptions, boundary conditions, element formulations, and material modelling choices are defined to obtain a stable and physically consistent numerical setup. Particular attention is given to numerical robustness, computational efficiency, and the ability to capture the delayed failure phenomena. The model is configured in a parametric manner such that selected governing variables are

explicitly parameterised, while all other variables are consistently defined. This parameterisation enables efficient comparison against experimental tests, as well as systematic numerical data generation for the subsequent surrogate modelling stage.

3. Validation of the numerical model against experimental data

The third phase consists of assessing the capability of the numerical model against available full-scale experimental tests on glued laminated timber columns subjected to real fire exposures. The experimental programme considered in this study was conducted by the CERIB institute, with an overview presented in Figure 3.3 and Figure 3.4. The tests represent furnace and natural fire exposures, covering different fire curves and member geometries.

For this purpose, the previously defined numerical modelling framework is applied to the individual experimental configurations in which corresponding test inputs are implemented accordingly. Key response metrics, such as failure time and measured member temperatures, are extracted from the simulations and compared with the experimental observations. Through this process, the numerical modelling framework is benchmarked against the experimental tests, providing insight into its applicability and limitations.

4. Formulation of the Gaussian Process Classification (GPC) surrogate model

The fourth phase focuses on the formulation of the GPC surrogate modelling framework. This includes outlining the general characteristics of the GPC model and its specific application to the burnout resistance metric. Within this framework, the assumptions and modelling choices relevant to the present application are defined, allowing the surrogate model to be tailored to the specific user case of this research.

5. Integration of an active learning strategy within the surrogate modelling framework

The fifth phase focuses on integrating an active learning strategy within the GPC surrogate modelling framework to efficiently select new simulation inputs. The objective is to concentrate sampling in regions that are most informative for describing the decision boundary between failure and survival within the considered design space. The active learning algorithm, also known as the acquisition strategy, is developed and verified through a stepwise increase in input dimensionality. In addition, synthetic trial functions are employed to assess whether the combined active learning and GPC framework can accurately identify the decision boundary, select new points efficiently, and determine when sufficient convergence has been achieved.

6. Data generation through numerical simulations and GPC model training

The sixth phase involves applying the defined GPC surrogate modelling framework in combination with the active learning strategy to the parametric numerical model. New simulation points are generated sequentially, with numerical simulations performed at selected input locations to provide training data for the GPC model. This iterative process enables efficient learning of the decision boundary in the physical parameter space, distinguishing between column failure and survival, while explicitly accounting for the full fire exposure, including the cooling phase.

7. Analysis and interpretation of results

The final phase focuses on analysing and interpreting the results obtained from the trained GPC surrogate models. The predictive capabilities of the surrogate models are explored with respect to both research-oriented analysis and design-oriented assessment. In particular, the potential of the proposed framework for efficient evaluation of burnout resistance and its applicability to practical design assessments are investigated and discussed.

4.1.2. General Modelling Framework

As described in the research approach, this study is structured around two interconnected modelling stages: a parametric numerical modelling stage and a surrogate modelling stage. The interaction between these stages is illustrated schematically in Figure 4.2. Within this framework, a set of modelling variables is defined and provided as input to the numerical simulations. For a given combination of input variables, the numerical model evaluates the structural response of the column under fire exposure and produces a binary outcome indicating either failure or survival.

The resulting input–output pairs form the training data for the supervised Gaussian Process Classification model. Subsequently, based on the fitted surrogate model, an active learning algorithm selects

a new combination of modelling variables that is most informative for refining the decision boundary between failure and non-failure within the design space. This procedure continues iteratively until sufficient convergence of the classification boundary is achieved.

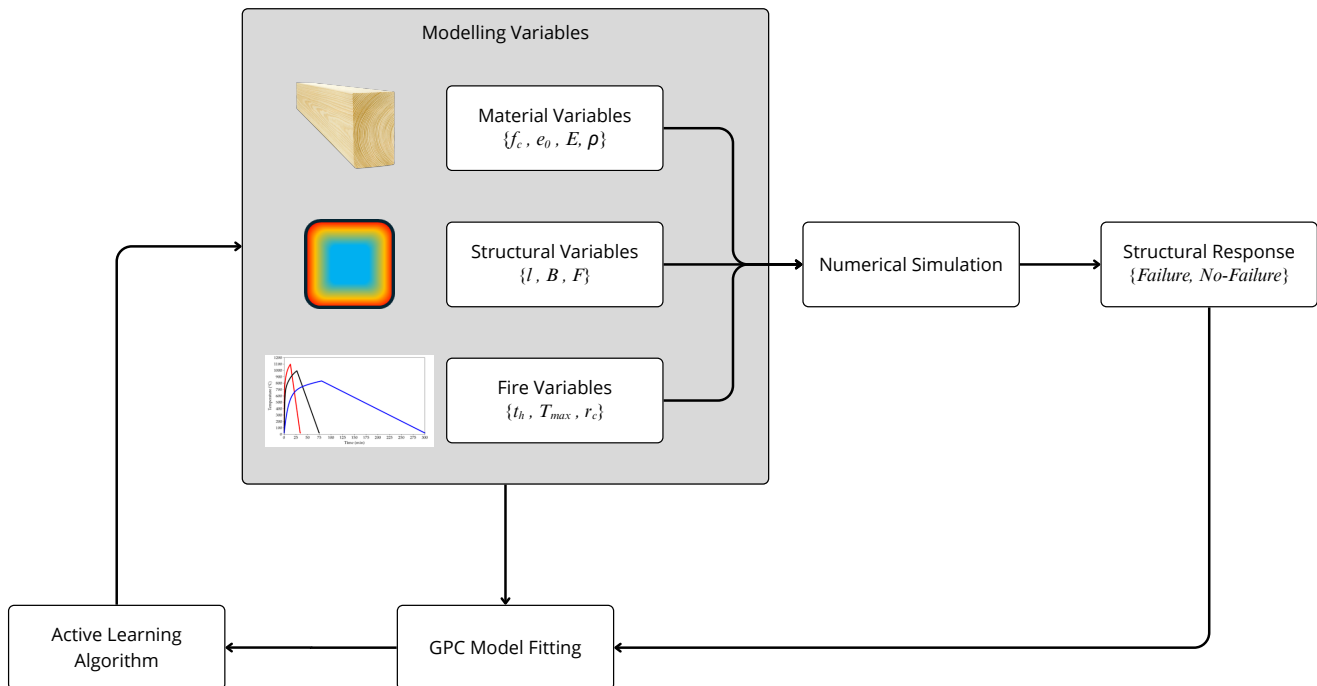


Figure 4.2: General modelling framework

The numerical modelling stage captures the governing physical processes associated with glued laminated timber columns subjected to realistic fire scenarios, while the surrogate modelling stage provides an efficient approximation of the ultimate structural response. In this framework, the numerical model acts as a data-generation engine, whereas the surrogate model together with the active learning algorithm provide efficient exploration and exploitation across the design space.

The modelling variables are grouped into material, structural, and fire-related parameters, reflecting the primary physical influences on burnout resistance.

The material variables describe the mechanical and thermal properties of the timber and are represented by the compressive strength f_c , eccentricity e_0 , Young's modulus E , and density ρ . While a single material design configuration is adopted in this study, these parameters are still treated as modelling variables to account for inherent material variability at the evaluation stage. As such, they represent stochastic properties that are not directly controlled by the designer but influence both the mechanical response and the thermal behaviour of the member throughout the fire exposure.

The structural variables define the column geometry and loading conditions and are treated as deterministic design parameters. These include the column length l , cross-sectional width B , and applied axial load F . Together these parameters characterise the structural configuration relevant for design-oriented assessment.

The fire exposure is described through a set of fire variables chosen to approximate a wide range of realistic fire scenarios, explicitly including both heating and cooling phases. The framework of the parametric fires described in Subsection 2.2.2 is used. The variables comprise the heating duration t_h , the maximum compartment temperature T_{max} , and the cooling rate r_c . A more detailed representation of the fire characteristics is provided in Subsection 4.2.1.

4.1.3. Justification of the Modelling Approach

The broader aim of this research is to improve the understanding and assessment of delayed failure and burnout resistance of glued laminated timber columns exposed to realistic fire scenarios, including both heating and cooling phases. With this objective in mind, an accurate modelling approach that is able

to provide the structural response across multiple influencing variables creates valuable opportunities. Such an approach enables extensive parametric studies, sensitivity analyses, and probabilistic assessments. However, the underlying physical problem is characterised by multiple interacting processes, including fire exposure, thermal response and mechanical behaviour. Consequently, the adopted modelling approach must be capable of capturing complex physical phenomena while allowing systematic exploration of a multi-dimensional parameter space.

A purely numerical investigation based on finite element simulations offers the necessary physical fidelity to represent temperature-dependent material behaviour and delayed failure mechanisms. However, such simulations are computationally expensive, which limits their suitability. Conversely, purely data-driven approaches provide computational efficiency but lack direct control over the underlying physical representation. The modelling framework adopted in this thesis therefore combines the complementary strengths of both approaches. High-fidelity numerical simulations are used to represent the governing physical processes, while a surrogate modelling layer enables efficient exploration of the resulting input–output relationships across the design space.

The burnout resistance problem naturally lends itself to a classification-based formulation, as the structural response can be expressed in terms of a binary failure state. Gaussian Process Classification is therefore a suitable surrogate modelling approach, as it is able to represent non-linear decision boundaries while providing a Bayesian framework that quantifies predictive uncertainty. This uncertainty information reflects the confidence of the surrogate model in its predictions and can be exploited for both model assessment and the implementation of active learning strategies. In addition, Gaussian Process models are well suited to low-data regimes, which is essential in the present application where each labelled sample corresponds to a computationally expensive numerical simulation. Hyperparameters can be inferred within the probabilistic framework through marginal likelihood optimisation, avoiding the need to reserve a separate validation dataset and thereby preserving scarce simulation data for model training.

The modelling variables considered in this framework are selected to represent the dominant physical influences on delayed failure. Accordingly, structural design variables, fire characteristics, and material properties are included. Structural variables reflect design-related influences and material properties account for inherent uncertainty. Fire exposure is parameterised to represent a wide range of realistic scenarios, explicitly including cooling phases, which remains adaptable to evolving insights in timber compartment fire behaviour.

Within the scope of this thesis, the combined numerical–surrogate modelling framework provides a balanced and efficient approach for addressing the overall research aim. Alternative modelling strategies were considered during the development of the research approach but were not adopted due to limitations related to efficiency, or applicability. These alternatives are discussed in Appendix A

4.2. SAFIR Numerical Modelling Framework

The numerical modelling framework adopted in this study is based on the non-linear finite element software SAFIR and is used to simulate the thermal and mechanical response of glued laminated timber columns exposed to fire. SAFIR is widely applied in structural fire engineering and employs a weakly coupled analysis approach, in which the thermal response is computed first and subsequently used as input for the mechanical analysis. A more detailed description of the underlying modelling principles is provided in Section 2.3.

This section describes the numerical modelling framework as implemented in SAFIR, including the adopted modelling assumptions, parameterisations, and configurations used to represent delayed failure behaviour. The following subsections present the setup of the numerical model in detail. An overall general overview of the numerical modelling framework is illustrated in Figure 4.3.

The code developed for this thesis is publicly available in the corresponding GitHub repository [38].

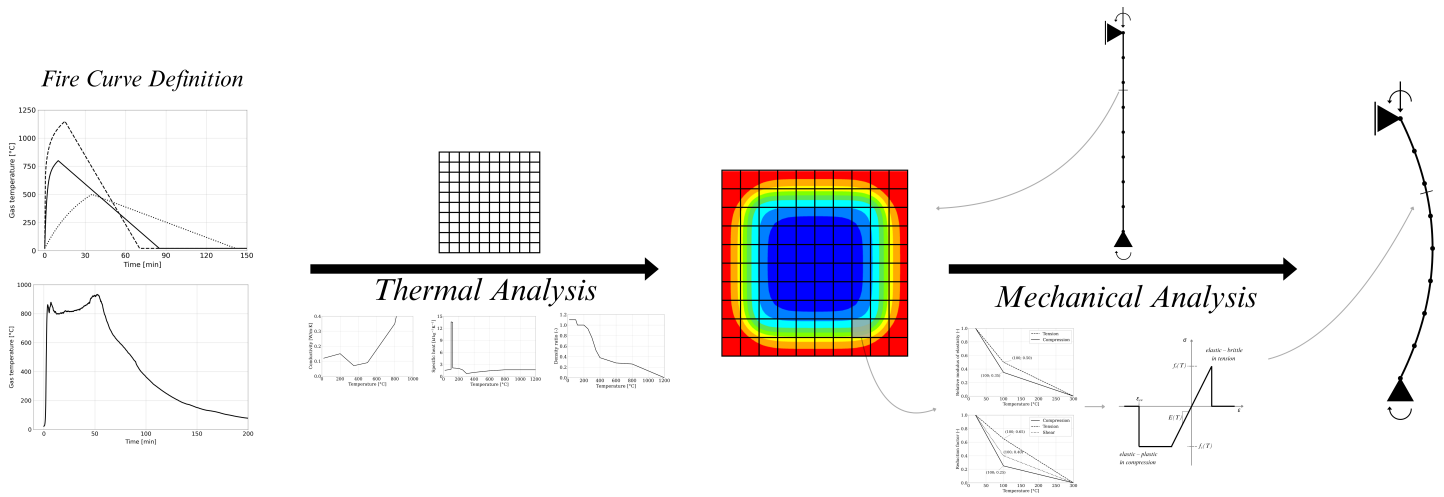


Figure 4.3: Numerical simulation framework

4.2.1. Fire Curve Definition

To investigate the influence of fire-curve characteristics on the structural response, and to approximate a wide variety of observed fire scenarios, the gas-temperature fire curve is parametrically defined. A parametric definition enables systematic variation of the key characteristics and thus facilitates the generation of a wide range of numerical simulations in a consistent and controlled manner. As a basis, the general shape of the Eurocode parametric fire curve is adopted, consisting of an exponential heating phase followed by a linear cooling phase.

The fire curve is governed by the following three primary parameters:

- Heating duration, t_h ;
- maximum gas temperature, T_{\max} ;
- and a linear cooling rate, r_c .

Together with the ambient temperature T_0 , these parameters fully define the fire curve. The ambient temperature is taken as 20 °C.

In addition, based on the primary parameters, the corresponding end of the cooling phase (t_c) is defined.

The end of the cooling phase is given by:

$$t_c = t_h + \frac{T_{\max} - T_0}{r_c}. \quad (4.1)$$

Following, a schematic overview of the defined parameters is shown in Figure 4.4.




Figure 4.4: Definition of the fire curve parameters.

The heating phase follows the Eurocode parametric definition and provides a physically grounded temperature–time relationship. It is expressed as:

$$T_g(t) = T_0 + 1325 \left(1 - 0.324e^{-0.2t^*} - 0.204e^{-1.7t^*} - 0.472e^{-19t^*} \right), \quad (4.2)$$

with

$$t^* = t \cdot \Gamma,$$

where $T_g(t)$ is the gas temperature ($^{\circ}\text{C}$), t the time (h), and Γ a non dimensional time-scaling parameter [20].

Since the heating phase in Equation 4.2 is a sum of exponential terms, Γ cannot be determined analytically. The scale parameter is therefore determined numerically with a robust root finding bisection method by enforcing:

$$T_g(t_h) = T_{\max}.$$

To solve for Γ , the corresponding scaled time t_h^* is determined first, after which:

$$\Gamma = \frac{t_h^*}{t_h}. \quad (4.3)$$

Defining the exponential part as described in Equation 4.2 as:

$$\phi(t^*) = 1 - 0.324e^{-0.2t^*} - 0.204e^{-1.7t^*} - 0.472e^{-19t^*},$$

and the target value as:

$$y = \frac{T_{\max} - T_0}{1325},$$

the applicable function becomes:

$$f(t^*) = \phi(t^*) - y,$$

for which the following root needs to be satisfied:

$$f(t_h^*) = 0.$$

As the bisection method requires an interval $[a, b]$ such that $f(a)f(b) < 0$, the lower bound is taken as $a = 0$, yielding $f(a) = -y < 0$, and b is incrementally increased until $f(b) > 0$. The bisection algorithm is then applied to determine t_h^* and hence Γ .

With Γ determined, the full fire curve can be expressed as:

$$T_g(t) = \begin{cases} T_0 + 1325 \cdot \phi(t^*), & 0 < t \leq t_h, \\ -r_c(t - t_h) + T_{\max}, & t_h < t < t_c, \\ T_0, & t > t_c. \end{cases} \quad (4.4)$$

As can be seen from Equation 4.2, the heating phase approaches an asymptote at $T_0 + 1325$ °C, causing that the choice of T_{\max} is restricted to the interval:

$$T_0 < T_{\max} < T_0 + 1325.$$

When $\Gamma = 1$, the heating branch corresponds closely to the ISO 834 standard fire curve. However one should note that the ISO 834 fire curve increases indefinitely, unlike the Eurocode parametric heating definition that asymptotically stabilises. Figure 4.5 illustrates the influence of Γ and shows the comparison with the ISO 834 curve.

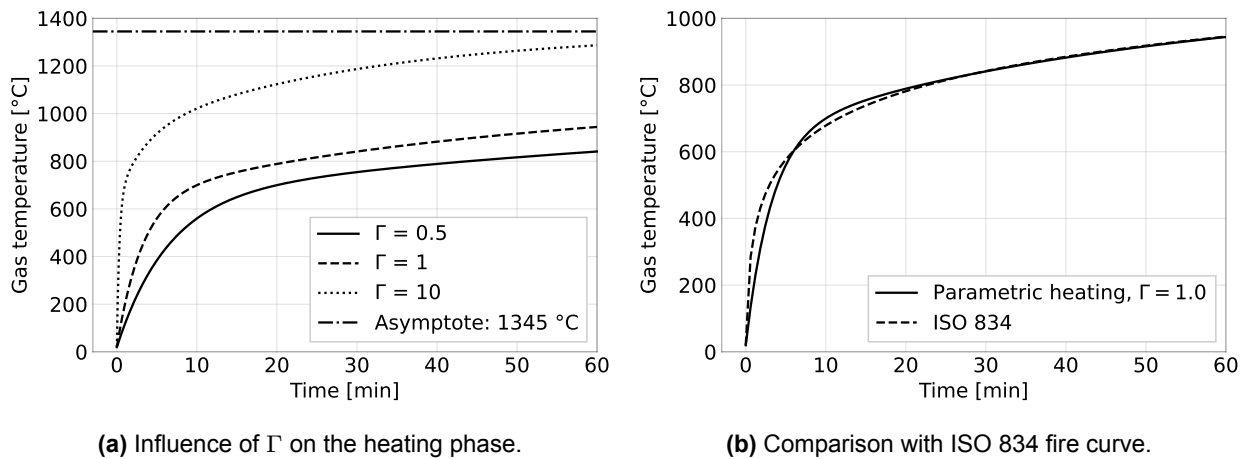


Figure 4.5: Effect of the Γ -factor on the fire curve heating phase.

4.2.2. Thermal Analysis

The thermal analysis is conducted on a two-dimensional representation of the column cross-section, where heat transfer is assumed to occur solely through conduction within the material. The governing conduction equation is presented in Subsection 2.1.2 by Equation 2.2. No additional heat source or sink is included.

The external boundary conditions are prescribed by the gas-temperature fire curve described in Subsection 4.2.1, with both radiation and convection heat transfer included at the exposed surfaces as described by Equation 2.9

The thermal analysis is defined parametrically with respect to the cross-sectional width, enabling consistent simulation of different geometries. A description of the underlying principles of the thermal analysis of SAFIR is provided in Subsection 2.3.1.

Setup of the Thermal Analysis

Considered is a glulam timber column of strength class GL24h made from spruce. The density is chosen based on mean values provided in Eurocode NEN-EN 14080:2013 Table 5 [39] and research conducted on the probabilistic mechanical properties of glued laminated timber made of softwood by Schilling et al. [40]. This implies a mean density of 420 kg m^{-3} . The moisture content is taken at 12% in accordance with information from two independent glulam suppliers [41, 42]. In SAFIR, the thermal material module WOODDEC5 is adopted as one of the thermal modelling options considered in this work.

It is available by default in the software and uses the effective thermal properties specified in Annex B of EN 1995-1-2:2004 [1], as illustrated in Figure 2.2.

Importantly, the thermal properties are assumed to be irreversible during cooling, meaning that each fibre retains the values corresponding to the maximum temperature previously attained.

The grain direction of the timber is assumed to be perpendicular to the cross-section considered. Timber is an anisotropic material, meaning that its thermal conductivity along the grain differs from that across the grain. The ratio of anisotropy in thermal conductivity is taken as 4. However, since the grain direction is oriented perpendicular to the analysed cross-section, this anisotropy has no effect on the present analysis. Consequently, all thermal material properties are considered transverse to the grain. The thermal conductivity of timber shown in Figure 2.2a therefore represents the conductivity transverse to the grain direction.

For the boundary conditions, the convective heat-transfer coefficient is set to $35 \text{ W m}^{-2} \text{ K}^{-1}$ when the surface is exposed to fire temperatures above ambient, and reduced to $4 \text{ W m}^{-2} \text{ K}^{-1}$ once the gas temperature returns to ambient conditions. A surface emissivity of 0.8 is adopted throughout. All values are taken in accordance with EN 1991-1-2 [20]. The column is taken to be fully exposed to fire on all four sides.

Additionally, the thermal analysis in SAFIR is controlled through several numerical parameters that define the time integration and convergence behaviour of the transient solution. A dynamic time-stepping scheme is adopted, with a shortest accepted time step of $\text{COMEBACK} = 3.75$ seconds, an initial time step of 7.5 seconds, and a maximum time step of 30 seconds. The integration parameter is set to $\text{TETA} = 0.9$, corresponding to an implicit time integration scheme, and the convergence tolerance of the iterative process is defined as $\text{PRECISION} = 1 \times 10^{-3}$. Output results are written every 60 seconds of simulated time. These numerical settings ensure a stable and accurate solution of the non-linear thermal problem while maintaining computational efficiency throughout the analyses.

Thermal Mesh Definition

The two-dimensional cross-section is discretised using rectangular finite elements consisting of four nodes per element. Each node has a single degree of freedom, representing the nodal temperature. Within each element, temperature gradients are evaluated at four integration points, two in each direction.

To accurately capture the higher temperature gradients that develop near the edges during fire exposure, the mesh is refined in the corner regions of the cross-section. A coarser mesh is used in the central region.

The cross-section is divided into subregions corresponding to quarters of the total width, as illustrated in Figure 4.6. Within these outer quarters, a fine mesh spacing, denoted as *fine_size*, is adopted. The inner half of the section is meshed with elements that are twice as large denoted as *coarse_size*. This arrangement results in a smooth transition between the fine and coarse regions while maintaining computational efficiency.

Special attention is given to the mesh definition since the column width varies in the parametric thermal analysis. The mesh is defined either in absolute dimensions or relative to the overall section width. In this study, an absolute mesh size is adopted to maintain accuracy for larger cross-sections, where a relative definition could otherwise lead to an overly coarse mesh discretisation and underestimation of thermal gradients. It should be noted, however, that the chosen values of *fine_size* and *coarse_size* must divide evenly within the quarter-width regions to prevent irregularities or boundary misalignment.

A schematic representation of the mesh layout and refinement pattern is shown in Figure 4.6a.

To reduce computational effort, only one quarter of the cross-section is modelled by exploiting geometrical symmetry. Along the symmetry lines, the boundary conditions are defined with no heat flux across the boundaries. This approach does not introduce any loss of accuracy, as the temperature field remains symmetric across the section. A visual representation of this setup is provided in Figure 4.6b. The use of symmetry is particularly beneficial since the algorithm sequentially performs multiple thermal analyses in the surrogate modelling phase, making computational efficiency essential. The achieved performance gain is further quantified in Appendix B, where it is shown that the symmetry model runs more than twice as fast as the full model, with even greater gains for non-failing columns.

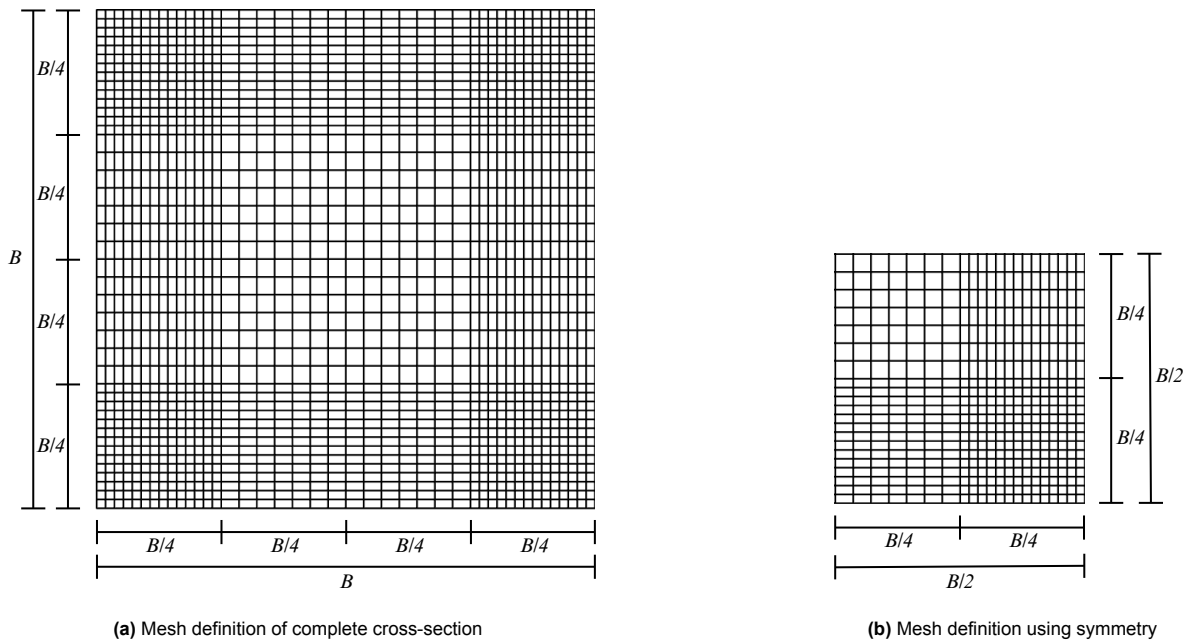


Figure 4.6: Mesh definition of the two-dimensional cross-section with refined corner regions.

To ensure sufficient accuracy in the thermal analysis, an appropriate mesh density was determined. The aim was to identify a mesh size that achieves adequate numerical accuracy while maintaining reasonable computational effort. A mesh sensitivity analysis was therefore conducted to define the `fine_size` parameter used in the thermal simulations. The complete study and results are presented in Appendix C. Based on these findings, the `fine_size` parameter was found to be set approximately to 5 mm, providing a balanced compromise between accuracy and computational efficiency.

However, as previously discussed, the thermal analysis is performed parametrically with respect to the cross-sectional width. Since an absolute mesh size is adopted, the `fine_size` value varies slightly for each cross-sectional width to ensure proper alignment of the mesh. Specifically, the `fine_size` and `coarse_size` values must fit exactly within the defined regions representing one quarter of the section width. To achieve this, an algorithm is applied that guarantees sufficient accuracy and consistent fitting of the mesh divisions.

The algorithm takes the given cross-sectional width, which may take any value, and repeatedly divides it by a factor of two until a value smaller than 5 mm is obtained. Subsequently, the value closest to the determined reference value is selected. This resulting value is then assigned to the `fine_size` variable, ensuring that the mesh resolution remains both accurate and properly structured across all cross-sectional dimensions.

Determination of Thermal Simulation Time

Determining the correct end time of the thermal analysis is crucial, as the subsequent mechanical analysis directly relies on its results. To capture potential delayed failure, the simulation must extend until the entire heat wave has propagated through the cross-section. The end time is therefore defined as the moment when the central node begins to cool, indicating that the maximum thermal penetration has been reached. If the column remains stable at this stage, it is expected that the column has survived the fire exposure, and no delayed failure occurs. This reasoning, however, includes the assumption of auto-extinguishment.

As SAFIR requires a predefined simulation time rather than a condition-based stopping rule, two methods were developed to capture this moment properly. A restart-based numerical approach and a simplified analytical model.

The restart method combined the restart functionality of SAFIR with a post-processing routine that detects cooling of the central node. However, repeated restarts introduced underestimation errors due to the loss of the irreversible thermal history of each element. For more detailed insights of these observations, see Appendix D.

The analytical approach used a one-dimensional heat-conduction model representing half of the cross-section, with a symmetry boundary at the centre and a prescribed fire temperature at the exposed surface. Constant thermal properties were assumed, and radiation at the boundary was excluded, as these introduce non-linearities that prevent an analytical solution. Radiation could, however, be incorporated in a linearised form, and cases with solely convection or with convection and linearised radiation were examined. However, no significant improvement was obtained compared with a simpler Dirichlet boundary condition. The applied fire curve was approximated using the fire curve shown in Figure 4.7. Although the analytical model showed deviations in predicting the moment at which the central node began to cool, applying a conservative factor of 1.3 to its outcome proved to give a reliable first estimate of the required thermal analysis time. Despite its simplifications, the analytical model uses the same physical inputs as the non-linear SAFIR simulations and therefore still provides a meaningful and practical basis for obtaining an initial estimate of the cooling time. A full derivation of the analytical approach and further discussion are provided in Appendix E.

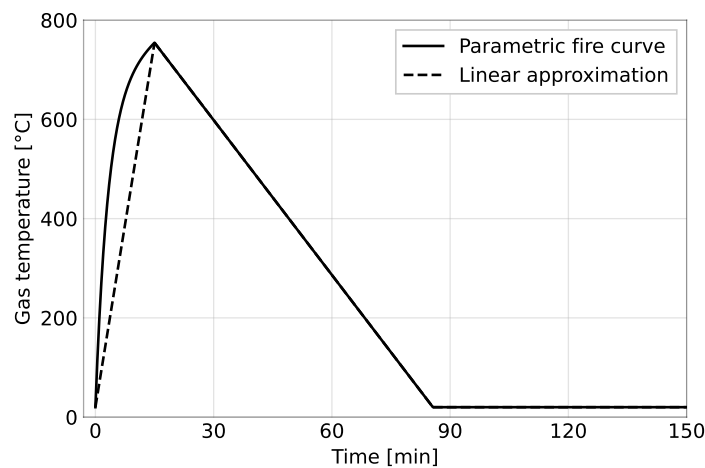


Figure 4.7: Approximated fire curve used as thermal boundary condition in the simplified analytical approach.

Because neither approach was fully sufficient, a hybrid strategy was adopted. As a first estimate, the simplified analytical approach is used to determine the expected time at which the central node starts to cool. Based on this estimate, the thermal analysis is performed with a post-processing procedure that checks whether cooling of the central node has occurred. If not, the analysis is extended with restarts of one-hour increments until cooling is detected. An iterative algorithm is then applied, in which the end time identified from the previous analysis with restarts serves as the new initial guess for the next thermal simulation. This process continues until a complete thermal analysis is executed without restarts. In this way, the final simulation reliably captures the full thermal response of the cross-section, providing a consistent basis for the subsequent mechanical analysis.

Expanded Thermal Model

The temperature-dependent thermal properties presented in Figure 2.2 are suitable for representing the heating phase of a standard ISO fire, corresponding to a Γ -factor equal to one. However, their applicability becomes uncertain under parametric fires, which include a cooling phase and may exhibit different heating phases corresponding to different Γ -factors.

To address these limitations, an expanded thermal model was implemented in this study, drawing on the heat-of-hydration approach introduced by Hopkin [43]. This expanded model improves the description of the thermal response of timber members by incorporating a modified conductivity model that accounts for accelerated heating, together with an additional heat contribution intended to compensate for the limitations of the temperature-dependent material properties during the cooling phase.

The thermal conductivity model proposed by Hopkin was adopted directly. In this model, a conductivity multiplier is introduced to account for the influence of the heating phase on the thermal response and is expressed as

$$k_{\lambda, \text{mod}} = 1.45 \Gamma^{-0.48}. \quad (4.5)$$

This multiplier is applied in combination with the temperature-dependent conductivity values shown in Table 4.1.

As in the original publication of Hopkin, the specific heat capacity of Cachim and Franssen and the Eurocode-based temperature-dependent density model were employed to ensure consistency with the underlying assumptions of the method [43].

Temperature (°C)	Conductivity (W m ⁻¹ K ⁻¹)
20	0.12
200	0.15
350	0.07
500	0.09 $k_{\lambda, \text{mod}}$
800	0.35 $k_{\lambda, \text{mod}}$
1200	1.50 $k_{\lambda, \text{mod}}$

Table 4.1: Temperature-dependent thermal conductivity from Hopkin et al. [44].

In addition, Hopkin's heat-of-hydration method introduces a volumetric heat source to represent the burning of timber and to improve predictions during the cooling phase. However, the numerical software SAFIR does not allow the inclusion of such a source term directly in the governing heat equation. For this reason, an alternative implementation strategy was adopted. Instead of adding a volumetric term in the governing equation, an equivalent additional boundary heat flux was introduced. This flux is applied at the position of the burning timber, which is taken to coincide with the 300 °C isotherm. As the member chars, this boundary moves inward, resulting in a "walking" boundary condition that mimics the energy release associated with material combustion.

The position of this moving boundary is determined from the cumulative-temperature charring equation (A.4.3.2(2)) from the draft of Eurocode 5 [45]. The equation is given by

$$d_{\text{char},t} = \left(\frac{\int_0^t (T^2) d\tau}{1.35 \times 10^5} \right)^{\frac{1}{1.6}}, \quad (4.6)$$

in which $d_{\text{char},t}$ is the charring depth (mm), T the compartment gas temperature (K), and t the time (min).

The integral is evaluated numerically using the trapezoidal rule, after which the charring rate $\beta(t)$ is obtained by differentiating the depth curve using a central-difference scheme.

To assess the validity of this approach, the calculated charring depth is compared with experimental measurements obtained from a set of natural fire tests conducted by Gernay et al. [30, 32, 33]. The predicted charring depth is derived from the natural fire curves recorded in these tests, while the experimental measurements represent the depth of the 300 °C isotherm. The resulting comparisons, presented in Appendix F, show that the model provides a sufficiently accurate approximation for the purposes of this study.

It is important to recognise that the cumulative-temperature model implies that charring continues indefinitely. Because the expression integrates temperature over time, the predicted charring depth will always increase as time progresses. In reality, however, it is questionable whether charring continues once the thermal exposure becomes too low to sustain further degradation. For this reason, a practical stopping criterion is introduced. Drawing on observations from the natural fire experiments of Gernay et al., charring is assumed to stop once the compartment gas temperature falls below 300 °C after previously exceeding this value. This reflects the physical understanding that, below this threshold, the timber does not receive enough heat to maintain continued charring. This approach implicitly relies on the assumption of auto-extinguishment and uses the 300 °C isotherm as the reference for the charring front.

The limited penetration of the 300 °C isotherm is particularly evident in Test 18 of the natural fire series by Gernay et al. Further details on these observations are provided in Appendix F.

Following, the additional heat flux applied at the burning front is defined as

$$-k(T) \frac{\partial T}{\partial n} \Big|_{\Gamma(t)} = q_{\text{comb}}(t), \quad (4.7)$$

where $\Gamma(t)$ denotes the moving boundary of the 300 °C isotherm, and q_{comb} the combustion-related flux (W m^{-2}), which is given by

$$q_{\text{comb}}(t) = \beta(t) \rho_d \Delta h_{\text{comb}} \eta. \quad (4.8)$$

Here, η is an efficiency parameter controlling the fraction of combustion heat transferred back into the member (–). Δh_{comb} represents the effective heat of combustion of timber, taken as 18.3 MJ kg^{-1} . The parameter ρ_d is the dry density of the timber (kg m^{-3}), calculated from the density including moisture ρ (kg m^{-3}) and the moisture content fraction w (–) as:

$$\rho_d = \frac{\rho}{(1 + w)}. \quad (4.9)$$

The resulting formulation provides a physically grounded way to incorporate the burning behaviour of timber into the thermal model, while remaining compatible with the implementation limitations of the SAFIR software.

The combustion heat flux in Equation 4.7 is applied as an additional contribution to the external heat flux generated by the fire curve. Initially, this flux acts on the exposed surface together with the convection and radiation boundary terms. As the charring depth increases, the location at which the flux is applied moves inward. The position of this moving boundary is determined through linear interpolation of the computed charring depth with time, and the boundary is shifted to the next mesh element once the charring front reaches the outer face of that element.

In SAFIR, this expanded thermal model is implemented using the user-defined material functionality. All temperature-dependent thermal properties are incorporated within this material definition to ensure that the numerical analysis reflects the full behaviour described above. In the remainder of this work, the term HOPKINS refers to the thermal module in which only the modified conductivity model of Hopkin is applied together with its corresponding material properties. The module HOPKINS+FLUX denotes the same implementation extended with the additional inclusion of the walking boundary heat flux.

4.2.3. Mechanical Analysis

After completion of the thermal analysis in SAFIR, the resulting transient temperature field is used as input for the subsequent mechanical analysis. This mechanical stage determines the time-dependent deformation response of the structural column under the combined effects of heating and applied loading. The sequential procedure enables the analysis to account properly for temperature-dependent material behaviour and the effects of large displacements, both of which are required to capture the evolving structural response accurately. A detailed description of the underlying mechanical formulation is provided in Subsection 2.3.2.

The mechanical problem is solved as a quasi-static analysis that includes both geometric and material non-linearity. Large-displacement effects are taken into account, ensuring that the additional destabilising moments generated by lateral deformation under axial load are represented. Material non-linearity is included through the temperature-dependent constitutive model, which reduces the stiffness and strength of the cross-section as fibres yield or soften with increasing temperature. Together, these mechanisms allow column buckling to develop naturally when the remaining stiffness is insufficient to sustain the applied axial load. The equilibrium equations are solved incrementally in time, with each increment resolved using a Newton based procedure.

The mechanical analysis is performed with one-dimensional Euler–Bernoulli beam elements, in which each node has three degrees of freedom. These include axial displacement, transverse displacement and in-plane rotation. In the present column model, the boundary conditions correspond to a pin–roller configuration. Both translational degrees of freedom at the base node are restrained, while at the top node only the transverse displacement is fixed.

The beam formulation in SAFIR employs a fibre-discretised cross-section. The temperature of each thermal finite element in the preceding heat-transfer analysis is transferred directly to a corresponding

fibre in the mechanical model. Consequently, the temperature-dependent stiffness and strength of every fibre in the cross-section evolve according to the local heating. Timber is modelled with zero thermal strain, meaning that all fibre deformation arises solely from the mechanical strains. At each stage of the analysis, the axial strain and curvature obtained from the nodal displacements are used to determine the strain in each fibre, from which the stress is computed using the non-linear, temperature-dependent constitutive model for timber. This allows both crushing in compression and brittle failure in tension to be reproduced naturally.

The overall analysis is defined parametrically with respect to the column length, cross-sectional geometry, and applied axial load, enabling consistent simulation of different column configurations.

Setup of the Mechanical Analysis


As mentioned earlier, a glued laminated timber column of strength class GL24h manufactured from spruce is assumed. In this study, two sets of ambient-temperature mechanical properties are considered.

The first set is based on the values provided in NEN-EN 14080:2013, Table 5 [39]. These properties correspond to standardised design values for GL24h and include a mean modulus of elasticity (E) of 11.5 GPa, a compressive strength (f_c) of 24 MPa, and a tensile strength (f_t) of 19.2 MPa. The strength values specified in the Eurocode represent characteristic values, corresponding to the 5-percentile of the strength distribution.

The second set of mechanical properties is derived from the probabilistic study on glued laminated soft-wood timber conducted by Schilling et al. [40]. This dataset represents material property distributions obtained from experimental characterisation and includes a mean modulus of elasticity (E) of 12.8 GPa, a mean compressive strength (f_c) of 40.4 MPa, and a mean tensile strength (f_t) of 26.4 MPa.

The mechanical behaviour of timber in SAFIR is modelled using two material modules, WOODDEC5 and WOODPRBWE, both of which account for temperature-dependent degradation of stiffness and strength. The WOODDEC5 module applies the deterministic reduction factors specified in Annex B of EN 1995-1-2:2004 [1]. These reduction factors describe the reduction of mechanical properties as described in Subsection 2.1.3 and are illustrated in Figure 2.3

The WOODPRBWE module extends this formulation by introducing probabilistic strength reduction through the application of a Weibull distribution, as proposed by Garcia-Castillo et al. [46]. In this formulation, an additional parameter, which is the Weibull quantile, governs the selection of the tensile and compressive strength reduction curves, allowing variability in material strength to be represented. The influence of this probabilistic formulation on the strength reduction factors is illustrated in Figure 4.8. For stiffness degradation, WOODPRBWE adopts the same temperature-dependent reduction factors as WOODDEC5. It is noted that, for a fixed Weibull quantile, the resulting material behaviour in a given simulation remains fully deterministic.



(a) Weibull reduction factor distribution for compression strength as a function of temperature

(b) Weibull reduction factor distribution for tensile strength as a function of temperature

Figure 4.8: Weibull distribution reduction factors for compression and tensile strength [46].

Mechanical degradation due to heating is assumed to be irreversible. This means that the reduction factors in Figure 2.3 and Figure 4.8 are applied according to the maximum temperature previously reached by each fibre, and that the mechanical properties do not recover during the cooling phase.

The mechanical analysis is carried out parametrically with respect to the column length (l), the applied vertical force (F), and the cross-sectional width (B). The loading applied to the column is defined to represent both axial force and bending effects in a consistent manner. The self-weight of the column is modelled as a distributed load acting along the central axis of the member and is calculated as the product of the initial density, the cross-sectional area, and a gravitational acceleration of 10 m s^{-2} .

In addition to the self-weight, the applied vertical force is associated with an initial eccentricity of 20 mm, which is represented by introducing a first-order bending moment. This moment is applied at both the top and bottom of the column, resulting in a constant moment distribution along the member. Consequently, the critical mid-height section experiences the prescribed initial bending moment. The axial load itself is applied at the centre of the top cross-section. A comparison was made with an alternative loading approach in which the vertical force was applied eccentrically rather than through explicit end moments. No meaningful differences in the structural response were observed. The results of this comparison are presented in Appendix G. For simplicity and clarity, the approach using initial end moments is therefore adopted in the analyses.

The load application is implemented using the SAFIR function `FLOAD`. The load is ramped up linearly from zero to its full value over the first 20 seconds by applying a factor of $t/20$ for $t \leq 20\text{s}$, and a factor of one thereafter. This gradual load application is adopted to ensure numerical stability, particularly for columns in which the applied load would already induce failure in the absence of fire exposure.

The number of fibres required in the mechanical analysis is taken directly from the mesh definition of the thermal analysis. A variable time-stepping scheme is adopted, with a minimum time step of `COMBACK` = 1×10^{-9} seconds, an initial step of 1 second, and a maximum step of 30 seconds. Results are written every 60 seconds. Convergence within each increment is controlled using `PRECISION` = 1×10^{-3} . Thermal expansion (`EPSTH`) is activated, although it has no practical effect as timber is modelled with zero thermal strain.

Mechanical Mesh Definition

The column is discretised into beam elements, each defined by three nodes and integrated along its length using three Gauss points.

To establish an appropriate discretisation size, a mesh sensitivity study was performed with the aim of achieving sufficient numerical accuracy while maintaining reasonable computational effort. The study, presented in Appendix H, showed that an element length of approximately 0.25 metres provides an adequate balance between accuracy and efficiency.

Because the mechanical analysis is performed parametrically over different column lengths, the mesh must adapt to each length while remaining close to the reference element size. Since the mesh must also fit exactly within the column length, a slightly adjusted mesh size is selected for each configuration. In addition, the number of elements is required to be odd so that both a node and a Gauss integration point are positioned exactly at the mid-height of the column. This ensures that all necessary analysis results can be extracted at this critical position.

To achieve this, an algorithm was applied that selects the most suitable mesh size for each column length. The algorithm first determines the ideal (non-integer) number of elements by dividing the column length by the target element size. The two nearest odd integers around this ideal value are then identified, and the corresponding mesh sizes are calculated. By comparing these two options, the mesh size whose deviation from the reference value is smallest is selected. In this way, the column length is divided into an odd number of elements while keeping the mesh size as close as possible to the prescribed value of approximately 0.25 metres.

Failure Determination

In the mechanical analysis, structural failure is defined as the point at which SAFIR is no longer able to reach equilibrium within a given time step. This occurs when the iterative procedure fails to satisfy the prescribed convergence tolerance, and the time step has been reduced to the minimum permitted value. Loss of equilibrium therefore marks the onset of global structural instability and is adopted as the numerical failure criterion in this study.

SAFIR models timber using the temperature-dependent constitutive law shown in Figure 4.9, which is elastic–brittle in tension and elastic–plastic in compression. The extent of compressive plastic deformation is governed by a prescribed limit value for the compressive plastic strain ε_{cu} , where a value of zero

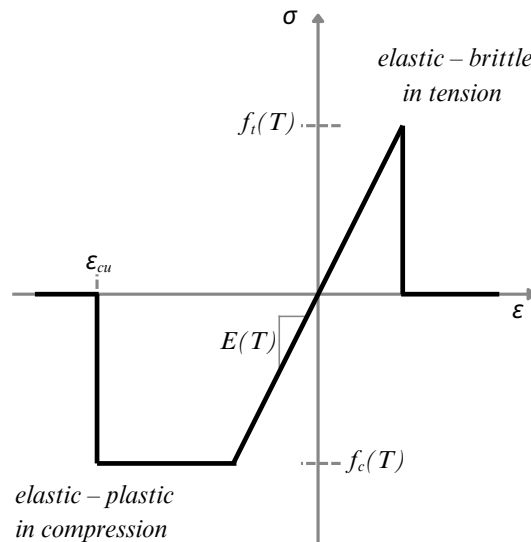


Figure 4.9: Constitutive stress–strain relationship for timber used in the mechanical analysis.

represents purely elastic–brittle behaviour and larger values permit increasing plastic deformation. The importance of the compressive plastic strain limit was investigated through a sensitivity study presented in Appendix I, in which simulated failure times were compared with four standard fire exposure tests on heated glulam columns reported by Stanke et al. [35]. Simulations were carried out using compressive plastic strain limits of zero and one, in combination with alternative strength reduction models. The results show that both the selected strength reduction curve and the chosen value of ε_{cu} influence the predicted failure times. However, for all investigated cases, the predicted failure times remained reasonable within the range of the experimental observations, with no pronounced outlier behaviour observed for large or low values of compressive plastic strain. The appropriate setting of ε_{cu} was therefore found to be problem-dependent and not uniquely defined. On this basis, both limiting cases of zero and full compressive plasticity are adopted in the study.

4.3. GPC Surrogate Modelling Framework

The GPC surrogate modelling framework is developed in a stepwise manner to ensure validity for increasing input dimensionality. The development started with defined trial functions in two dimensions and was subsequently extended to three, six, and twelve dimensions. These trial functions provide a known ground truth, allowing the performance of the surrogate modelling and active learning strategy to be systematically assessed. The results of these trial-function studies are presented in Appendices J–M. Based on these analyses, the framework is found to perform reliably up to six input dimensions, beyond which increased limitations and difficulty is experienced (see Appendix M for further insight). Consequently, the surrogate modelling framework is applied up to six dimensions in the remainder of this study.

Accordingly, two modelling cases are defined in this study. The first case focuses on structural and fire parameters and includes the structural variables l , F , and B , together with the fire characteristics t_h , T_{\max} , and r_c . The second case combines geometric and fire exposure variables with material properties and considers B , t_h , E , ρ , f_c , and e_0 .

The applicability of the surrogate framework to the physical problem is further assessed by progressively increasing the input dimensionality when applied to numerical simulation data. This stepwise expansion allows the accuracy of the surrogate predictions to be evaluated to a certain extent as additional physically relevant variables are introduced.

For the first modelling case, the analysis starts with a two-dimensional input space defined by B and t_h . The input space is then expanded to three dimensions by including the applied load F , and finally extended to a six-dimensional configuration comprising all structural and fire variables.

For the second modelling case, the framework is first applied to a two-dimensional input space defined by t_h and ρ . This is followed by a three-dimensional case including t_h , ρ , and E . The final configuration considers a six-dimensional input space defined by t_h , B , ρ , E , f_c , and e_0 .

All surrogate modelling is implemented in Python, using the GPy library for the Gaussian Process formulation. The implementation builds on the methodological framework and reference code provided by Zhao et al. [28], which was adapted to the requirements of this study. The complete code developed and used in this thesis is publicly available in the corresponding GitHub repository [38].

4.3.1. GPC Surrogate Model Setup

The Gaussian Process Classification surrogate model represents the structural response as a binary outcome corresponding to column failure or survival. Accordingly, the model is formulated using a Bernoulli likelihood with a probit link function to map the latent Gaussian process values to class probabilities. Expectation propagation is employed for inference, as it provides an efficient posterior approximation and integrates effectively with the active learning framework.

GPC Inputs

As Gaussian Process Classification is a supervised modelling approach, it requires both input variables and corresponding output labels obtained from the numerical simulations. For the surrogate modelling stage, all input variables are bounded in order to define a physically meaningful and computationally reasonable design space. The bounds are selected to reflect realistic design and structural fire scenarios.

The bounds of the structural variables are defined based on two representative building typologies, for which burnout resistance is expected to be particularly relevant. The first typology corresponds to a residential mid-rise building, while the second represents an office high-rise building. The resulting variable ranges are additionally evaluated to provide insight into which conventional fire resistance classes are covered. In particular, the bounds are confirmed to span relevant standard fire resistance ratings from approximately R60 to beyond R120, as shown in Table N.1. This evaluation is performed through inverse application of Eurocode fire resistance calculations and allows for eventual comparison to standard fire resistance classes when required.

The bounds for the fire characteristics are defined based on a review of experimental studies on timber compartment fires, explicitly including both heating and cooling phases [13–15, 32]. These studies provide guidance for realistic ranges of heating duration, peak temperature, and cooling rate.

Given that the material variables are subsequently treated as stochastic inputs in the trained surrogate model, their bounds are defined to cover 99.7% of the corresponding probability density functions. The probability density distributions reported by Schilling et al. [40] are adopted for this purpose. For the stiffness size correction factors depending on the column width are taken into account, whereas for the eccentricity corrections are applied according to the length of the column. Consequently, the resulting bounds are determined by considering the most restrictive and most favourable geometric configurations for which the 99.7% criterion remains satisfied.

For the detailed derivation and calculation of the defined bounds, reference is made to Appendix N.

An overview of all input variables and their corresponding bounds is provided in Table 4.2.

For variables that are not varied in a given surrogate modelling case, fixed reference values are adopted. The reference values for the material variables are based on the mean values from the study by Schilling et al. [40], whereas the reference structural and fire variables are derived from the configuration of experimental tests 3 and 6 reported by Gernay et al. [31]. The adopted reference values are $F = 322$ kN, $B = 0.28$ m, $l = 3.65$ m, $t_h = 15$ min, $r_c = 10.4$ °C min⁻¹, $\rho = 420$ kg m⁻³, $E = 12.8$ GPa, $f_c = 40.4$ MPa, and $e_0 = 0.02$ m. When not prescribed explicitly, the maximum temperature T_{\max} is determined by the fire curve as a function of the heating duration with a Γ -factor equal to one.

The output of each numerical simulation provides a binary response indicating column survival or failure, which is encoded as numerical class labels of 0 for no failure and 1 for failure.

Table 4.2: Values and Bounds of input variables for simulation data generation

Type	Variable	Symbol	Value/Bounds	Unit
Fire	Duration of heating phase	t_h	[1, 60]	min
	Peak temperature	T_{\max}	[600, 1200]	°C
	Cooling rate	r_c	[6, 20]	°C min ⁻¹
Structural	Column length	l	[3, 5]	m
	Column width	B	[250, 500]	mm
	Applied load	F	[300, 1800]	kN
Material	Compressive strength	f_c	[35.7, 45.8]	MPa
	Eccentricity	e_0	[13.275, 26.865]	mm
	Modulus of elasticity	E	[9.44, 17.37]	GPa
	Density	ρ_d	[360, 490]	kg m ⁻³

Normalisation

Prior to surrogate model training, the physical input variables are normalised to ensure numerically stable and efficient Gaussian Process fitting. Normalisation removes the influence of differing units and scales across input variables, allowing all dimensions to contribute comparably to the surrogate model.

Normalisation of the inputs is performed using a min–max transformation, defined as

$$x' = \frac{x - a}{b - a},$$

where x denotes the original variable value and a and b represent the lower and upper design bounds of the corresponding variable (see Table 4.2). This transformation maps each input variable to the interval $[0, 1]$, ensuring that all features are dimensionless and lie within a consistent range.

GPC Kernel Selection

The covariance structure of the Gaussian Process Classification model is defined using a radial basis function (RBF) kernel with automatic relevance determination (ARD). The RBF kernel is chosen due to its smoothness properties and its ability to represent non-linear decision boundaries in the input space. This is consistent with the physically expected behaviour of the failure boundary.

The ARD RBF kernel is defined as

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2} \sum_{i=1}^d \frac{(x_i - x'_i)^2}{\ell_i^2}\right),$$

where \mathbf{x} and \mathbf{x}' denote two input vectors, σ_f^2 is the kernel variance or kernel scaling, d is the number of input dimensions, and ℓ_i are the characteristic length scales associated with each input variable.

The use of ARD allows the surrogate model to adaptively weight the influence of individual input dimensions through their corresponding length scales.

In addition, the choice of the RBF kernel is consistent with the importance-based sampling strategy employed by the active learning framework as presented in Subsection 2.4.3.

4.3.2. Setup of the Active Learning Algorithm

Active learning is employed within the GPC framework to efficiently identify the decision boundary between failure and non-failure regions. The primary motivation for using active learning is to minimise the number of expensive simulations required to accurately locate the boundary. Rather than uniformly sampling the design space, the active learning strategy iteratively selects new simulation points that are expected to be most informative for improving the classifier, with a particular focus on regions near the decision boundary. As a result, the decision boundary can be resolved with significantly fewer simulations compared to passive sampling approaches, leading to substantial computational savings.

The active learning framework developed in this work is tailored to the considered specific engineering problem and leverages several problem-specific characteristics and assumptions. These assumptions

are not generic to all classification problems but are grounded in the physical interpretation of the underlying numerical simulations and are exploited to increase sampling efficiency.

First, simulations corresponding to the non-failure region are computationally more expensive than those in the failure region. While both failure and non-failure cases require the full thermal analysis, non-failure cases additionally require the complete mechanical response to be simulated up to the prescribed end time. In contrast, failure simulations typically terminate earlier once the failure criteria are met, leading to reduced computational costs. Consequently, reducing the number of non-failure simulations is particularly beneficial from a computational perspective.

Second, the classification boundary is deterministic, as the class labels are generated by numerical simulations without stochastic variability. Each input configuration therefore maps to a unique failure outcome, implying that the failure and non-failure regions in the design space are strictly separable and do not exhibit class overlap.

A further key assumption concerns the structure of the decision boundary in the design space. The problem is assumed to exhibit a single, clean separation between the failure and non-failure regions, without disconnected regions, enclosed class behaviour, or isolated failure islands. Geometrically, this implies that the decision boundary forms a continuous hypersurface that partitions the design space into two distinct regions. From a functional perspective, this behaviour can be interpreted as monotonicity with respect to each design variable. This implies that once all other variables are held fixed, variation of a single variable leads to at most one transition between non-failure and failure. This assumption is physically motivated by the underlying system behaviour and is consistent with the deterministic nature of the numerical simulations.

Finally, these characteristics are efficiently used in the design of the active learning algorithm. Sampling is initiated from the failure region and proceeds towards the decision boundary, rather than exploring the entire design space uniformly. A prior belief favouring the non-failure class is adopted, reflecting the assumption that designs are safe unless evidence of failure is observed. The initial simulation point is chosen such that failure is guaranteed, based on prior physical insight into the system behaviour. From this starting point, the active learning procedure iteratively refines the classifier by selectively sampling near the estimated decision boundary.

Active Learning Method

The active learning strategy applied in this study is based on the NR-SMOCU-SGD algorithm. A detailed theoretical description of this method is provided in Subsection 2.4.3. To improve computational efficiency in higher-dimensional input spaces, the original formulation is extended with a pre-selection step.

Rather than relying on uniform random Monte Carlo sampling over the entire design space for evaluating the SMOCU acquisition function, an initial Sobol sampling is performed over the domain. For each Sobol sample, the posterior failure probability predicted by the current GPC model is evaluated, and a boundary score $|p - 0.5|$ is computed, where lower values indicate that points lie closer to the current decision boundary. Subsequently, a pre-described fraction of the best-scoring points is retained and used as the candidate set for subsequent SMOCU evaluation.

This boundary-filtering step explicitly guides the SMOCU acquisition towards regions of interest and significantly improves efficiency in higher-dimensional spaces. In such settings, uniform random sampling would require an increasing number of candidate points to locate informative regions, while simultaneously increasing the approximation points of the SMOCU integral for each candidate. Since posterior probability evaluation is computationally less expensive compared to SMOCU integral approximation, the Sobol-based pre-selection allows a more dense initial exploration of the design space. The SMOCU acquisition is then evaluated only for the retained boundary-focused candidate set, from which the highest-scoring point is selected. This remaining point is then refined using the gradient-based NR-SMOCU-SGD optimisation procedure.

It should be noted that this approach is valid in the present application, as the failure and non-failure regions are expected to be separated by a single, continuous decision boundary within the considered region of the design space.

Stochastic gradient descent is used to optimise the acquisition function, with a learning rate of 0.001 and a maximum of 200 optimisation steps. The number of Sobol samples, retained candidate points, and SMOCU integral approximation samples are adjusted according to the dimensionality of the input space.

For the two-dimensional cases, 20,000 Sobol samples are generated, of which 1 % is retained, resulting in 200 candidate points. The SMOCU integral is approximated using 2,000 samples.

For the three-dimensional cases, 200,000 Sobol samples are generated with a retention fraction of 0.1 %, again yielding 200 candidate points. The SMOCU integral is approximated using 2,000 samples.

For the six-dimensional cases, 1,000,000 Sobol samples are generated with a retention fraction of 0.01 %, resulting in 1,000 candidate points. In this case, 4,000 samples are used to approximate the SMOCU integral.

The initial hyperparameters of the covariance kernel function are set to a kernel variance of 3 and length scales of 0.4 for all input dimensions. The prior mean of the latent Gaussian Process is fixed to a value of -2, corresponding to an initial failure probability of $P(y = 1) = 0.1587$ under the probit link function, see Figure 2.10 and Equation 2.36. This reflects the initial setting of favouring the non-failure class.

Hyperparameter Optimisation

During the active learning procedure, the kernel hyperparameters of the GPC model are optimised. For the ARD RBF kernel considered in this work, these hyperparameters specifically include the kernel variance and the characteristic length scales associated with each input dimension. Hyperparameter optimisation, however, influences both the expressiveness of the fitted model even as the behaviour of the acquisition of the active learning algorithm.

At the same time, hyperparameter optimisation is inherently problem dependent and involves a degree of arbitrariness. This is particularly the case in sequential active learning settings, where the training set evolves continuously as new samples are selected adaptively based on the current state of the model.

In low-dimensional problems, it is often possible to assess visually whether hyperparameter optimisation is appropriate. In higher-dimensional design spaces, however, such visual diagnostics are no longer feasible, and quantitative criteria are required to determine when optimisation should be performed. Based on the problem-specific insight applied in this work, hyperparameter optimisation is intentionally avoided at an early stage and is only enabled once the decision boundary is sufficiently enclosed. First of all, this is because optimising the kernel variance too early causes the prior belief favouring one class to diminish, as dictated by the probit link formulation in Equation 2.36. In addition, premature optimisation of the kernel length scales is also undesirable, since when the decision boundary is not yet enclosed in each dimension, length scales may increase excessively, again weakening the approach of an intended bias towards one class.

At the same time, delaying hyperparameter optimisation for too long is also undesirable, as improved hyperparameter estimates lead to a more accurate representation of the decision boundary and enable the selection of more informative query points. Therefore, to determine the appropriate moment for hyperparameter optimisation and to assess whether the decision boundary is sufficiently enclosed, a global boundary movement metric is introduced. Prior to the start of the active learning algorithm, a fixed Sobol monitoring set is generated over the design space. At each iteration, the number of monitoring points classified as failure is evaluated and compared with the previous iteration. This difference is averaged by the total number of monitoring points, yielding a scalar measure of the change in the global prediction field.

Unlike relying solely on the behaviour of the acquisition function, which becomes increasingly spiky in higher-dimensional settings, this metric remains more stable and provides a more reliable indicator of decision boundary stabilisation.

The decision boundary is considered to be sufficiently enclosed once the boundary movement metric has been observed to fall twenty times below 25 % of the maximum recorded movement. From this point onward, hyperparameter optimisation is enabled. When optimisation is performed a BFGS optimiser with six random restarts and a maximum of 500 iterations per restart is applied. The latent mean function is kept fixed at its initially specified prior during optimisation. The global boundary movement metric is evaluated using a fixed Sobol monitoring set of which the size is based on the dimensionality of the design space. Specifically, 20,000 monitoring points are used for two-dimensional problems,

200,000 points for three-dimensional problems, and 1,000,000 points for six-dimensional problems.

Special care is taken with respect to kernel variance optimisation. While optimising the kernel variance substantially improves the overall fit of the GPC model, it is also able to adversely affect the behaviour of the active learning algorithm. In particular, increasing the kernel variance rescales the latent function values, which in turn alters the magnitude of changes in the predictive field. Since the SMOCU acquisition function explicitly quantifies expected changes in prediction, this rescaling is able to change the sampling behaviour. The effect is further amplified in the presence of strongly anisotropic kernel length-scales, as in this way correlations are able to extend over large regions in specific directions. In such cases, query points that are locally confident but globally influential may be favoured over samples located near the decision boundary, particularly once the boundary has already been approximated with reasonable accuracy.

Although the application of a boundary filter already reduces the risk of such misbehaviour, kernel variance optimisation remains a sensitive component of the active learning framework. For this reason, multiple optimisation strategies are considered with respect to the kernel variance, which are discussed in the following.

First, as a default approach, once hyperparameter optimisation is enabled, both the kernel variance and the kernel length scales are optimised every ten active learning iterations. This strategy typically yields the best overall model fit, as it allows all hyperparameters to adapt to the evolving training set.

If this approach leads to undesirable acquisition behaviour, a more conservative strategy is adopted. In this case, hyperparameter optimisation for both the kernel variance and the length scales is still performed every ten iterations, but updates to the kernel variance are only accepted if the relative improvement in the log marginal likelihood exceeds a predefined threshold. Kernel length scales are optimised regardless, as they primarily govern the geometric representation of the decision boundary.

If misbehaviour persists, kernel variance optimisation is disabled entirely, and only the kernel length scales are optimised every ten iterations. While this approach ensures stable and informative sampling behaviour, it typically results in an eventual less accurate final model fit. Consequently, convergence of the active learning process must be assessed using alternative criteria, as discussed in the following section. Once sufficient convergence is established, the resulting dataset in this case can be reused to fit a new GPC model to the data, in which optimisation of both the kernel variance and the length scales are enabled to improve the overall model fit.

Convergence Determination

As with hyperparameter optimisation, determining convergence of the active learning algorithm is inherently problem dependent and involves a degree of arbitrariness. In contrast to classical supervised learning, convergence is not determined using a validation dataset, as the objective of the active learning procedure is to accurately identify a deterministic decision boundary, which represents a local characteristic in the global design space whose position is not known in advance. Applying a validation-based approach would therefore require a very large initial simulation dataset, which is computationally infeasible.

Based on insight obtained from the trial functions described in Appendices J-M, where the true misclassification error can be evaluated, a set of practical convergence metrics is defined. These include the acquisition curve, the boundary movement metric, a boundary spread metric, and a misclassification metric, all of which are monitored throughout the active learning process.

The boundary movement metric, introduced previously, quantifies changes in the global prediction field by tracking variations in the number of monitoring points classified as failure between iterations. This metric provides a robust indication of whether the decision boundary continues to evolve or has stabilised. In addition, a boundary spread metric is defined to assess the sharpness of the fitted decision boundary. This metric is evaluated using the same fixed Sobol monitoring set as the boundary movement metric. At each iteration, all monitoring points with predicted failure probabilities in the interval $p(y = 1 | x_i) \in [0.3, 0.7]$ are identified, and the fraction of these points relative to the total number of monitoring points is computed. A decreasing boundary spread indicates progressive sharpening of the decision boundary and a global model fit improvement.

A further convergence indicator is the misclassification metric, which measures the fraction of present

simulated data points that are misclassified by the current GPC model. Since the underlying problem is deterministic, the true misclassification rate is expected to be zero. In practice, a small nonzero threshold is tolerated to account for model approximation error and the increased sampling density near the decision boundary. Elevated misclassification values indicate that the decision boundary is not yet sufficiently resolved, reflecting that the current hyperparameter configuration is not yet able to adequately represent its structure.

Under the default active learning strategy, where both kernel variance and length scales are optimised once the boundary is enclosed, convergence is determined using a combination of the boundary spread and misclassification metrics. Convergence is declared when the boundary spread metric falls below 0.01 and the misclassification metric has been remained below 0.05 for eleven consecutive iterations. However, as discussed in the previous subsection on hyperparameter optimisation, scenarios exist in which kernel variance optimisation is restricted or disabled to prevent undesirable sampling behaviour. In these cases, the boundary spread metric can no longer be relied upon, as it typically decreases only when kernel variance is optimised. Consequently, an alternative convergence strategy is applied. When kernel variance optimisation is limited, convergence is assessed primarily based on stabilisation of the acquisition curve and the boundary movement metric. Once these indicators exhibit consistently flat behaviour, the decision boundary is considered converged. This assessment however remains based on qualitative insight into the observed trends. In addition, in such cases the misclassification metric can be used as a complementary indicator. Because restricting kernel variance optimisation limits the model's fitting ability the misclassification metric may increase once convergence is approached. This increase reflects the growing concentration of samples near the decision boundary and occurs because the model under the chosen hyperparameter optimisation strategy is no longer capable of accurately fitting these increasingly fine-scale data.

After convergence is declared based on boundary stabilisation, the collected dataset is reused to fit a final GPC model to the data, with both optimisation of kernel variance and length scales to improve the overall model fit.

4.3.3. GPC Surrogate Model Evaluation

Once the GPC surrogate has been trained using the active learning procedure and convergence has been determined, an evaluation stage is carried out. In this stage, the trained surrogate is able to be used as a computationally efficient representation of the finite element model, enabling extensive post-processing without additional numerical simulations.

For the second trained six-dimensional model, a central aspect of the evaluation concerns the propagation of material uncertainty. Probabilistic descriptions of the GL24h material properties, including density, stiffness, and compressive strength, are adopted from Schilling et al. [40], whose distribution parameters are presented in Table 4.3. The corresponding physical means are 420 kg m^{-3} for ρ , 12.8 GPa for E , and 40.4 MPa for f_c . In addition, the reference provides a geometry-dependent correction factor for the stiffness, accounting for member size effects through the member depth h . For the eccentricity, Blaß [47] is used, in which a column-length-dependent normal distribution is provided.

Table 4.3: Probabilistic descriptions of GL24h material properties adopted from Schilling et al. [40] and Blaß [47].

Variable	Distribution	Distribution parameters
Density (ρ)	Lognormal	$\lambda = 6.04, \zeta = 0.0521$
Modulus of elasticity (E)	Lognormal	$\lambda = 2.55, \zeta = 0.0513$
Compressive strength (f_c)	Lognormal	$\lambda = 3.70, \zeta = 0.0416$
Eccentricity (e_0)	Normal	$\mu = 1.4 \cdot 10^{-5} \times l$ $\sigma = 4.53 \cdot 10^{-4} \times l$

The stiffness correction factor is applied as

$$k_{E,h,\zeta} = \left(\frac{0.60}{h} \right)^{0.78},$$

which is to be multiplied with the corresponding material distribution factor ζ .

Based on the specified column geometry, the correction factor and stochastic material distributions are created.

Specifically for the eccentricity, a global reference mean of $e_0 = 0.02$ m is adopted. A geometry-dependent mean eccentricity $\mu(l)$ and standard deviation $\sigma(l)$ are obtained from the length-dependent distribution reported in Table 4.3. For each Monte Carlo realisation, the geometry-dependent mean $\mu(l)$ is added or subtracted from the reference mean e_0 , with both options selected with a probability of 0.5. The resulting mean is then used together with the standard deviation $\sigma(l)$ to define a normal distribution, from which a single eccentricity value is sampled for that realisation.

Finally, a Monte Carlo evaluation is performed on a structured grid of the deterministic variables B and t_h , comprising 150×150 grid points. For each grid point, 200 Monte Carlo realisations are generated by sampling the stochastic material variables from their respective distributions. Each realisation is subsequently evaluated using the trained GPC surrogate model to determine the failure state.

In the absence of sufficient information to define a complete dependence structure between the material parameters, all material variables are sampled from their marginal distributions under an independence assumption.

5

Results

This chapter presents the results obtained in this study and is structured in two main parts, corresponding to the two core components of the thesis.

First, the results of the numerical modelling are presented in Section 5.1. This section examines the ability of the finite element simulations to reproduce the fire-induced response of the experimental investigations from Gernay et al. [31] and Renard et al. [32] on glulam timber columns. The results are first presented for standard ISO 834 fire exposure, followed by those obtained under standardised parametric fire scenarios. Finally, the response of the structures subjected to natural fire exposures is presented.

Subsequently, Section 5.2 presents the results of the surrogate modelling approach. This section evaluates the capability of the GPC framework to approximate the failure domain and to efficiently represent the outcomes of the numerical simulations. The results are first presented for a six-dimensional model comprising the structural design and fire variables. Second the six-dimensional model including material variables is shown.

5.1. Numerical Modelling Results

This section presents the results of the numerical modelling study. The simulations aim to reproduce the experimental tests described in Section 3.1, for which overviews are provided in Figure 3.3 and Figure 3.4. The results are organised according to the type of fire exposure considered. First, outcomes corresponding to the standard ISO 834 fire tests are presented. This is followed by the results obtained under standardised parametric fire exposures. Finally, the structural response under natural fire scenarios are presented.

Throughout this section, attention is given to several modelling aspects in addition to the variables directly associated with the experimental test configurations. In particular, the influence of material parameters is examined, with specific emphasis on stiffness and strength properties. Two material property configurations are considered. In the first configuration, material properties are defined in accordance with EN 14080 [39]. In the second configuration, mean material property values reported by Schilling et al. [40] are adopted.

Furthermore, detailed insight is provided into the evolution of the plastic compression strain, as illustrated in Figure 4.9. Differences resulting from the application of alternative temperature-dependent strength reduction models are also investigated. These include the deterministic WOODDEC5 model in accordance with Eurocode provisions (Figure 2.3a) and the probabilistic WOODPRBWE model employing a Weibull quantile of 0.5 (Figure 4.8).

For the standard and parametric fire scenarios, the results are primarily evaluated in terms of failure time. In contrast, for the natural fire scenarios, both structural response and measured thermal data from thermocouples are presented.

5.1.1. Standard ISO 834 Fire Exposure

This subsection presents the numerical modelling results for the standard ISO 834 fire exposure, based on the experimental furnace tests conducted by Gernay et al. [31]. In particular, Tests 1, 2, and 5 are

considered. Test 1 features a hinged–fixed boundary condition, whereas Tests 2 and 5 are characterised by hinged–hinged end restraints.

All columns share an identical geometric configuration. The column length is 3.65 m and the square cross-sectional width is 0.28 m. In addition, a timber density of 420 kg m^{-3} is adopted, an initial load eccentricity of 0.02 m is applied, and the axial load is fixed at 322 kN.

Test 1: Hinged–Fixed Configuration

The numerical results for Test 1 are summarised in Table 5.1. The experimentally observed failure time for this test is 78 minutes. Sensitivity in the predicted failure times is observed across the simulations, reflecting variations in material strength level, the inclusion of plastic compression strain, and the selected strength reduction model. Overall, the predicted failure times are of the same order of magnitude as the experimental result, with some degree of doubt associated with Simulation 5, which exhibits the largest deviation.

Among the evaluated cases, Simulations 2 and 8 provide the closest agreement with the experimentally observed failure time. Both simulations adopt the mean material property values reported by Schilling et al. [40]. Simulation 2 includes plastic compression strain and applies the deterministic Eurocode-based WOODDEC5 strength reduction curve. In contrast, Simulation 8 neglects plastic compression strain and employs the probabilistic WOODPRBWE strength reduction model using a Weibull quantile of 0.5.

Table 5.1: Numerical simulation results for Test 1 of Gernay et al. [31].
Thermal module: WOODDEC5
Failure = 78 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (–)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODDEC5	65.42
2	12.8	40.4	26.4	1.0	WOODDEC5	78.44
3	11.5	24.0	19.2	1.0	WOODPRBWE	72.76
4	12.8	40.4	26.4	1.0	WOODPRBWE	85.65
5	11.5	24.0	19.2	0.0	WOODDEC5	50.67
6	12.8	40.4	26.4	0.0	WOODDEC5	66.96
7	11.5	24.0	19.2	0.0	WOODPRBWE	67.30
8	12.8	40.4	26.4	0.0	WOODPRBWE	82.77

Tests 2 & 5: Hinged–Hinged Configuration

The numerical results corresponding to Tests 2 and 5, are presented in Table 5.2. The experimentally observed failure times are 55 minutes for Test 2 and 58 minutes for Test 5. Compared to the hinged–fixed configuration, the hinged–hinged boundary conditions result in systematically shorter failure times, in accordance with the experimental results.

Similar trends are observed with respect to sensitivity in material strength, plastic strain inclusion, and the selected strength reduction model. Simulation 5 shows the greatest deviation, whereas all the other simulations are considered within the same order of magnitude as the experimental failure times.

Table 5.2: Numerical simulation results for Tests 2 & 5 of Gernay et al. [31].
Thermal module: WOODDEC5
Test 2: **Failure = 55 min**
Test 5: **Failure = 58 min**

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (–)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODDEC5	46.36
2	12.8	40.4	26.4	1.0	WOODDEC5	60.60
3	11.5	24.0	19.2	1.0	WOODPRBWE	54.63
4	12.8	40.4	26.4	1.0	WOODPRBWE	64.86
5	11.5	24.0	19.2	0.0	WOODDEC5	35.32
6	12.8	40.4	26.4	0.0	WOODDEC5	52.26
7	11.5	24.0	19.2	0.0	WOODPRBWE	51.05
8	12.8	40.4	26.4	0.0	WOODPRBWE	63.25

Key Results

Across the ISO 834 furnace tests, the numerical model predicts failure in all configurations and reproduces the experimentally observed failure times within a comparable range. For Test 1 (hinged–fixed, failure at 78 min), the closest agreement is obtained in Simulations 2 (78.44 min) and 8 (82.77 min), while the predicted failure times across all configurations range from 50.67 min to 85.65 min. For Tests 2 and 5 (hinged–hinged, failures at 55 min and 58 min), the predicted failure times are systematically shorter than for the hinged–fixed case, consistent with the experimental trend, with results ranging from 35.32 min to 64.86 min. Across both boundary condition cases, variations in strength level, inclusion of plastic compression strain, and the selected strength reduction curve lead to noticeable shifts in predicted failure time, while the simulation failure times range covers the observed experimental failure times.

5.1.2. Standardised Parametric Fire Exposure

This subsection presents the numerical modelling results for the standardised parametric fire exposures, based on the experimental furnace tests reported by [31]. Specifically, Tests 3, 4, 6, and 7 are considered. Tests 3 and 6 are subjected to a heating duration of 15 minutes, whereas Tests 4 and 7 feature a shorter heating phase of 10 minutes.

Both fire scenarios are characterised by a Γ -factor equal to one and a linear cooling rate of $10.4\text{ }^{\circ}\text{C min}^{-1}$. All columns share an identical geometric configuration. The column length is 3.65 m with a square cross-sectional width of 0.28 m. In addition a timber density of 420 kg m^{-3} , an initial load eccentricity of 0.02 m, and an axial load of 322 kN are applied. Numerical simulations are performed with the WOODDEC5 thermal module.

Test 3 & 6: 15 Minutes Heating

The numerical modelling results for Tests 3 and 6 are presented in Table 5.3. The experimentally observed failure times are 98 minutes for Test 3 and 153 minutes for Test 6.

The simulation results exhibit a limited degree of consistency. Only Simulations 1 and 6 yield failure times that are of the same order of magnitude as the experimental observations, although even for these cases the agreement may already be questioned. In contrast, multiple simulations predict no failure. This outcome is not consistent with the experimental observations, as both tests exhibited failure behaviour.

Overall, the results indicate a strong sensitivity of the predicted response to the modelling assumptions for the standardised parametric fire exposure with a 15 minutes heating phase, and highlight difficulties in reliably reproducing the experimentally observed failure times for these scenarios.

Table 5.3: Numerical simulation results for Tests 3 & 6 of Gernay et al. [31].

Thermal module: WOODDEC5

Test 3: **Failure = 98 min**

Test 6: **Failure = 153 min**

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (–)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODDEC5	96.72
2	12.8	40.4	26.4	1.0	WOODDEC5	False
3	11.5	24.0	19.2	1.0	WOODPRBWE	False
4	12.8	40.4	26.4	1.0	WOODPRBWE	False
5	11.5	24.0	19.2	0.0	WOODDEC5	43.63
6	12.8	40.4	26.4	0.0	WOODDEC5	172.60
7	11.5	24.0	19.2	0.0	WOODPRBWE	False
8	12.8	40.4	26.4	0.0	WOODPRBWE	False

Test 4 & 7: 10 Minutes Heating

The numerical modelling results for Tests 4 and 7 are presented in Table 5.4. For both experimental tests, no structural failure was observed during the fire exposure or the subsequent cooling phase.

Overall, the numerical simulation results are largely comparable with the experimental observations, as the majority of configurations predict no failure. This outcome is observed across different material

property sets, strength reduction models, and assumptions regarding plastic compression strain. One exception is Simulation 5, which predicts failure despite the absence of experimental failure for both tests. This result deviates from the general trend of the numerical outcomes and suggests a higher sensitivity of this configuration under short heating durations.

Table 5.4: Numerical simulation results for Tests 4 & 7 of Gernay et al. [31].

Thermal module: WOODEC5

Test 4: **Failure = False**

Test 7: **Failure = False**

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	False
2	12.8	40.4	26.4	1.0	WOODEC5	False
3	11.5	24.0	19.2	1.0	WOODPRBWE	False
4	12.8	40.4	26.4	1.0	WOODPRBWE	False
5	11.5	24.0	19.2	0.0	WOODEC5	56.61
6	12.8	40.4	26.4	0.0	WOODEC5	False
7	11.5	24.0	19.2	0.0	WOODPRBWE	False
8	12.8	40.4	26.4	0.0	WOODPRBWE	False

Key Results

For the 15 minute heating exposure (Tests 3 and 6; failures at 98 min and 153 min), the numerical results show a mixed outcome across the investigated configurations, including both predicted failure and no-failure cases. Predictions range from 43.63 min to multiple no failure cases, while both experimental tests resulted in failure. For the 10 minute heating exposure (Tests 4 and 7; no failure observed), the numerical simulations predominantly predict no failure across material configurations and strength reduction models, with one configuration (Simulation 5) predicting failure at 56.61 min. Overall, the parametric-fire results exhibit a relatively large spread in predicted outcomes compared to the experimental results.

5.1.3. Natural Fire Exposure

This subsection presents the numerical modelling results for natural fire exposures, based on the experimental programme conducted by Renard et al. [32]. In these experiments, natural fires were generated using wood cribs.

The analysis focuses on Experiments 9, 10, 11, 12, 15, and 17, in which Test 10 served as the reference test within the experimental programme. Tests 9, 10, 11, and 12 showed failure during the cooling phase, whereas Tests 15 and 17 failed during the heating phase. Experiments 14, 18, 19, 20, and 21 are excluded from the analysis due to differences in testing, fire curve, and observed failure behaviour.

Several thermal modelling approaches are employed. The WOODEC5 module corresponds to temperature-dependent thermal properties as prescribed by the Eurocode [1]. The HOPKINS module adopts the conductivity model proposed by Hopkin et al. [44], as defined in Equation 4.5. The HOPKINS+FLUX module extends this approach by combining the conductivity model with a moving additional heat flux, as described in Subsection 4.2.2. The combustion efficiency of timber is assumed to be 3.5%, consistent with numerical findings of Šulc, Šmilauer, and Wald [48].

To obtain the Γ -factor required for the conductivity model of Hopkin et al. [44], the natural fire curves are approximated using parametrically defined fire curves, as described in Subsection 4.2.1. The resulting Γ -factor is subsequently applied within the conduction model.

TEST 10

Test 10 serves as the reference experiment within the natural fire test programme. In this test, failure of the column occurred shortly within the cooling phase of the fire, with an experimentally observed failure time of 47 minutes.

The natural fire was generated using timber wood cribs composed of sticks with a cross-sectional size of $90 \times 90 \text{ mm}^2$. The fire scenario was characterised by an opening factor of $0.065 \text{ m}^{1/2}$ and a fire load density of 780 MJ m^{-2} .

For the thermal analysis, the WOODDEC5 and HOPKINS thermal modules are first applied. Subsequently, the extended HOPKINS+FLUX thermal module is considered. For the WOODDEC5 and HOPKINS+FLUX thermal models, corresponding mechanical analyses are performed in which the failure time are reported.

The results of the thermal simulations for the WOODDEC5 and HOPKINS thermal modules are compared against the experimental thermocouple measurements in Figure 5.1. The numerical mechanical modelling results corresponding to the WOODDEC5 thermal analysis are presented in Table 5.5.

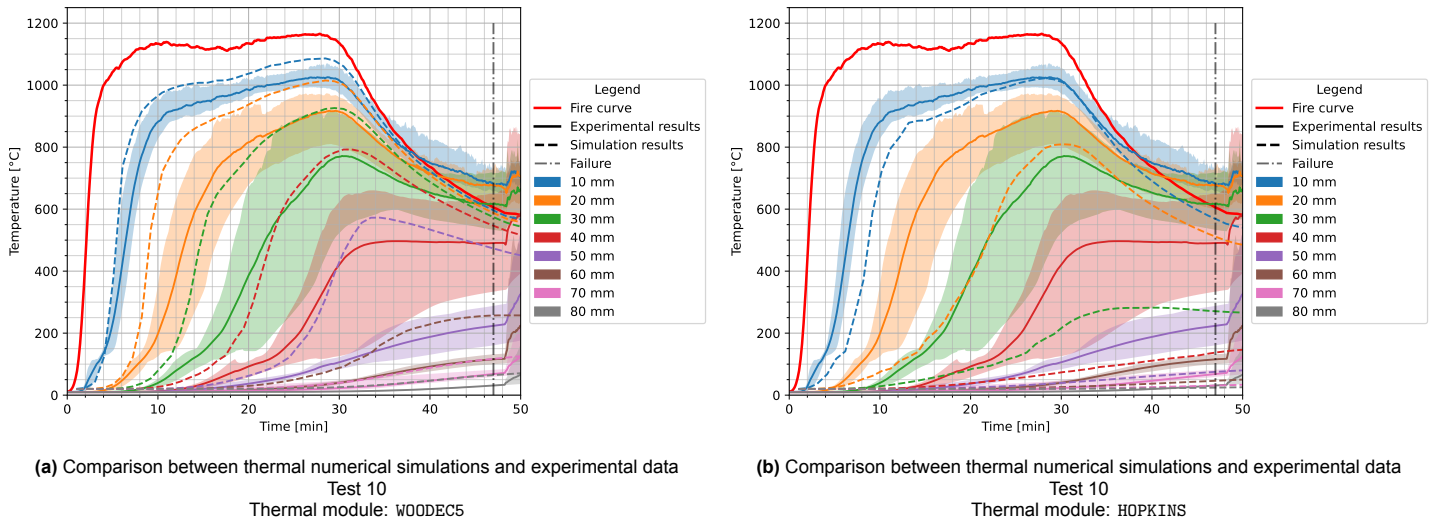


Figure 5.1: Comparison of thermal numerical simulations against experimental data for Test 10 by Renard et al. [32]. Thermal modules: WOODDEC5 and HOPKINS

Examining Figure 5.1a it is observed that that, for the WOODDEC5 thermal module, the temperature wave propagates more rapidly through the cross-section in the numerical simulations than in the experimental measurements. The deviations are particularly large at depths of 40 mm and 50 mm.

In contrast, the HOPKINS thermal module as illustrated in Figure 5.1b shows a markedly delayed thermal wave compared to the WOODDEC5 results. This delay increases towards the interior of the cross-section.

Consistent with the thermal observations, an assessment of the mechanical numerical modelling results obtained using the WOODDEC5 thermal module in Table 5.5 shows that the predicted failure times consistently underpredict the experimentally observed failure time of the column. Even the strongest numerical configurations remain relatively far from the experimental failure time.

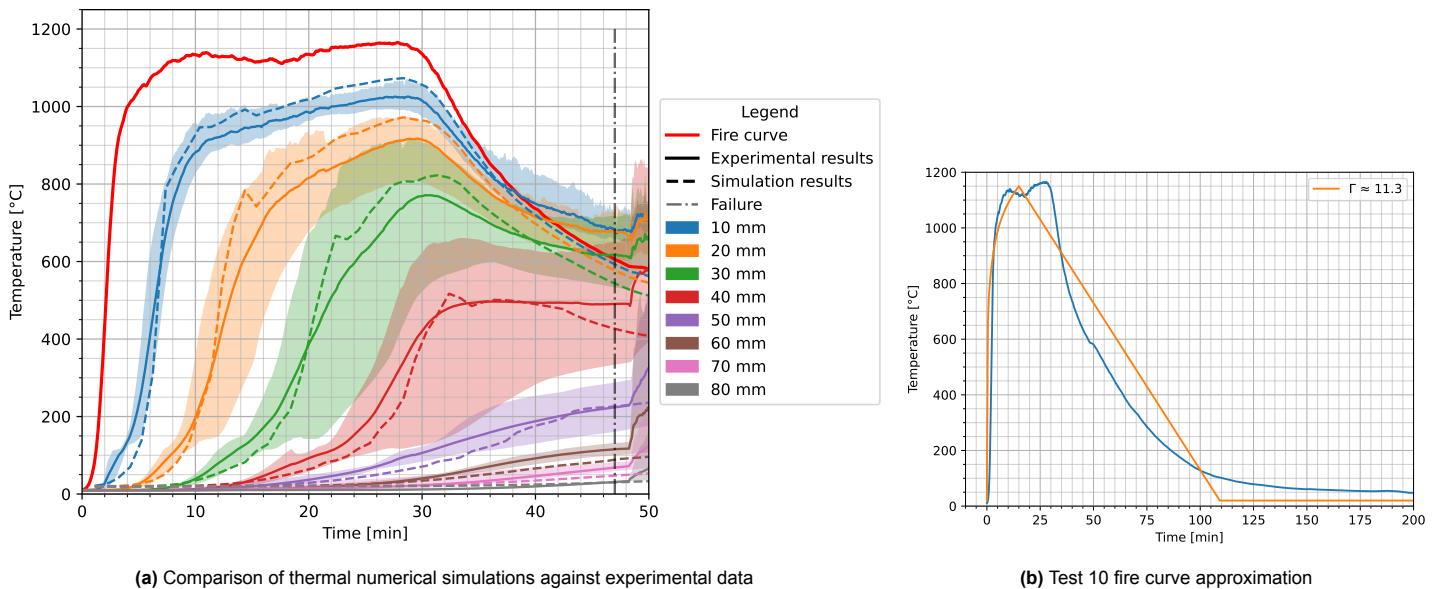
Table 5.5: Numerical mechanical simulation results for Test 10 of Renard et al. [32].

Thermal module: WOODDEC5
Test 10: Failure = 47 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODDEC5	22.12
2	12.8	40.4	26.4	1.0	WOODDEC5	27.38
3	11.5	24.0	19.2	1.0	WOODPRBWE	23.26
4	12.8	40.4	26.4	1.0	WOODPRBWE	28.29
5	11.5	24.0	19.2	0.0	WOODDEC5	18.29
6	12.8	40.4	26.4	0.0	WOODDEC5	24.80
7	11.5	24.0	19.2	0.0	WOODPRBWE	21.53
8	12.8	40.4	26.4	0.0	WOODPRBWE	27.52

Subsequently, the thermal simulation results obtained using the HOPKINS+FLUX thermal module are compared with the experimental data in Figure 5.2. Examination of these results shows improved agreement between the numerical simulations and the experimental measurements. An approximated Γ -factor of 11.30 provides good correspondence with the experimental data. A small spike is observed

in the numerical thermocouple results, which reflects the discrete advancement of the moving additional heat flux boundary within the cross-section.



(a) Comparison of thermal numerical simulations against experimental data

(b) Test 10 fire curve approximation

Figure 5.2: Comparison of thermal numerical simulation against experimental data for Test 10 of Renard et al. [32].
Thermal module: HOPKINS+FLUX

Comparison of the mechanical numerical modelling results corresponding to the HOPKINS+FLUX thermal module with the experimental failure time shows improved agreement relative to the other thermal modelling approaches. In particular, configurations employing mean material properties in combination with strength reduction curves based on a Weibull quantile of 0.5 provide the closest correspondence with the experimental failure time.

Table 5.6: Numerical mechanical simulation results for Test 10 of Renard et al. [32].
Thermal module: HOPKINS+FLUX
Test 10: Failure = 47 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	31.82
2	12.8	40.4	26.4	1.0	WOODEC5	42.29
3	11.5	24.0	19.2	1.0	WOODPRBWE	33.95
4	12.8	40.4	26.4	1.0	WOODPRBWE	46.81
5	11.5	24.0	19.2	0.0	WOODEC5	26.21
6	12.8	40.4	26.4	0.0	WOODEC5	36.28
7	11.5	24.0	19.2	0.0	WOODPRBWE	30.99
8	12.8	40.4	26.4	0.0	WOODPRBWE	44.19

TEST 11

Column failure for test 11 occurred shortly after the heating phase of the fire, with a experimentally observed failure time of 35 minutes. The natural fire was generated using timber wood cribs composed of sticks with a cross-sectional size of $90 \times 90 \text{ mm}^2$. The fire scenario is characterised by an opening factor of $0.065 \text{ m}^{1/2}$ and a fire load density of 950 MJ m^{-2} .

Thermal modelling results obtained using the HOPKINS+FLUX thermal module are presented in Figure 5.3. The natural fire curve is approximated using a Γ -factor of 16.28, which results in good agreement with the experimentally observed temperature measurements. Again as previously observed, a small spike is visible in the numerical thermocouple results, corresponding to the discrete advancement of the moving boundary heat flux within the cross-section.

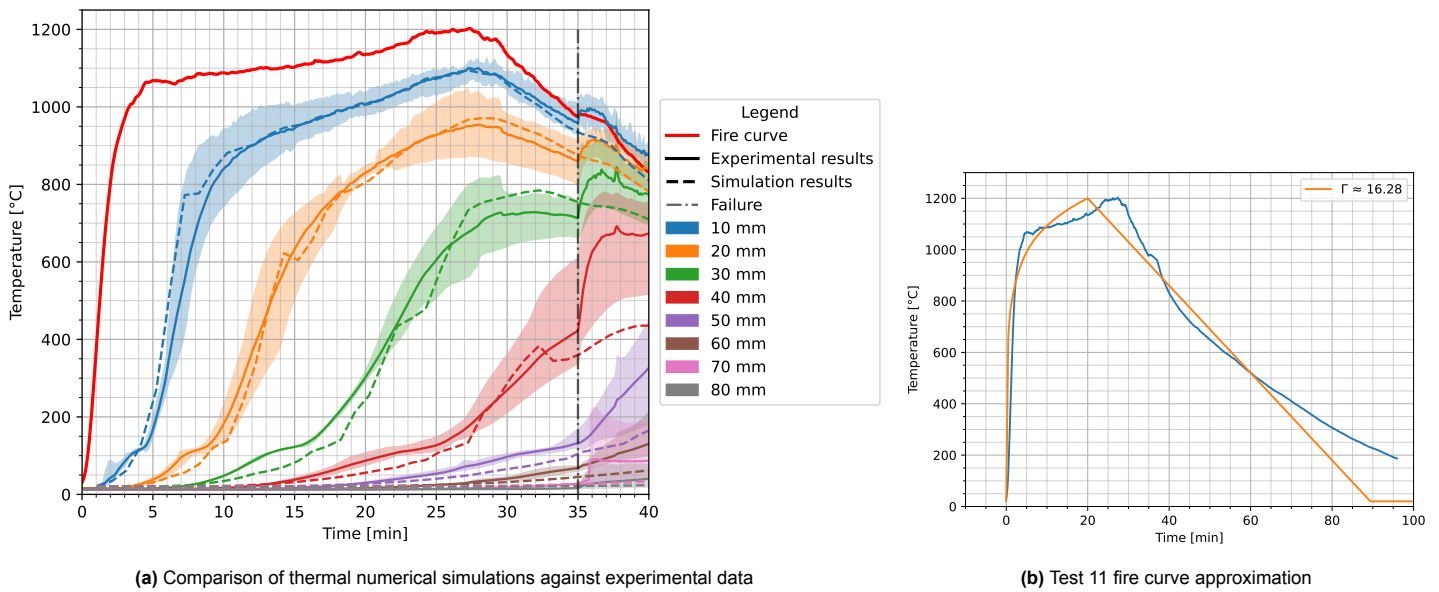


Figure 5.3: Comparison of thermal numerical simulation against experimental data for Test 11 of Renard et al. [32].
Thermal module: HOPKINS+FLUX

The mechanical numerical modelling results for Test 11 obtained using the HOPKINS+FLUX thermal module are summarised in Table 5.7. Overall, the predicted failure times show good agreement with the experimentally observed failure time of 35 minutes. The numerical results consistently fall within the same range as the experimental outcome across the different material property configurations and strength reduction models.

Table 5.7: Numerical mechanical simulation results for Test 11 of Renard et al. [32].
Thermal module: HOPKINS+FLUX
Test 11: Failure = 35 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	33.29
2	12.8	40.4	26.4	1.0	WOODEC5	43.65
3	11.5	24.0	19.2	1.0	WOODPRBWE	35.31
4	12.8	40.4	26.4	1.0	WOODPRBWE	48.00
5	11.5	24.0	19.2	0.0	WOODEC5	27.56
6	12.8	40.4	26.4	0.0	WOODEC5	37.71
7	11.5	24.0	19.2	0.0	WOODPRBWE	32.80
8	12.8	40.4	26.4	0.0	WOODPRBWE	46.05


TEST 15

For Test 15, column failure occurred during the heating phase of the fire, with an experimentally observed failure time of 37 minutes. The natural fire was generated using timber wood cribs consisting of sticks with a cross-sectional size of $90 \times 90 \text{ mm}^2$. The fire scenario is characterised by a reduced opening factor of $0.032 \text{ m}^{1/2}$ and a fire load density of 780 MJ m^{-2} .

Thermal simulation results obtained with the HOPKINS+FLUX thermal module are shown in Figure 5.4. The natural fire curve is approximated using a Γ -factor of 9.30, which yields a satisfactory correspondence with the experimentally measured temperature data. Similar to the previously analysed tests, a small spike is observed in the numerical thermocouple response, reflecting the discrete progression of the moving boundary heat flux through the cross-section.

Furthermore, the numerical simulations exhibit a clear increase in temperature as the moving boundary heat flux reaches the corresponding cross-sectional depth. At this stage, the numerical predictions intersect the experimental measurements, with a slight overestimation of temperature at the time of

failure. Nevertheless, the numerical results remain within the spread of the experimentally measured thermocouple data.



(a) Comparison of thermal numerical simulations against experimental data

(b) Test 15 fire curve approximation

Figure 5.4: Comparison of thermal numerical simulation against experimental data for Test 15 of Renard et al. [32].
Thermal module: HOPKINS+FLUX

The mechanical numerical modelling results for Test 15 obtained using the HOPKINS+FLUX thermal module are presented in Table 5.8. The predicted failure times show close agreement with the experimental result, with most simulations yielding values within a narrow range around the observed failure time. This level of agreement is consistently observed across the different material property sets, assumptions regarding plastic compression strain, and selected strength reduction models.

Table 5.8: Numerical mechanical simulation results for Test 15 of Renard et al. [32].
Thermal module: HOPKINS+FLUX
Test 15: Failure = 37 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	35.88
2	12.8	40.4	26.4	1.0	WOODEC5	42.60
3	11.5	24.0	19.2	1.0	WOODPRBWE	37.34
4	12.8	40.4	26.4	1.0	WOODPRBWE	43.72
5	11.5	24.0	19.2	0.0	WOODEC5	30.63
6	12.8	40.4	26.4	0.0	WOODEC5	39.46
7	11.5	24.0	19.2	0.0	WOODPRBWE	35.12
8	12.8	40.4	26.4	0.0	WOODPRBWE	42.75

TEST 17

Test 17 showed an experimental failure time of 45 minutes, corresponding to a column failure within the heating phase of the fire. The natural fire was generated using timber wood cribs composed of sticks with a cross-sectional size of $120 \times 120 \text{ mm}^2$. In addition is the fire scenario characterised by an opening factor of $0.065 \text{ m}^{1/2}$ and a fire load density of 780 MJ m^{-2} .

Thermal simulation results obtained using the HOPKINS+FLUX thermal module are presented in Figure 5.5. The natural fire curve is approximated using a Γ -factor of 1.28. The numerical results show an increase in temperature in regions where the additional boundary heat flux is applied, resulting in a slight overprediction of temperatures when compared with the experimental measurements.




Figure 5.5: Comparison of thermal numerical simulation against experimental data for Test 17 of Renard et al. [32].
Thermal module: HOPKINS+FLUX

The mechanical numerical modelling results for Test 17 obtained using the HOPKINS+FLUX thermal module are summarised in Table 5.9. The numerical predictions show proper agreement with the observed experimental failure time, as most of the simulations yield a failure time within a comparable range.

Table 5.9: Numerical mechanical simulation results for Test 17 of Renard et al. [32].
Thermal module: HOPKINS+FLUX
Test 17: Failure = 45 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	41.15
2	12.8	40.4	26.4	1.0	WOODEC5	50.84
3	11.5	24.0	19.2	1.0	WOODPRBWE	44.04
4	12.8	40.4	26.4	1.0	WOODPRBWE	53.15
5	11.5	24.0	19.2	0.0	WOODEC5	32.80
6	12.8	40.4	26.4	0.0	WOODEC5	44.72
7	11.5	24.0	19.2	0.0	WOODPRBWE	41.42
8	12.8	40.4	26.4	0.0	WOODPRBWE	51.81

TEST 9

Test 9 corresponds to an experimental failure time of 71 minutes, and occurred deep in the cooling phase of the natural fire. The natural fire was generated using timber wood cribs composed of sticks with a cross-sectional size of $90 \times 90 \text{ mm}^2$. In addition, the fire scenario is characterised by a relatively large opening factor of $0.140 \text{ m}^{1/2}$ and a fire load density of 780 MJ m^{-2} .

Thermal simulation results obtained using the HOPKINS+FLUX thermal module are presented in Figure 5.6. The natural fire curve is approximated using a Γ -factor of 10.24. The numerical results show reasonable agreement with the experimental temperature measurements during the heating phase. However, during the cooling phase, the simulated temperatures decrease more rapidly than observed experimentally, leading to an underprediction of temperature at the time of failure. As observed in the other simulation results, a small spike is present in the numerical thermocouple response, which is attributed to the discrete change of the moving boundary heat flux.




Figure 5.6: Comparison of thermal numerical simulation against experimental data for Test 9 of Renard et al. [32].
Thermal module: HOPKINS+FLUX

The mechanical numerical modelling results for Test 9 obtained using the HOPKINS+FLUX thermal module are summarised in Table 5.10. In contrast to the other natural fire tests, the predicted failure times for Test 9 exhibit a large spread across the different simulation configurations. Both relatively early failure and no failure within the analysed time frame are predicted, depending on the assumed material properties, the inclusion of plastic compression deformation, and the selected strength reduction model. This variability indicates a pronounced sensitivity of the mechanical response to the modelling assumptions for this test.

Table 5.10: Numerical mechanical simulation results for Test 9 of Renard et al. [32].
Thermal module: HOPKINS+FLUX
Test 9: Failure = 71 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (–)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	40.68
2	12.8	40.4	26.4	1.0	WOODEC5	65.60
3	11.5	24.0	19.2	1.0	WOODPRBWE	51.33
4	12.8	40.4	26.4	1.0	WOODPRBWE	False
5	11.5	24.0	19.2	0.0	WOODEC5	30.69
6	12.8	40.4	26.4	0.0	WOODEC5	45.59
7	11.5	24.0	19.2	0.0	WOODPRBWE	42.19
8	12.8	40.4	26.4	0.0	WOODPRBWE	122.45

TEST 12

Test 12 exhibited an experimentally observed failure time of 87 minutes, occurring deep into the cooling phase of the natural fire. The column has a cross-sectional width of 0.34 m, in contrast to the other experiments. The natural fire was generated using timber wood cribs composed of sticks with a cross-sectional size of $90 \times 90 \text{ mm}^2$. The fire scenario is characterised by an opening factor of $0.065 \text{ m}^{1/2}$ and a fire load density of 780 MJ m^{-2} .

Thermal simulation results obtained using the HOPKINS+FLUX thermal module are presented in Figure 5.7. The natural fire curve is approximated using a Γ -factor of 23.72. The numerical simulations show great agreement with the experimental temperature measurements during the heating phase. During the cooling phase, however, the simulated temperatures decrease more rapidly than observed experimentally, resulting in an underprediction of the temperature at the time of failure.




Figure 5.7: Comparison of thermal numerical simulation against experimental data for Test 12 of Renard et al. [32].
Thermal module: HOPKINS+FLUX

The mechanical numerical modelling results for Test 12 obtained using the HOPKINS+FLUX thermal module are summarised in Table 5.11. The majority of the numerical simulations predict no failure. Only two simulations result in column failure. Both of these failure cases correspond to a weaker material configuration combined with the Eurocode-based strength reduction curve.

Table 5.11: Numerical mechanical simulation results for Test 12 of Renard et al. [32].
Thermal module: HOPKINS+FLUX
Test 12: Failure = 87 min

Simulation	E (GPa)	f_c (MPa)	f_t (MPa)	ε_c (-)	Strength reduction curve	Failure time (min)
1	11.5	24.0	19.2	1.0	WOODEC5	334.80
2	12.8	40.4	26.4	1.0	WOODEC5	False
3	11.5	24.0	19.2	1.0	WOODPRBWE	False
4	12.8	40.4	26.4	1.0	WOODPRBWE	False
5	11.5	24.0	19.2	0.0	WOODEC5	87.05
6	12.8	40.4	26.4	0.0	WOODEC5	False
7	11.5	24.0	19.2	0.0	WOODPRBWE	False
8	12.8	40.4	26.4	0.0	WOODPRBWE	False

Key Results

For Test 10 (failure at 47 min), the thermal response predicted by WOODEC5 shows faster heat penetration than measured, while HOPKINS produces a delayed thermal wave relative to the experiments. The corresponding mechanical analyses for Test 10 yield substantially different failure times depending on the thermal module. For WOODEC5, predicted failure occurs between 18.29 min and 28.29 min, whereas for HOPKINS+FLUX predicted failure occurs between 26.21 min and 46.81 min. For Tests 11, 15, and 17 (failures at 35 min, 37 min, and 45 min), the HOPKINS+FLUX simulations show good agreement with the experimental failure times, with predicted failure times spanning 27.56–48.00 min (Test 11), 30.63–43.72 min (Test 15), and 32.80–53.15 min (Test 17). For tests with failure deep into the cooling phase (Tests 9 and 12; failures at 71 min and 87 min), the predicted failure times show larger variability across configurations, including both early failure and no-failure outcomes (Test 9), and predominantly no-failure predictions (Test 12). Across all natural-fire tests, the thermal simulations using HOPKINS+FLUX reproduce the heating-phase thermocouple measurements relatively close, while larger deviations are observed in for Tests 9 and 12 during the cooling phase.

5.2. GPC Surrogate Modelling Results

This section presents the results of the GPC surrogate modelling study. Two six-dimensional input spaces are considered. In both cases, the surrogate model is built up progressively from two- to three- and finally six-dimensional formulations, enabling a consistent assessment of how efficiently the failure domain is to be identified as the problem dimensionality increases.

All surrogate model results reported in this section are based on SAFIR simulations using the WOODDEC5 thermal module and the WOODDEC5 structural module, with an applied plastic compression strain of $\varepsilon_c = 1$. In addition, mean material properties are adopted in accordance with Schilling et al. [40].

The first part of the section considers a six-dimensional model comprising the structural design and fire variables. The second part the section consists a six-dimensional setting that incorporates material variables, and provides a representation of how the GPC model is able to represent resulting stochastic influences to the classification boundary.

5.2.1. 6D Structural Design and Fire Variables

The first surrogate model considers a six-dimensional input space comprising the heating duration (t_h), column width (B), applied axial load (F), column length (l), maximum fire temperature (T_{\max}), and the linear cooling rate (r_c). The bounds adopted for each variable are summarised in Table 4.2.

The results are presented progressively. First, the two-dimensional case (t_h, B) is evaluated, as this extends directly the framework of Gernay [12] by enabling the Duration of Heating Phase (DHP) to be assessed across a full range of column widths. Next, the model is extended to three dimensions by including the applied load (F). Finally, the full six-dimensional formulation is considered by adding l , T_{\max} , and r_c .

For variables not treated as inputs in the considered model, the fixed design parameters are as specified in Subsection 4.3.1. Simulations for the structural design and fire variable models were executed on a laptop with an Intel Core i7-8750H CPU (base frequency 2.20 GHz) and 16 GB RAM.

2D Surrogate Model of Structural Design and Fire Variables (t_h, B)

The two-dimensional surrogate model is first constructed using the heating duration (t_h) and column width (B) as input variables. The simulation procedure is initialised with a physically motivated failure case corresponding to a column width of 0.26 m and a heating duration of 50 minutes.

Convergence of the active learning algorithm is stated after 58 iterations, at which point the boundary spread metric falls below 1%. The total active learning time for this model was 9 hours and 51 minutes, corresponding to an average simulation time of 10.19 minutes per iteration. An overview of the convergence behaviour and monitoring metrics associated with the active learning process is provided in Appendix O.

Figure 5.8 shows the resulting GPC prediction field after application of the probit link function, as defined in Equation 2.36. The dashed contour denotes the decision boundary corresponding to a failure probability of 0.5. It is observed that deep within the non-failure region, the predicted probabilities exhibit a slight reduction away from zero. Further, no simulations are performed deep inside the non-failure domain, and most samples are taken close to the defined decision boundary.

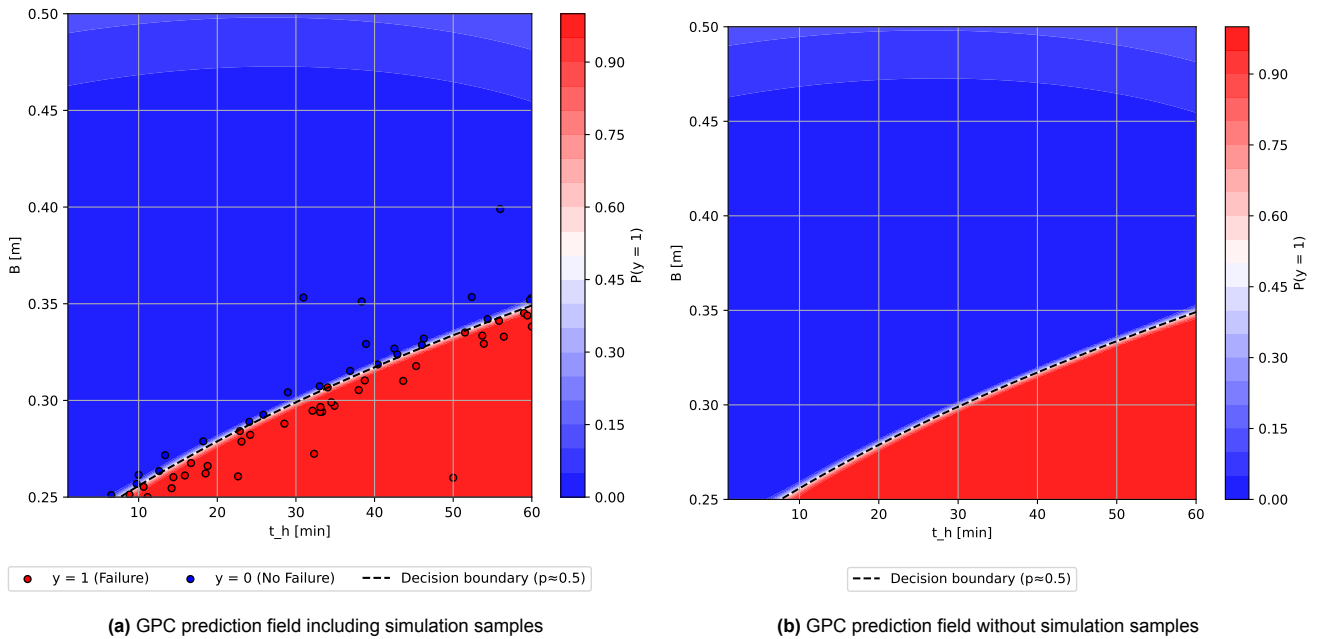


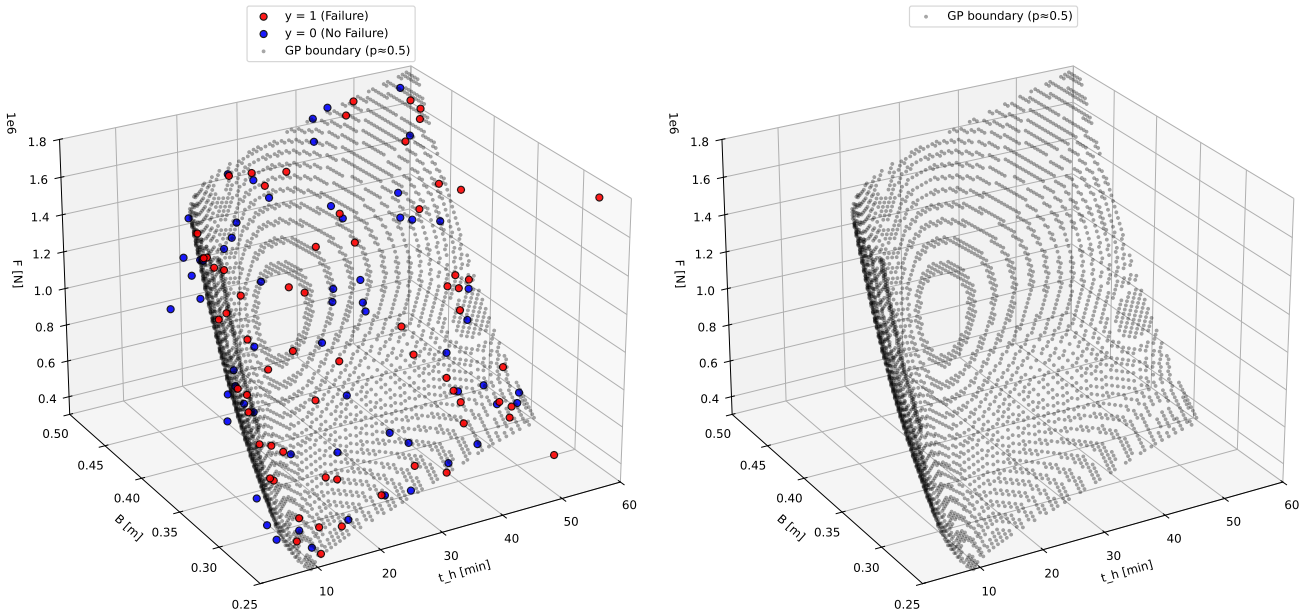
Figure 5.8: Two-dimensional GPC surrogate model for (t_h, B) , illustrating the decision boundary between failure (red) and no failure (blue).

3D Surrogate Model of Structural Design and Fire Variables (t_h, B, F)

Next, the surrogate model is extended to three dimensions by including the applied axial load (F) as an additional input variable. The simulation procedure is initialised with a physically motivated failure case corresponding to a column width of 0.26 m, a heating duration of 50 minutes, and an applied load of 500 kN.

Convergence of the active learning algorithm is achieved after 127 iterations, at which point the boundary spread metric drops below 1%. The total active learning time for this model was 41 hours and 30 minutes, corresponding to an average simulation time of 19.61 minutes per iteration. An overview of the convergence behaviour and monitoring characteristics for this three-dimensional model is provided in Appendix O.

A visualisation of the two dimensional decision boundary surface is presented in Figure 5.9. Consistent with the two-dimensional case, it is observed that no simulation points are located deep within the non-failure region. Instead, the majority of samples are concentrated closely to the classification boundary.




(a) Three-dimensional GPC model, including decision boundary surface, with simulation samples (b) Three-dimensional GPC model including decision boundary surface, without simulation samples

Figure 5.9: Three-dimensional GPC surrogate model for (t_h, B, F) , illustrating the two-dimensional decision boundary separating failure (red) and non-failure (blue).

For the three-dimensional surrogate model, it is no longer possible to directly visualise the full prediction field using colour-coded failure probabilities. However, it is possible to extract two-dimensional slices from the three-dimensional model to support interpretation. These slices are obtained by fixing the non-displayed variable at its standard value, as defined in Subsection 4.3.1.

Figure 5.10 presents the resulting two-dimensional slices for all pairwise combinations of the input variables. The slice in the (t_h, B) plane is consistent with the prediction field and decision boundary of the original two-dimensional surrogate model, with a marginally increase in boundary spread.



(a) Two-dimensional slice in the (t_h, B) plane with F fixed at its standard value (b) Two-dimensional slice in the (t_h, F) plane with B fixed at its standard value (c) Two-dimensional slice in the (F, B) plane with t_h fixed at its standard value

Figure 5.10: Two-dimensional slices of the three-dimensional GPC surrogate model for (t_h, B, F) , illustrating the projected decision boundaries between failure (red) and non-failure (blue).

6D Surrogate Model of Structural Design and Fire Variables ($t_h, B, F, l, T_{\max}, r_c$)

In the six-dimensional surrogate model, all design and fire variables are considered simultaneously, including the heating duration (t_h), column width (B), applied axial load (F), column length (l), maximum fire temperature (T_{\max}), and linear cooling rate (r_c).

The simulation procedure is initialised with a physically motivated failure case corresponding to a heating duration of 50 minutes, a column width of 0.26 m, an applied axial load of 500 kN, a column length of 3.2 m, a maximum fire temperature of 1150 °C, and a cooling rate of 11.5 °C min⁻¹.

Convergence of the active learning algorithm for the six-dimensional surrogate model was achieved after 360 iterations, at which point the boundary spread metric dropped below 1%. The total active learning time for this model was 146 hours and 59 minutes, corresponding to an average simulation time of 24.50 minutes per iteration. An overview of the convergence behaviour and monitoring characteristics for this six-dimensional model is provided in Appendix O.

Due to the high dimensionality of the input space, direct visualisation of simulation samples and the five-dimensional decision hyper-surface is no longer presentable. Instead, interpretation of the trained surrogate model relies on the extraction of lower-dimensional projections and slices from the fitted six-dimensional hypercube model. Once trained, the surrogate model allows arbitrary dependencies between input variables to be examined through such projections.

In this formulation, T_{\max} is treated as an independent input variable. For comparison with the lower-dimensional surrogate models, the dependency of T_{\max} against the heating duration t_h is reintroduced. In particular, projections are considered in which the relationship between t_h and T_{\max} follows the heating-phase fire curve defined in Equation 4.2. For this purpose, t_h is bounded between 6 and 60 minutes. After approximately 6 minutes of heating, a temperature of 600 °C is exceeded, corresponding to the lower bound of the considered T_{\max} range, while a heating duration of 60 minutes results in a temperature of 944.14 °C, which remains within the training bounds of the surrogate model.

To assess consistency with the lower-dimensional surrogate models, two- and three-dimensional slices are extracted from the six-dimensional model, with the remaining variables fixed at their standard values as defined in Subsection 4.3.1. First, Figure 5.11 compares the two-dimensional slice in the (t_h, B) plane obtained from the six-dimensional model with the prediction field of the originally trained two-dimensional surrogate model. A very close correspondence between the two decision boundaries is observed, indicating that the six-dimensional model reproduces the behaviour of the lower-dimensional formulation.

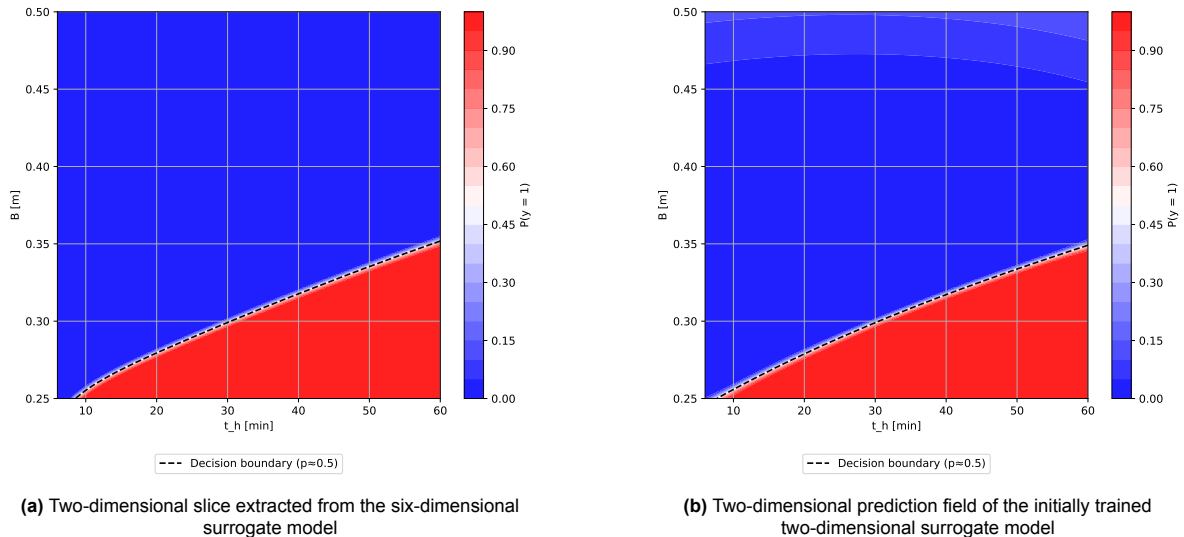


Figure 5.11: Comparison of the decision boundary in the (t_h, B) plane obtained from the six-dimensional surrogate model for structural design and fire variables and the originally trained two-dimensional model.

A similar comparison is performed for the (t_h, F) projection in Figure 5.12, where the two-dimensional slice of the six-dimensional model is compared against the corresponding slice extracted from the

trained three-dimensional surrogate model. Again, a high degree of similarity in the prediction field and decision boundary is observed. An equivalent level of agreement is found for the (F, B) projection, as shown in Figure 5.13.

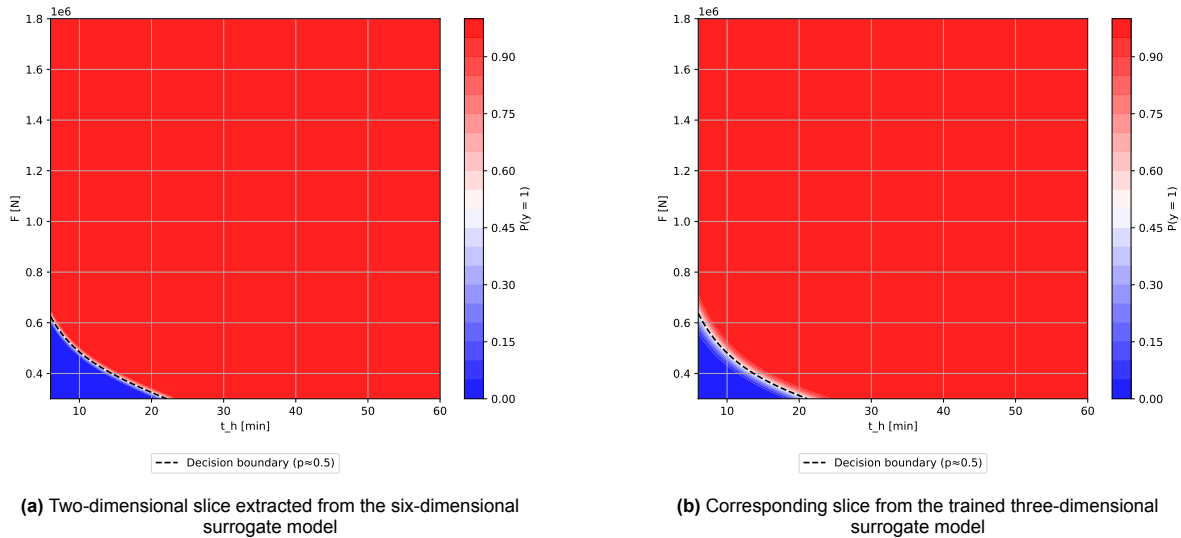


Figure 5.12: Comparison of the decision boundary in the (t_h, F) plane obtained from the six-dimensional surrogate model for structural design and fire variables and the trained three-dimensional model.

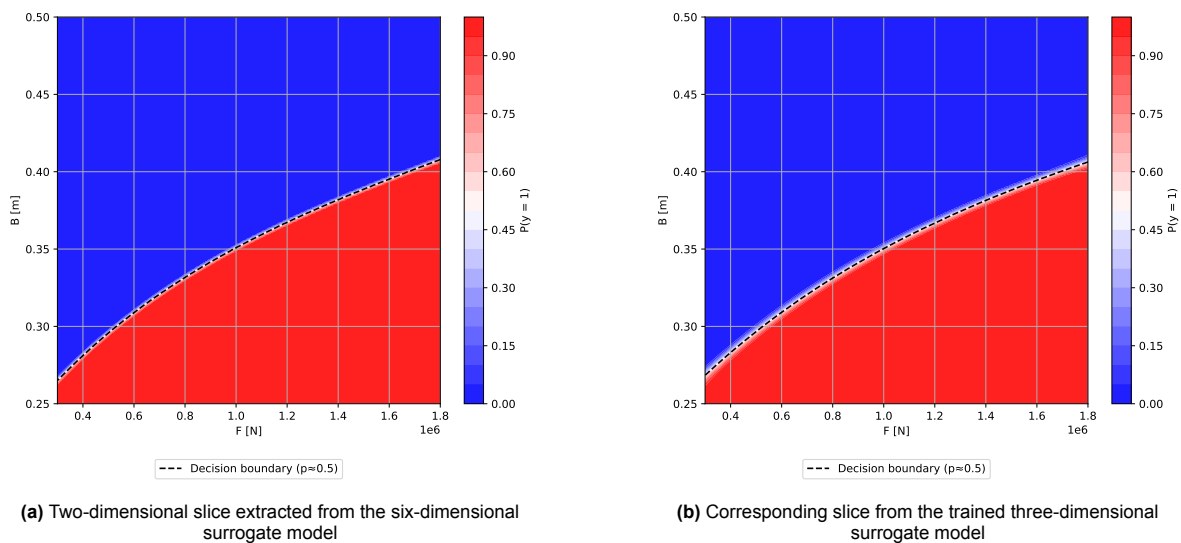


Figure 5.13: Comparison of the decision boundary in the (F, B) plane obtained from the six-dimensional surrogate model for structural design and fire variables and the trained three-dimensional model.

Lastly, a three-dimensional slice of the six-dimensional surrogate model in the (t_h, B, F) space is compared directly with the trained three-dimensional surrogate model in Figure 5.14. The resulting decision boundary surfaces coincide closely, confirming that the six-dimensional surrogate model consistently embeds the lower-dimensional classification behaviour.

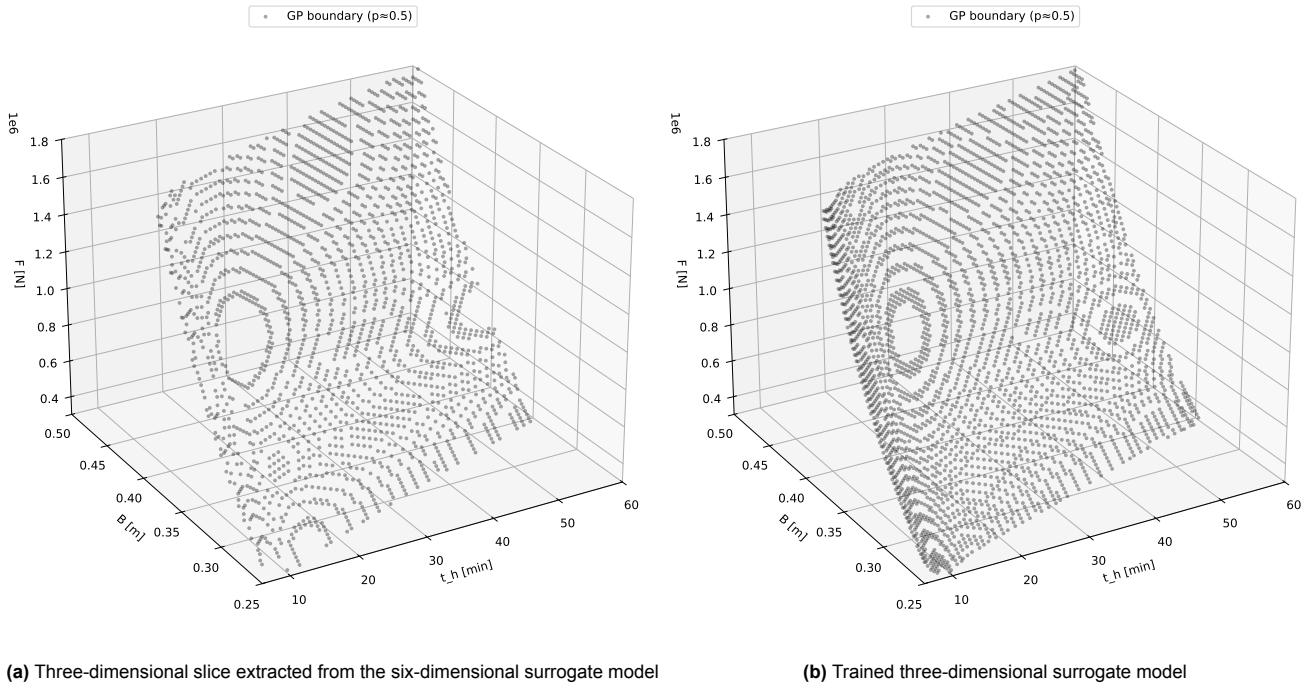


Figure 5.14: Comparison of three-dimensional decision boundary surfaces in the (t_h, B, F) space obtained from the six-dimensional surrogate model for structural design and fire variables and the trained three-dimensional model.

To illustrate, to some extent, how the projected decision boundary of the six-dimensional surrogate model changes when the non-displayed variables are varied, a set of two-dimensional slices in the (t_h, B) plane is presented in Figure 5.15. For consistency with the lower-dimensional models, T_{\max} is again not treated as an independent input in these plots, but is prescribed as a function of the heating duration t_h using the heating-phase fire curve defined in Equation 4.2. The remaining variables (F, l, r_c) are fixed to values corresponding to the 0%, 25%, 50%, 75%, and 100% fractiles of their respective input bounds. This yields a small set of plausible parameter combinations (among many possible), and provides a concise visual indication of how the decision boundary in the (t_h, B) plane shifts when these fixed variables take on more conservative or more severe values.

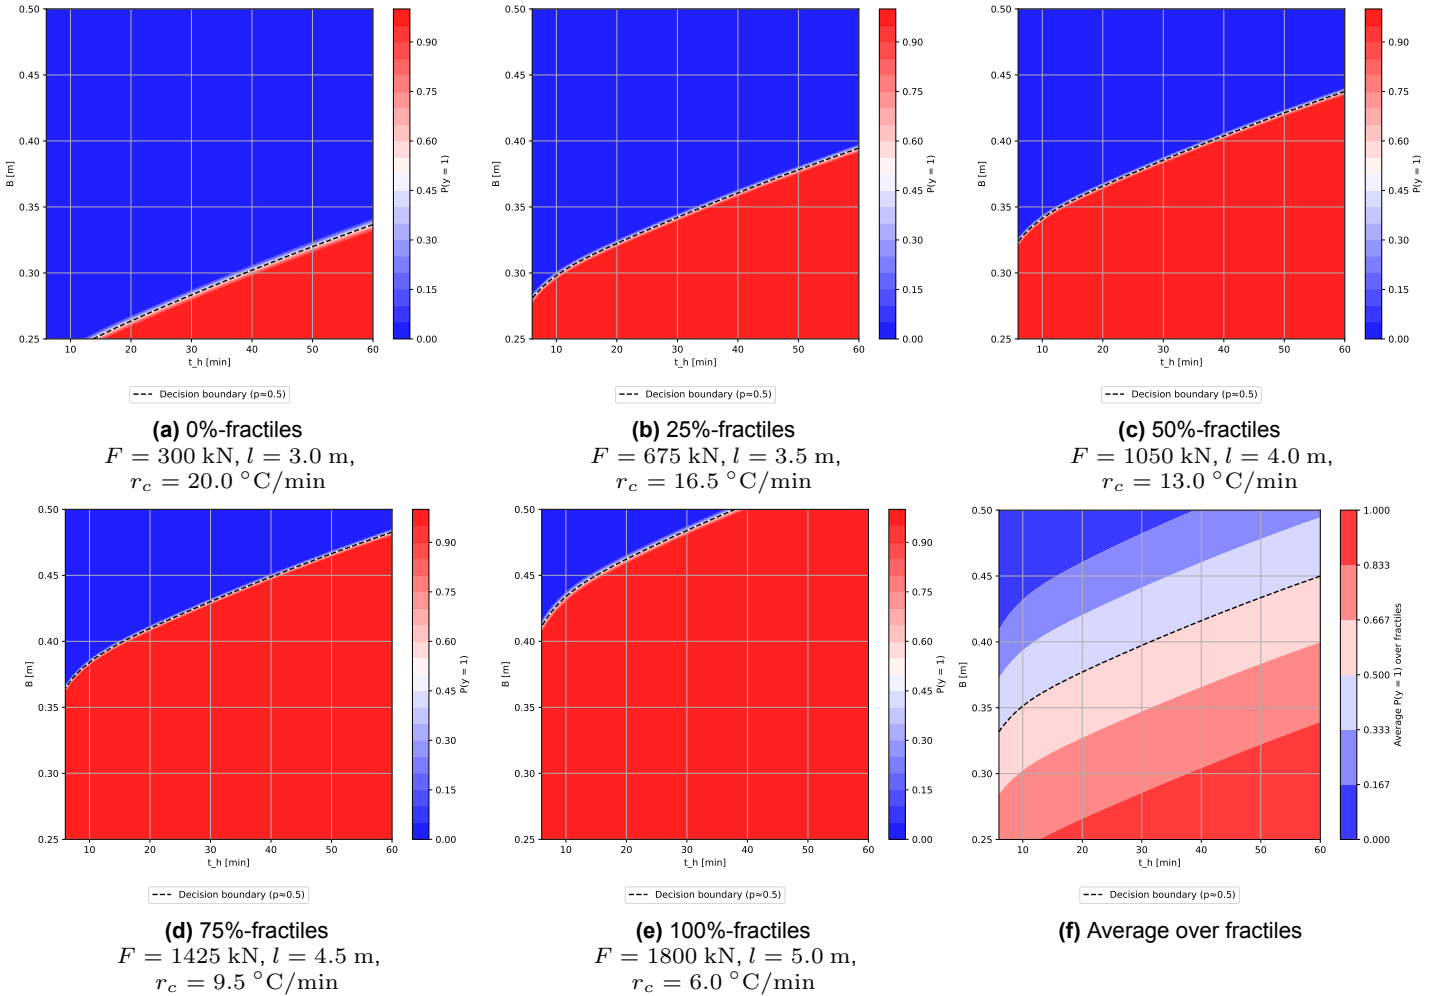


Figure 5.15: 2D-slices of the (t_h, B) plane with quarter fractiles for fixed variables obtained from the six-dimensional surrogate model for structural design and fire variables.

Key Results

The GPC surrogate models for the structural design and fire variables converge for all investigated dimensionalities, with convergence reached after 58 iterations (2D: t_h, B), 127 iterations (3D: t_h, B, F), and 360 iterations (6D: $t_h, B, F, l, T_{\max}, r_c$). Across these models, kernel hyperparameters were periodically re-optimised during active learning using the same initial default strategy adopted. The trained models show the expected sampling behaviour of the active learning procedure, with simulation samples concentrated near the decision boundary and no samples located deep within the non-failure region in the presented 2D and 3D cases. Two- and three-dimensional slices extracted from the six-dimensional model match closely with the corresponding independently trained lower-dimensional surrogate models in the (t_h, B) , (t_h, F) , and (F, B) projections, and in the (t_h, B, F) decision-boundary surface.

5.2.2. 6D Material Variables

This subsection presents the results of the six-dimensional GPC surrogate model incorporating material variables. The input space comprises the heating duration (t_h), column width (B), timber density (ρ), modulus of elasticity (E), compressive strength (f_c), and load eccentricity (e_0).

The results are presented progressively. First, the two-dimensional case (t_h, ρ) is evaluated to assess the influence of material density on the classification behaviour. Next, the model is extended to three dimensions by including the modulus of elasticity (E). Finally, the full six-dimensional formulation is considered by adding the remaining material and geometric parameters (B, f_c , and e_0).

For variables not treated as inputs in a given surrogate model, the remaining design parameters are fixed at their standard values as specified in Subsection 4.3.1. Simulations for the material variable models were executed on a laptop with an Intel Core i7-1185G7 CPU (base frequency 3.0 GHz) and 32 GB RAM.

2D Surrogate Model with Material Variables (t_h, ρ)

The two-dimensional surrogate model is constructed using the heating duration (t_h) and timber density (ρ) as input variables, providing a baseline for assessing the influence of material variables. The configuration of the initial simulation consists of a heating duration of 50 minutes and a timber density of 380 kg m^{-3} . The active learning procedure converges after 68 iterations, at which point the boundary spread metric drops below 1%. The total active learning time for this model was 2 hours and 34 minutes, corresponding to an average simulation time of 2.26 minutes per iteration. Additional monitoring metrics associated with the active learning process are reported in Appendix O.

The resulting GPC prediction field is shown in Figure 5.16. The decision boundary exhibits an approximately linear relationship between t_h and ρ for the considered column configuration.

The distribution of simulation samples are concentrated near the decision boundary, even as in regions extending deep into the corners of the failure domain. In contrast, comparatively fewer samples are located within the non-failure region.

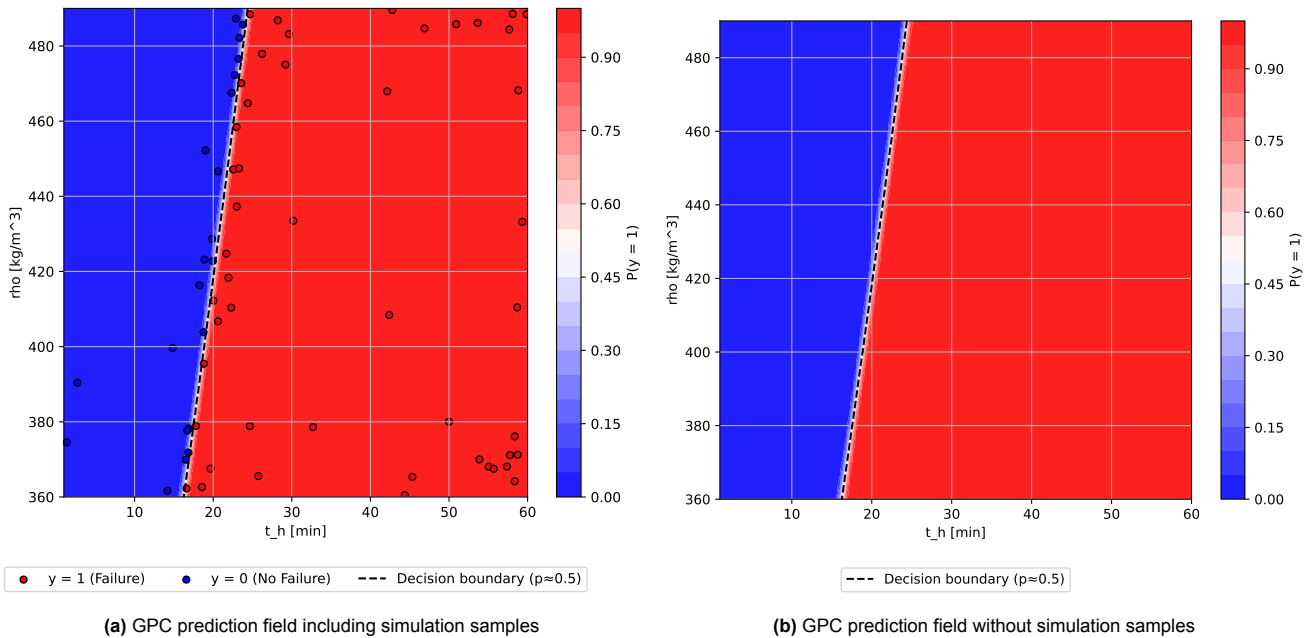


Figure 5.16: Two-dimensional GPC surrogate model for (t_h, ρ) , illustrating the decision boundary between failure (red) and no failure (blue).

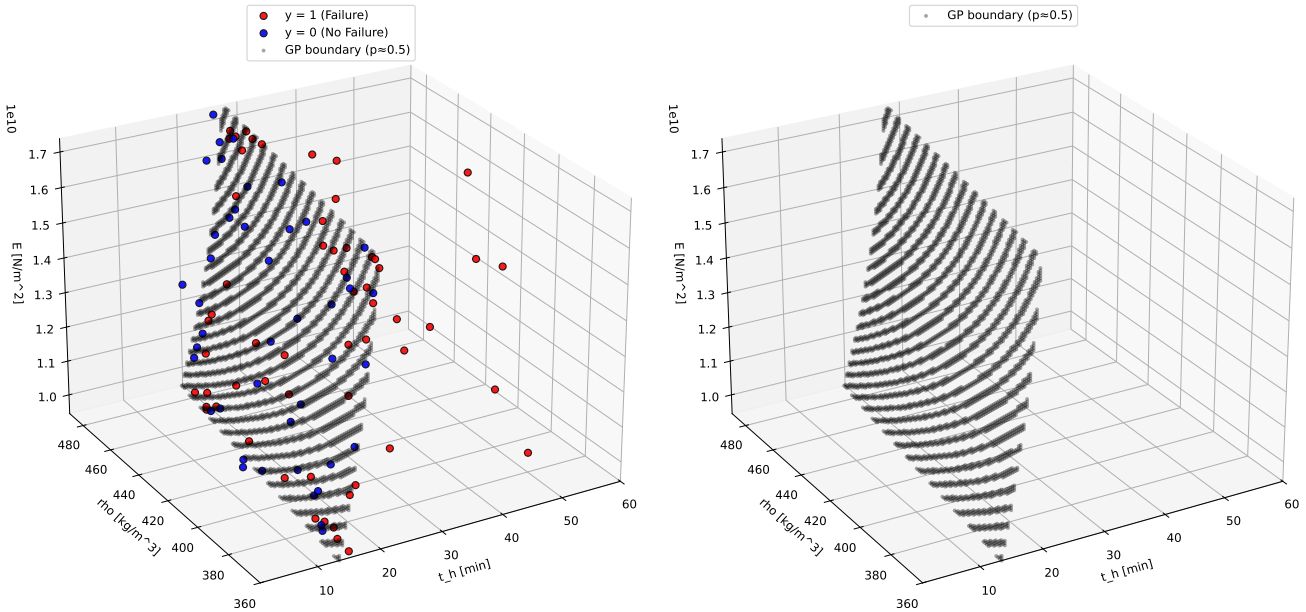
3D Surrogate Model with Material Variables (t_h, ρ, E)

The three-dimensional surrogate model incorporating material variability is constructed using the heating duration (t_h), timber density (ρ), and modulus of elasticity (E) as input variables. The simulation procedure is initialised with a physically motivated failure case corresponding to a heating duration of 50 minutes, a density of 380 kg m^{-3} , and a modulus of elasticity of 10.0 GPa.

For this model, an alternative hyperparameter optimisation strategy is adopted. During the active learning process, only the kernel length scales are updated, while the kernel variance is kept fixed. Convergence of the surrogate model is achieved after 102 iterations, as indicated by the stabilisation of the acquisition values, boundary change, and boundary spread metrics, supported by an increase of the misclassification measure. After convergence, a final optimisation step is performed in which both the kernel variance and length scales are jointly updated to improve the overall model fit. The total active learning time for this model was 4 hours and 27 minutes, corresponding to an average simulation time

of 2.62 minutes per iteration. The measured convergence metrics corresponding to the active learning algorithm are presented in Appendix O.

The resulting three-dimensional decision boundary surface is shown in Figure 5.17. Simulation samples are observed to be concentrated close to the decision boundary, while no samples are located deep within the non-failure region.




(a) Three-dimensional GPC model, including decision boundary surface, with simulation samples

(b) Three-dimensional GPC model, including decision boundary surface, without simulation samples

Figure 5.17: Three-dimensional GPC surrogate model for (t_h, ρ, E) , illustrating the decision boundary surface separating failure (red) and non-failure (blue).

In addition, two-dimensional slices are extracted from the three-dimensional surrogate model for each pairwise combination of the input variables, as shown in Figure 5.18. The slice in the (t_h, ρ) plane is consistent with the prediction field obtained from the previously fitted two-dimensional surrogate model. Observed is a relatively larger boundary spread for the two-dimensional slice of (ρ, E) in comparison to the other projections.



(a) Two-dimensional slice in the (t_h, ρ) plane

(b) Two-dimensional slice in the (t_h, E) plane

(c) Two-dimensional slice in the (ρ, E) plane

Figure 5.18: Two-dimensional slices extracted from the three-dimensional GPC surrogate model for (t_h, ρ, E) , illustrating projected decision boundaries between failure (red) and non-failure (blue).

6D Surrogate Model with Material Variables ($t_h, \rho, E, B, f_c, e_0$)

The six-dimensional surrogate model for material variability incorporates the heating duration (t_h), timber density (ρ), modulus of elasticity (E), column width (B), compressive strength (f_c), and load eccentricity (e_0). The simulation procedure is initialised with a physically motivated failure case corresponding to a heating duration of 50 minutes, a column width of 0.26 m, a density of 380 kg m^{-3} , a modulus of elasticity of 10.0 GPa, a compressive strength of 38.0 MPa, and a load eccentricity of 0.024 m.

Even as for the three-dimensional material model, an alternative hyperparameter optimisation strategy is adopted for this model. During the active learning process, only the kernel length scales are updated, while the kernel variance is kept fixed. Convergence of the surrogate model is achieved after 400 iterations, as indicated by stabilisation of the acquisition values, boundary change, and boundary spread metrics, together with an increase in the misclassification measure. After convergence, a final optimisation step is performed in which both the kernel variance and length scales are jointly optimised to improve the overall model fit. The total active learning time for this model was 31 hours and 59 minutes, corresponding to an average simulation time of 4.80 minutes per iteration. An overview of the monitoring metrics associated with the active learning procedure is provided in Appendix O.

To assess consistency with the lower-dimensional surrogate models, two-dimensional slices are extracted from the six-dimensional model. In particular, the (t_h, ρ) , (t_h, E) , and (ρ, E) projections are considered and compared against the corresponding results obtained from the previously trained two- and three-dimensional surrogate models. The extracted slices show close agreement with the lower-dimensional prediction fields. The projections are presented in Figure 5.19.




Figure 5.19: Two-dimensional slices extracted from the six-dimensional GPC surrogate model incorporating material variables, illustrating projected decision boundaries between failure (red) and non-failure (blue).

In addition, a three-dimensional slice is extracted from the six-dimensional surrogate model in the (t_h, ρ, E) space. The resulting decision boundary surface closely matches that obtained from the independently trained three-dimensional surrogate model. The result is presented in Figure 5.20




Figure 5.20: Three-dimensional slice extracted from the six-dimensional GPC surrogate model incorporating material variables for (t_h, ρ, E) .

Finally, by incorporating material uncertainty explicitly, the six-dimensional surrogate model enables a stochastic evaluation of failure behaviour using Monte Carlo sampling, as described in Subsection 4.3.3. Figure 5.21 presents the results of the stochastic GPC analysis, showing the mean decision boundary together with the associated 5% fractile bounds.




Figure 5.21: Stochastic evaluation of the six-dimensional GPC surrogate model incorporating material variability, showing the mean decision boundary and 5% fractile bounds.

Key Results

For the surrogate models incorporating material variables, convergence is reached after 68 iterations (2D: t_h, ρ), 102 iterations (3D: t_h, ρ, E), and 400 iterations (6D: $t_h, \rho, E, B, f_c, e_0$). For the three- and six-dimensional material-variable models, convergence is not stated solely based on the boundary spread threshold, but based on stabilisation of the active learning monitoring metrics including acquisition values and boundary change, together with an increase in the misclassification measure. For these models, a modified hyperparameter optimisation strategy is applied during active learning in which only the kernel length scales are updated while the kernel variance is kept fixed. After convergence, a final optimisation step is performed in which the kernel variance and length scales are jointly optimised. Two-dimensional slices extracted from the six-dimensional model show close agreement with the decision boundaries obtained from the corresponding lower-dimensional models in the (t_h, ρ) , (t_h, E) , and (ρ, E) projections, and a three-dimensional slice in (t_h, ρ, E) closely matches the independently trained three-dimensional model.

6

Discussions

This chapter discusses the main findings of the study and places them in the context of the research objectives and methodological choices. The discussion is structured in line with the two principal components of the research.

First, the numerical modelling results are examined. This includes an interpretation of the observed numerical modelling response under fire exposure, followed by a critical reflection on the limitations and assumptions inherent to the modelling approach.

Second, the results obtained from the GPC surrogate modelling are discussed. The performance and behaviour of the surrogate models are interpreted, and their limitations with respect to accuracy, robustness, and applicability are addressed.

Finally, the implications of the combined numerical and surrogate modelling framework for fire engineering practice are discussed. This section evaluates the relevance of the findings and outlines the potential practical value of the research.

6.1. Numerical Modelling Discussions

This section discusses the numerical modelling results presented in Section 5.1, with the aim of providing insight into the governing structural and thermal behaviour under fire exposure. The interpretation of the results is addressed first, followed by a discussion of the limitations of the numerical modelling approach

6.1.1. Interpretation of the Numerical Modelling Results

The numerical modelling results are first interpreted for each fire exposure type considered in this study. Specifically, the structural response under standard ISO 834 fire exposure, standardised parametric fire exposure, and natural fire exposure is discussed. Subsequently, an overall interpretation is provided in which the results are compared and common trends and observations are identified across the different fire scenarios.

Standard ISO 834 Fire Exposure

The thermal material properties applied in the numerical model are calibrated for standard fire exposures. As the experimental tests considered in this study comprise ISO 834 fire scenarios, a reliable representation of the thermal response is therefore expected.

This expectation is reflected in the numerical predictions of failure time, which show proper agreement with the experimental observations. The numerical results are of comparable magnitude to the experimentally observed failure times, indicating that the thermo-mechanical response of the column is overall adequately captured.

With the assumption of an accurate thermal response of the column, which is to be expected for standard ISO 834 fire exposures, it follows that simulations based on mean material states combined with mean strength reduction behaviour are expected to show the closest agreement with the experimental results.

An examination of the numerical simulations indicates that the Eurocode-based strength reduction

curve, when combined with a mean material state and the inclusion of compressive plastic strain, provides the closest agreement with the experimentally observed failure times for both ISO fire tests. At the same time, comparable agreement is also achieved using a mean strength reduction curve, represented by the WOODPRBWE model with a Weibull quantile of 0.5. When allowing for plausible variations in material state, compressive strain development, and strength reduction behaviour within the considered parameter ranges, the predicted failure times remain consistent with the experimental observations. This is reflected by the fact that the numerical results enclose the experimentally observed failure times.

In contrast, when both the Eurocode-based material state and the Eurocode strength reduction curve are applied simultaneously, the numerical simulations consistently underestimate the experimentally observed failure times. This underestimation is observed for all ISO fire tests and occurs irrespective of whether compressive plastic strain is included in the mechanical model.

Standardised Parametric Fire Exposure

For the standardised parametric fire exposure, a noticeably wider spread in the experimentally observed failure times is observed for Tests 3 and 6. This increased variability is reasonable to expect, as these fire scenarios are close to the transition between failure and non-failure. In this regime, relatively small variations in material properties, geometric configuration, or mechanical response are likely to have more influence on whether failure occurs and on the associated failure time.

As a result, the predicted failure times for these tests are expected to exhibit a higher variability compared to the ISO 834 fire scenarios. Nevertheless, it is important to note that failure was experimentally observed for both columns subjected to the standardised parametric fire exposure.

In contrast, the numerical modelling results show a substantially larger spread, ranging from early failure at approximately 43.6 minutes to cases in which no failure is predicted for the majority of simulations. The numerical results therefore exhibit limited consistency and do not clearly reproduce the experimentally observed failure behaviour. This discrepancy raises questions regarding the adequacy of the thermal modelling approach when applied to standardised parametric fire exposures.

Natural Fire Exposure

For the natural fire exposure scenarios, notable differences are observed between the applied thermal modelling approaches. When using the WOODDEC5 thermal module, heat is observed to propagate too rapidly through the cross-section. Conversely, when applying only the adapted conduction model proposed by Hopkin et al. [44], combined with an approximated fire curve through the use of a Γ -factor, the thermal response does not align well with the experimental thermocouple data. In this case, heat propagation through the cross-section is overly restricted.

The numerical simulations employing the WOODDEC5 thermal module consequently show a consistent underestimation for the predicted failure times in comparison to the experimental test. This behaviour reinforces the observed fast heat penetration into the cross-section, resulting in a prematurely weakened structural response of the timber column.

Introducing an additional moving heat flux improves the agreement with the experimentally measured thermocouple data. When the thermal response is approximated more accurately in this manner, the numerical simulations yield consistent and reasonable structural failure times that lie within the range of the experimental observations. This applies both to columns that fail during the heating phase (Tests 15 and 17) and to those that fail shortly after the onset of the cooling phase (Tests 10 and 11). The variability observed in these results is comparable to that previously noted for the standard ISO 834 fire exposure scenarios.

More significant difficulties arise for experiments in which failure occurs deep into the cooling phase of the fire. For these cases, specifically Tests 12 and 9, the thermal response of the column is not captured accurately. As for the thermal modelling, the sustained high temperatures within the core of the cross-section during the cooling phase are not adequately represented and are underpredicted. When this thermal behaviour is not accurately captured, the predicted structural failure times deviate substantially from the experimental observations. The discrepancies observed for these cases therefore might be a direct consequence of limitations in the thermal modelling of the cooling phase.

Comparative Interpretation

A comparison across the different fire exposure types reveals consistent patterns in the predictive capability of the numerical modelling approach. For standardised parametric fire exposures and for natural fire scenarios in which column failure occurs deep into the cooling phase, the numerical predictions show limited reliability. In these cases, a wide spread in predicted failure times is observed, ranging from early failure to scenarios in which no failure is predicted at all, indicating a lack of consistency in the structural response. This behaviour is most likely amplified by limitations in the thermal modelling during the cooling phase, as sustained elevated internal temperatures are not captured accurately and therefore the predicted structural response becomes highly sensitive.

In contrast, for fire scenarios in which column failure occurs during the heating phase, the numerical modelling results show substantially better agreement with the experimental observations. This behaviour is observed for the standard ISO 834 fire exposure when using the WOODDEC5 thermal module, as well as for natural fire exposures modelled with the HOPKINS+FLUX thermal module. In these cases, the predicted structural failure times align well with the experimentally observed failure behaviour.

A similar level of agreement is also obtained for natural fire exposures in which column failure occurs shortly after the onset of the cooling phase. This indicates that, provided the thermal response of the column is captured with sufficient accuracy during both the heating phase and the early cooling phase, the numerical model is capable of producing reliable predictions of structural failure time.

6.1.2. Limitations of the Numerical Modelling

The numerical modelling approach adopted in this study provides a structured and potentially accurate framework for approximating the complex thermo-mechanical behaviour of timber columns exposed to fire. Nevertheless, as with any numerical model, it remains a simplified representation of reality and is subject to several limitations arising from modelling assumptions, idealisations, and computational constraints.

First of all, a primary fundamental limitation in the modelling approach is the use of a weakly coupled thermal–mechanical analysis. While this approach allows the individual physical processes to be described in detail, it neglects the direct interaction between thermal, mechanical, and physical-chemical processes. In reality, heat transfer, moisture transport, material degradation, and mechanical response are correlated and evolve simultaneously.

Secondly, several simplifications are introduced in the thermal analysis. Moisture transport within the timber cross-section and changes in mass transfer are neglected. These processes influence thermal conductivity, heat capacity, and density. Their omission introduces simplifications in the thermal modelling and, consequently, limitations in the predicted thermal response.

A further major limitation arises from the fact that the applied thermal material properties are effectively calibrated for standard fire curves. This implies multiple physical processes to be accounted for within these properties. As a result, their applicability to faster heating fire scenarios and to the cooling phase of fires is limited.

To improve agreement with the experimental thermocouple data, the conductivity model proposed by Hopkins and an additional walking heat flux boundary condition are introduced. These measures provide an improvement to describe the thermal response. However, their physical interpretation remains open for discussion.

Although their implementation is consistent in terms of physical units, heat released by combustion of the timber to itself is expected to be mainly relevant during the cooling phase of the fire. Combustion requires sufficient oxygen, which during the heating phase is largely consumed by the surrounding burning materials. During this phase, the structural column therefore mainly acts as a responding specimen. During the cooling phase, oxygen levels in the compartment typically increase. At the same time, the timber column remains at elevated temperatures. Under these conditions, heat release due to combustion of the specimen itself is expected to become more relevant, contributing to the sustained thermal response observed experimentally.

A more physically consistent representation of pyrolysis and combustion could therefore be achieved by introducing an additional heat source term in the governing heat conduction equation using an Arrhenius-based formulation, as proposed by Hopkin [43]. This approach provides a more complete physical description by accounting for reaction kinetics and activation energies. Within this framework,

released energy is not limited to the actively charring front but is also able to occur at lower timber temperatures that persist during the cooling phase.

Despite this improved physical formulation, it should be noted that the calibration of the additional heat source in this modelling approach still depends on thermal material properties that are themselves calibrated for standard fire conditions. The introduced heat term is not independently derived from combustion physics, but instead inferred from modelling variables that already embed specific fire behaviour. In turn, this raises questions regarding the appropriateness of this approach.

Moreover, the Hopkins conductivity model relies on an approximate representation of the fire curve through the use of a Γ -factor. This parameter represents an order-of-magnitude approximation and is inherently uncertain, as it depends on variations in compartment oxygen availability and fuel characteristics. In addition, the combustion efficiency associated with the applied additional heat flux boundary remains an arbitrary and uncertain parameter. Both aspects may vary significantly between individual cases.

Despite these limitations, the approach provides a physically motivated means to assess whether the structural model is capable of reproducing failure times when the thermal response is matched more closely.

Finally, more general thermal simulation limitations arise from the spatial discretisation and the applied boundary conditions. A trade-off is made between mesh resolution and computational efficiency, which is required for surrogate modelling applications. While finer meshes would likely improve convergence and accuracy, it would also significantly increase computational cost. In addition, boundary conditions in the thermal analysis are simplified and treated as constants. In reality, convective heat transfer coefficients and radiative exchange vary with temperature and fire dynamics, which is not captured in the present modelling approach.

Further, the mechanical analysis introduces additional simplifications that introduces limitations in the structural response. To begin with, the constitutive behaviour of timber is represented using a bi-linear material model, and plastic strain development is treated in a simplified manner. While this enables efficient numerical implementation, it does not fully capture the complex, temperature-dependent non-linear behaviour of timber under compression.

Subsequently, the structural response is modelled using Euler–Bernoulli beam theory, which inherently neglects three-dimensional effects. As a result, phenomena such as local buckling are not represented. This restricts the ability of the model to capture more detailed evolving failure mechanisms that may occur under fire exposure.

In addition, changes in self-weight due to moisture evaporation and mass loss are not considered. In reality, the progressive reduction in density during fire exposure would lead to a gradual decrease in axial load. This may influence second-order effects and, consequently, the overall stability of the column.

Furthermore, initial geometric imperfections are represented through an imposed eccentricity. Although this approach captures the dominant bending behaviour at the critical mid-height section, alternative representations, such as an initial sinusoidal curvature, could provide a more realistic description of the initial deformation state and its influence on stability behaviour.

Lastly, additional limitations arise from simplifications in the applied fire load. The thermal exposure is assumed to be uniform along the height of the column. In reality, temperature gradients may develop within the compartment, with higher thermal exposure near the ceiling level.

Furthermore, auto-extinguishment is implicitly assumed, and the column is considered to remain stable after the onset of cooling at the cross-sectional centre. In practice, continued local smouldering may occur, leading to ongoing degradation of stiffness and strength and, consequently, potentially delayed failure.

6.2. GPC Surrogate Modelling Discussions

This section discusses the results of the Gaussian Process Classification (GPC) surrogate modelling presented in Section 5.2. First, the behaviour and performance of the surrogate models are interpreted. Subsequently, the limitations and assumptions associated with the adopted surrogate modelling approach are discussed.

6.2.1. Interpretation of the GPC Surrogate Modelling Results

An examination of the GPC surrogate modelling results indicates that both the six-dimensional model incorporating structural design and fire variables and the model describing material variables are able to reproduce the deterministic numerical decision boundary in a consistent manner.

Lower-dimensional slices extracted from the higher-dimensional surrogate models exhibit behaviour that is comparable to that of independently trained lower-dimensional models. This suggests that the dominant response characteristics are preserved across different dimensional representations.

Furthermore, introducing dependent variables through transformations of the included inputs yields projections with no noticeable loss in accuracy. This suggests that interpretation-driven dimensionality reduction does not substantially affect the qualitative formulation of the classification boundary.

The monotonicity assumption imposed in advance also appears to be reasonable for the considered cases. Across the analysed results, no strong curvature or abrupt changes in the decision boundary are observed, suggesting that the assumed monotonic behaviour is consistent with the numerical simulation data.

With respect to the hyperparameter optimisation strategy, different behaviours are observed for the two surrogate modelling cases. For the six-dimensional model incorporating design and fire variables, periodic optimisation of the kernel hyperparameters was performed every ten active learning iterations once the decision boundary had been determined to be sufficiently enclosed. This strategy appeared to be effective in maintaining a stable description of the classification boundary.

In contrast, for the six-dimensional model incorporating material variables, the same optimisation strategy proved less consistent. As a result, optimisation of the kernel variance was not applied during the active learning process. Convergence for this model was instead assessed based on other active learning metrics. Despite the absence of kernel variance optimisation, the resulting surrogate model still produced reasonable classification behaviour, although a larger number of active learning iterations was applied.

These observations suggest that kernel scaling optimisation becomes more challenging when variables with differing spatial influence are combined. In the six-dimensional model including material variables, the heating duration and column width are defined over relatively wide bounds and therefore influence the response across a broad region of the input space. On the other hand, the material variables considered are confined to narrower bounds and primarily affect the response in more localised regions. It seems that when the governing variables exert a more uniform, global influence, the active learning process appears smoother and intermediate hyperparameter optimisation is more effective. By contrast, when material variables are included, the decision boundary becomes more localised, requiring both short and long characteristic length scales to capture variable interactions, which may hinder stable kernel variance optimisation. It is therefore possible that if material variables were defined over a wider range, such that their influence extends more globally across the input space, this behaviour could become less pronounced.

When observing the stochastic decision boundary associated with material uncertainty, a modest increase in the spread between the 5% fractile boundaries and the mean decision boundary is to be identified with increasing column width. This behaviour appears reasonable, as the influence of material variability is expected to become more pronounced when the effective area over which material properties act increases. It should be noted that the material variables are treated as statistically independent in this evaluation. In reality, material properties of timber exhibit correlation. If such correlation were included, a larger spread around the mean decision boundary could be expected, as correlated variables tend to vary jointly and therefore cause a wider spread of probability mass.

The computational advantage of the surrogate model becomes particularly evident for the stochastic evaluation with material uncertainty. The plot in Figure 5.21 is evaluated on a grid of 150 by 150 points in the (t_h, B) space, with 200 Monte Carlo samples per grid point, which corresponds to 4.5 million realisations. Using only SAFIR and adopting the average simulation time of 4.80 minutes observed for the six-dimensional material-variable model would result in about 41.10 years of computation. In contrast, the surrogate model required 31 hours and 59 minutes of active learning, in which the additional subsequent evaluation of the GPC model takes only minutes.

Notably, this estimate is already conservative, and the true computational cost would in fact be even higher. This is because simulations become substantially slower for larger cross-sections, since mesh resolution must be maintained for accuracy and the thermal wave requires more time to propagate

through the section. This effect is observed in the six-dimensional structural design and fire-variable model, where additional input dimensions shift parts of the decision boundary towards larger cross-sections, which forces the active learning procedure to sample these more computationally demanding cases. As a result, the average simulation time increases to 24.50 minutes and individual runs can take up to hours. Although these simulations were executed on a slower machine, this further supports that the 4.80-minute average used above is a conservative baseline, and a full SAFIR-based stochastic grid evaluation would be even less feasible.

6.2.2. Limitations of the GPC Surrogate Modelling

While the GPC surrogate modelling framework proves effective for learning classification boundaries in high-dimensional spaces, several limitations and trade-offs arise from modelling assumptions, optimisation choices, and computational considerations.

A key limitation concerns the optimisation of kernel hyperparameters during the active learning process, in which particularly the kernel scaling (variance). The effectiveness of kernel scaling optimisation appears to be problem dependent. Optimising the kernel scaling improves the fit of the surrogate model and accelerate convergence towards an accurate decision boundary. However, it may also adversely affect the acquisition function by distorting the relative importance of regions in the input space, potentially leading to suboptimal sampling behaviour. In contrast, omitting kernel scaling optimisation and restricting optimisation to kernel length scales provides a more robust and stable active learning process, at the expense of slower convergence. Importantly, the results indicate that the decision boundary can still be captured reliably without kernel scaling optimisation, provided that a final complete hyperparameter optimisation is performed once convergence has been established. This introduces a trade-off between robustness and efficiency during the active learning algorithm, which considers careful monitoring using appropriate convergence metrics.

A further limitation relates to how convergence is defined. The current monitoring metrics mainly provide a global indication of whether the decision boundary has stabilised, but they may be less sensitive to local changes. A more informative approach could be to directly evaluate the local effect of adding a new data point to the data set, using a procedure similar to the acquisition function but applied after the point has selected. This could also be implemented using importance sampling strategy and would allow the actual influence of new samples on the decision boundary to be quantified. Compared to the current boundary movement metric, such an approach would offer a more local and intuitive measure of convergence.

The surrogate modelling approach is further challenged by increasing dimensionality. In higher dimensional spaces, defining convergence becomes more difficult as the monitoring samples become increasingly sparse relative to the size of the input space, while the decision boundary itself occupies a larger and possibly more complex region.

In addition, in this study, two six-dimensional surrogate models are constructed, representing a practical stopping point in terms of efficiency and computational feasibility. However, the structural response of timber columns exposed to fire is governed by interacting physical processes involving a larger set of variables, which are not all included within a single surrogate model. A more comprehensive representation of the underlying physical interactions would therefore require combining these variables within a higher-dimensional input space, and potentially incorporating additional influencing parameters. Such an extension would be more representative to the real system, but would also necessitate a substantially larger number of training samples to achieve a comparable level of confidence of the learned decision boundary.

An additional consideration lies in the choice of the prior and the initial sampling strategy. In this study, the active learning process is biased towards one class, effectively approaching the decision boundary from the failure domain. This choice is motivated by computational constraints and exploits the boundary-seeking behaviour of the SMOCU-based acquisition strategy. While effective in the present context, this introduces a dependence on the assumed monotonic behaviour of the response. With greater computational resources and less disruption of computational time, an alternative strategy could involve initial sampling throughout the design space with a neutral prior, allowing the SMOCU algorithm to locate the decision boundary without bias towards a specific class. In such a setting, the monotonicity assumption would no longer be required, although it seems to remain valid for the considered applica-

tion.

More generally, the GPC framework inherits the fundamental assumptions of Gaussian process models. In particular, it is assumed that any finite subset of the input space follows a joint Gaussian distribution. While this assumption is common and often effective in practice, it may not fully reflect the true complexity of the underlying physical processes. In addition, the choice of kernel function implicitly imposes assumptions on the smoothness and shape of the decision boundary. Although the selected kernel appears suitable for the present problem, it nevertheless restricts the class of functions that can be represented.

Finally, Gaussian process surrogate models are inherently limited in their ability to extrapolate beyond the range of the training data. Reliable predictions can only be expected within the region of the input space that has been sufficiently sampled. Moreover, as the number of training samples increases, the computational cost of training and updating the surrogate model grows rapidly. This limits the practical scalability of the approach, particularly for higher-dimensional problems and still fine resolution of the decision boundary is required.

6.3. Implications for Fire Engineering

The findings of this study invite reflection on how delayed failure and burnout resistance are approached within fire engineering. Structural response under fire exposure is governed by complicated interacting thermal, mechanical, and material degradation processes that evolve over extended time periods. This is especially the case when the cooling phase of the fire starts to be included. This complexity poses challenges for both modelling and design-oriented assessment.

The complexity of fire-induced structural behaviour often leads to an approach in which individual physical processes are modelled separately. Fire dynamics are described using fire models, heat transfer using thermal models, and structural response using mechanical models. Numerical modelling is commonly adopted to represent these processes in as much physical detail as needed. Nevertheless, each of these models already constitutes an approximation of reality.

When surrogate models are subsequently constructed on top of numerical simulations, an additional level of abstraction is introduced. This layering of models increases the importance of accuracy and physical consistency at each modelling level. Any uncertainty or inaccuracy in the underlying numerical representation directly affects the reliability of the surrogate model. As a result, the need for carefully calibrated and validated numerical models becomes even more pronounced when surrogate modelling is employed.

At the same time, the surrogate modelling offers a clear advantage in the fire engineering context. Once trained, the surrogate model allows the influence of multiple interacting variables and uncertainties to be explored across a wide parameter space, which would be impractical using detailed numerical simulations alone. This is particularly relevant for early stage design-oriented assessments, where iterative design procedures require rapid insight into trends and sensitivities. Within this setting, the proposed framework provides a way to engage with model uncertainty explicitly rather than implicitly accepting it. By combining surrogate modelling with probabilistic descriptions of input variability and classification uncertainty, modelling assumptions and their effects can be examined more transparently. Beyond interpretation, these analyses can be synthesised into simplified design-oriented representations that are more accessible for practical use, such as design charts, lookup tables, or equation-based rules. Conceptually, this is analogous to the formation of some classical design codes (e.g. current Eurocode buckling curves), but extended to include fire-exposure variables and cooling-phase effects.

In addition, this layered modelling approach places increased importance on how the fire exposure is defined. When delayed failure and cooling-phase effects are considered, the representation of the fire load becomes particularly influential. Variations in heating phase characteristics and fire decay behaviour can lead to significant differences in the predicted structural response, indicating that fire parametrisation is not only a modelling detail but an integral part of the burnout resistance assessment. Within the proposed framework, alternative fire representations are able to be explored in a consistent manner, for example by incorporating exponentially decaying fire curves or other idealised descriptions of fire exposures. While such parametrisations do not eliminate the limitations of the fire curve

parametric approximation, they provide a structured way to examine how fire exposure characteristics propagate through the modelling chain and affect structural response.

Since burnout resistance assessment inherently requires specification of the full fire curve, the assumed fire exposure is closely linked to the applied building typology in design practice. This is because compartment characteristics, fuel load, and ventilation conditions largely govern the fire scenario. Therefore, a framework in which representative building typologies are directly associated with corresponding idealised fire curves would support first-order design assessments.

From a practical perspective, the combined numerical and surrogate modelling approach can therefore be interpreted as a workflow for translating simulation-based assessments into design-oriented guidance, as illustrated in Figure 6.1. As the design progresses and more detailed information becomes available, a full fire analysis combined with thermal and structural modelling would typically be carried out in any case. The primary practical value of the surrogate-based approach therefore lies in supporting early-stage decision-making, where typology-linked idealised fire curves and surrogate-derived design rules can provide a first quantitative indication of burnout resistance.

As discussed, the extent to which proper design recommendations are to be formulated remains conditional on the fidelity of the underlying numerical model. The framework should therefore be viewed as demonstrating methodological potential, while improved calibration and validation are required before the workflow can be generalised with confidence. Finally, while the present work focuses on timber columns, the implications are not limited to a single material system, and similar challenges have been identified for other structural materials in which delayed failure may occur, as discussed by Gernay and McNamee [2].

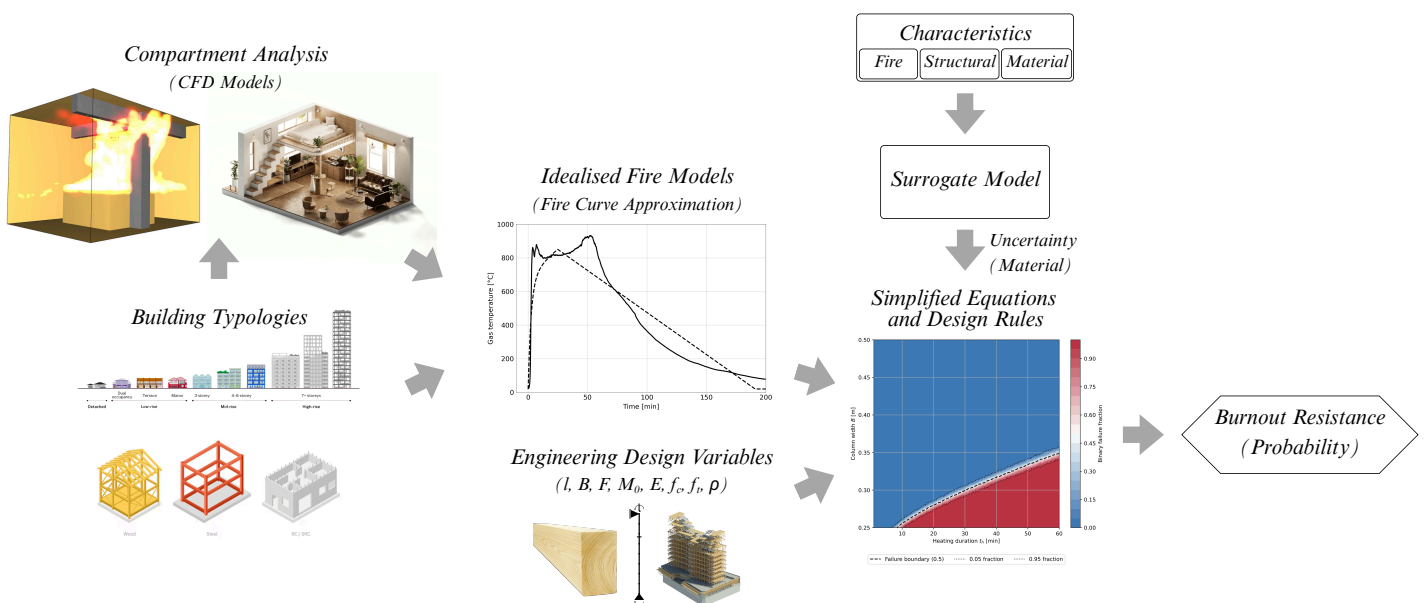


Figure 6.1: Workflow for implications into fire engineering
(Image sources: Thornton Tomasetti, ArchDaily, 2022, Shutterstock, and ArchitectureAU, 2025)

7

Conclusions and Recommendations

This chapter presents the main conclusions of the study and outlines recommendations for further research. To maintain a clear link with the aim and structure defined in Section 1.5, the conclusions are structured around the two principal objectives and components of the work. First, conclusions are drawn on the applicability of non-linear thermo-mechanical finite element modelling in SAFIR for describing delayed failure behaviour of glulam timber columns under natural fire exposure. Second, conclusions are formulated on the applicability of Gaussian Process Classification as a surrogate modelling framework for burnout resistance assessment. Finally, recommendations for further research are provided for both the numerical and surrogate modelling components.

Research Questions

The main research question was formulated as follows.

How are the delayed failure behaviour and burnout resistance of glulam timber columns under natural fire exposure, including both heating and cooling phases, modelled and predicted in a reliable and computationally efficient manner?

This question was addressed through four sub-questions.

- Which physical processes and governing variables influence the thermo-mechanical response of glulam timber columns subjected to realistic fire exposures, including both the heating and cooling phases?
- To what extent is non-linear thermo-mechanical finite element modelling in SAFIR capable of reproducing experimentally observed delayed failure behaviour of glulam timber columns under natural fire exposure?
- To what extent is Gaussian Process Classification a reliable and useful surrogate modelling approach for approximating burnout resistance outcomes predicted by numerical simulations of glulam timber columns?
- How can Gaussian Process Classification be applied to support efficient parametric and probabilistic assessment of burnout resistance for glulam timber columns?

Summary of Main Conclusions

The main conclusions of this study are summarised in the following list. They reflect the findings for both the numerical modelling component and the surrogate modelling component, and highlight the key limitations and strengths identified throughout the thesis.

- Non-linear thermo-mechanical finite element modelling in SAFIR is able to reproduce failure times reasonably for heating-dominated scenarios, but it is not fully capable of reliably predicting delayed failure when the cooling phase governs the response.
- The dominant limitation of the numerical approach appears to be the thermal modelling during cooling, where sustained elevated core temperatures are systematically underpredicted, leading to large errors and inconsistent failure or no-failure outcomes for cooling-dominated tests.

- The results indicate that continued pyrolysis and smouldering effects during cooling are likely key mechanisms driving delayed failure and burnout behaviour, and they must be represented explicitly to achieve reliable predictions.
- Gaussian Process Classification coupled with active learning is shown to be a reliable and efficient surrogate framework for learning the deterministic failure boundary produced by the numerical simulations across parametrically defined input spaces.
- Once trained, the surrogate model demonstrates the ability to support rapid parametric studies and uncertainty propagation, enabling computationally feasible probabilistic burnout resistance assessment.
- The reliability and physical interpretability of the surrogate results remain limited by the fidelity of the underlying numerical model, particularly the accuracy of the cooling-phase thermal response.

7.1. Numerical Modelling Conclusions

This section summarises the main conclusions related to the numerical modelling component of the study in line with the first objective and the first two sub-questions defined in Section 1.5. The conclusion addresses the applicability of non-linear thermo-mechanical finite element modelling in SAFIR for describing delayed failure and burnout resistance of glulam timber columns under realistic fire exposures. Specific attention is given to the governing physical processes and variables, and to the ability of the numerical modelling approach to reproduce experimentally observed delayed failure behaviour during both the heating and cooling phases of fire exposure.

Overall, it is concluded that the applied non-linear thermo-mechanical finite element modelling approach is not fully capable of accurately reproducing delayed failure behaviour of glulam timber columns under realistic fire exposure. The primary limitation is found to lie in the modelling of the thermal response during the cooling phase of the fire.

The results show that the available thermal material properties are insufficient to represent the sustained elevated temperatures within the cross-section during cooling. While reasonable agreement can be achieved during the heating phase, and partially improved through adapted modelling strategies, the thermal response during cooling remains poorly captured. In particular, the persistence of high internal temperatures, which strongly governs delayed failure behaviour, cannot be reproduced reliably with the adopted thermal models. Introducing an additional heat source term in the governing heat conduction equation is likely to improve the representation of this behaviour to some extent. However, such an approach does not remove the underlying limitation that the available thermal material properties are not sufficient to accurately describe cooling-phase behaviour. Which is especially visible in the thermal couple mismatch made for the higher temperatures during cooling.

This conclusion is supported by the observation that, when the experimentally measured thermocouple data are matched more closely, the predicted structural failure times become consistent and show good agreement with experimental observations across multiple tests. Conversely, for standardised parametric fires and natural fire scenarios in which failure occurs deep into the cooling phase, the numerical predictions deteriorate significantly. This confirms that inaccuracies in the thermal modelling of the cooling phase are the dominant source of error.

In response to the first research question, it is therefore concluded that continued pyrolysis and smouldering processes during the cooling phase are likely to be the governing mechanisms driving delayed failure of timber columns. With respect to the second research question, the numerical modelling approach is shown to be inadequate for columns that fail deep into the cooling phase, primarily due to limitations in the representation of temperature-dependent thermal properties during this stage of fire exposure. These limitations become increasingly critical once the cooling phase plays a dominant role in the structural response.

7.2. GPC Surrogate Modelling Conclusions

This section presents the main conclusions related to the GPC surrogate modelling component of the study, in line with the second objective and the third and fourth sub-questions defined in Section 1.5. The conclusions focus on the applicability and reliability of GPC as a surrogate modelling framework

for approximating burnout resistance outcomes, and on its potential to support efficient parametric and probabilistic assessment based on numerical simulation results.

Overall, it is concluded that Gaussian Process Classification is a reliable and efficient surrogate modelling approach for capturing the decision boundary between failure and non-failure outcomes as predicted by deterministic numerical simulations. The GPC framework is shown to accurately approximate complex classification boundaries, even when multiple interacting variables are considered.

The results also demonstrate that increasing model dimensionality introduces additional complexity, which requires careful consideration of both the number of variables included and their representation within the surrogate model. These choices directly influence the assumptions underlying the learned decision boundary and therefore affect its interpretation.

A key strength of the proposed framework lies in the rapid prediction capability of the trained surrogate models. Once constructed, the GPC surrogate enables efficient evaluation of failure and non-failure outcomes across the considered ranges of the model input parameters. This enables systematic sensitivity analysis and probabilistic assessment of burnout resistance in the presence of inherent uncertainties, which would be computationally impractical to achieve using numerical simulations alone.

It is further concluded that the validity of the surrogate modelling results is directly linked to the accuracy of the underlying numerical simulations. As the numerical modelling approach is currently limited in its ability to reproduce delayed failure behaviour during the cooling phase, the present surrogate models cannot be used to draw definitive physical conclusions regarding delayed failure mechanisms. Nevertheless, the study demonstrates that, once the numerical modelling is sufficiently accurate, the GPC framework provides a powerful and computationally efficient means of assessing failure probability and burnout resistance across a high-dimensional parameter space.

7.3. Recommendations for Further Research

The study demonstrates the potential of combined numerical and surrogate modelling approaches for improving the assessment of delayed failure and burnout resistance of glulam timber columns. At the same time, the results highlight that the underlying physical processes and modelling strategies remain complex and require further investigation. As a result, the work opens up several directions for further research that could build on the presented framework and contribute to a more comprehensive understanding of delayed failure behaviour and burnout resistance. Accordingly, this section presents recommendations for further research on both the numerical modelling and the GPC surrogate modelling components of the study.

7.3.1. Numerical Modelling Recommendations

The most significant opportunities for further research are within improving the thermal modelling of structural timber under fire exposure. In particular, additional research is required to better represent the sustained elevated temperatures of timber columns during the cooling phase. This behaviour within the interior of the timber member is likely governed by continued smouldering and pyrolysis processes, which are not yet adequately captured by current thermal models.

A first step towards improving the thermal representation could be by the introduction of an additional heat source term in the governing heat conduction equation using an Arrhenius-based formulation as discussed in Hopkin [43]. Such an approach would allow reaction kinetics and delayed heat release to be incorporated more explicitly. Future studies should evaluate this modelling strategy against full-scale structural fire experiments on timber columns, rather than relying solely on small-scale material tests.

More broadly, further experimental research is recommended to strengthen calibration and validation of thermal models for natural-fire scenarios. In particular, additional full-scale compartment furnace tests on glulam columns would be valuable, in which the heating phase is varied through different Γ -factors and the cooling phase through different cooling rates. A furnace setting enables controlled compartment gas temperatures, which supports systematic calibration of the thermal response. These results would then allow natural fire exposures to be approximated more reliably using parametrically defined fire curves, and the resulting structural response to be predicted with higher confidence. In

addition, measuring column deflections throughout the fire and cooling phases is recommended, so that numerical mechanical analyses can be validated against global response data.

Further research is also recommended to reassess the temperature-dependent thermal properties of timber with specific attention to cooling-phase behaviour. Timber undergoes phase changes, especially during heating, and these transitions should be reflected more explicitly in the thermal material models. The introduction of an additional material state variable, as outlined in Appendix A through the proposed inverse PINN-based approach, offers a promising direction in this regard.

In addition, the temperature-dependent evolution of mechanical stiffness and its potential recovery during the cooling phase could be explored in further research. While stiffness recovery is not expected to be the primary driver of delayed failure, it may become relevant when assessing post-fire structural performance together with the phenomena of auto-extinguishment.

Finally, further research should address this process of auto-extinguishment in more depth. Burnout resistance of timber buildings does not depend solely on the behaviour of individual members, but also on whether the load-bearing structure contributes sufficient fuel to sustain the fire. Investigating fire–structure interaction at the level of structural assemblies and whole buildings is therefore essential for a comprehensive assessment of burnout resistance in timber construction.

7.3.2. GPC Surrogate Modelling Recommendations

Further research on the surrogate modelling framework should primarily focus on extending the dimensionality of the input space to better represent the full set of variables influencing delayed failure and burnout resistance. In the present study, the surrogate models are limited to six dimensions. Extending the framework towards a combined, higher-dimensional formulation would allow interacting design, fire, and material variables to be represented within a single model. A preliminary investigation into such an extension, with the use of a twelve-dimensional trial function, is provided in Appendix M.

Combining all relevant influencing variables into a single surrogate model would result in a twelve-dimensional input space. This model could include geometric variables (column width, column length, applied axial load, and load eccentricity), fire-related variables (heating duration, maximum temperature, and cooling rate), and material properties (density, stiffness, compressive strength, tensile strength, and moisture content).

In addition, statistical variability in material strength reduction could be incorporated by including a parameter such as the Weibull quantile, which governs temperature-dependent strength reduction behaviour. The compressive plastic strain parameter could also be included to account for variability in plastic deformation. Adding these variables would provide a more comprehensive description of burnout resistance, but would further increase the dimensionality of the surrogate model.

On the other hand, increasing the dimensionality of the surrogate model will require a substantially larger number of numerical simulations to reliably capture the decision boundary. This raises questions regarding the applicability of the standard Gaussian Process Classification formulation, as computational cost increases rapidly with the number of training samples. Further research is therefore recommended to investigate when alternative surrogate modelling approaches become more suitable. Promising directions include sparse Gaussian Processes, which offer improved scalability, or small neural network classifiers, which may be more efficient for large datasets.

Another potential direction lies in feature selection and feature engineering. Identifying which variables exert the strongest influence on delayed failure behaviour could allow less influential parameters to be excluded, thereby reducing dimensionality. Alternatively, physically motivated feature engineering could be used to combine multiple variables into lower-dimensional representations. While such approaches may lead to some loss of information, they could significantly improve learning efficiency. The feasibility and effectiveness of these strategies would need to be assessed on a case-by-case basis.


Finally, the surrogate modelling framework developed in this study is not limited to timber structures. Delayed failure has also been observed in other structural materials, such as concrete and steel. Extending the approach to these materials provides an opportunity to investigate burnout resistance across different structural designs and to further evaluate the general applicability of the applied surrogate modelling techniques in fire engineering.

References

- [1] *NEN-EN 1995-1-2, Eurocode 5 - Design of Timber Structures. Part 1-2: General – Structural Fire Design*. Tech. rep. Brussels: CEN, 2004.
- [2] T. Gernay and R. McNamee. *Burnout resistance based on the Duration of Heating Phase concept - Literature review and Roadmap (TVBB; No. 3260)*. Tech. rep. Lund: Division of Fire Safety Engineering, Lund University, 2024.
- [3] C. D. Oliver et al. “Carbon, Fossil Fuel, and Biodiversity Mitigation With Wood and Forests”. In: *Journal of Sustainable Forestry* 33.3 (2014), pp. 248–275. ISSN: 1540756X. DOI: 10.1080/10549811.2013.839386.
- [4] F. Richter et al. “Thermal Response of Timber Slabs Exposed to Travelling Fires and Traditional Design Fires”. In: *Fire Technology* 57.1 (2021), pp. 393–414. ISSN: 15728099. DOI: 10.1007/s10694-020-01000-1.
- [5] G. Churkina et al. “Buildings as a global carbon sink”. In: *Nature Sustainability* 3.4 (2020), pp. 269–276. ISSN: 23989629. DOI: 10.1038/s41893-019-0462-4.
- [6] D. Barber. “Tall Timber Buildings: What’s Next in Fire Safety ?” In: *Fire Technology* 51.6 (2015), pp. 1279–1284. ISSN: 15728099. DOI: 10.1007/s10694-015-0497-7.
- [7] M. Williams. *Nottingham university fire destroys new multimillion-pound chemistry building*. The Guardian. Sept. 2014. URL: <https://www.theguardian.com/uk-news/2014/sep/13/nottingham-university-fire-police-investigate-significant-blaze>.
- [8] BBC News. *University of Nottingham laboratory fire caused by electrical fault, says report*. BBC. Jan. 2015. URL: <https://www.bbc.com/news/uk-england-nottinghamshire-30751431>.
- [9] H. Shakib et al. “Evaluation of Plasco Building fire-induced progressive collapse”. In: *Structures* 28 (2020), pp. 205–224. ISSN: 23520124. DOI: 10.1016/j.istruc.2020.08.058.
- [10] S. S. Sunder et al. *Final Report on the Collapse of World Trade Center Building 7*. NIST NCSTAR 1A. Gaithersburg, MD: National Institute of Standards and Technology, 2008. DOI: 10.6028/NIST.NCSTAR.1A.
- [11] T. Gernay and J. M. Franssen. “A performance indicator for structures under natural fire”. In: *Engineering Structures* 100 (2015), pp. 94–103. ISSN: 1873-7323. DOI: 10.1016/j.engstruct.2015.06.005.
- [12] T. Gernay. “Fire resistance and burnout resistance of timber columns”. In: *Fire Safety Journal* 122 (2021). ISSN: 03797112. DOI: 10.1016/j.firesaf.2021.103350.
- [13] R. McNamee et al. “Enclosure fire dynamics with a cross-laminated timber ceiling”. In: *Fire and Materials* 45.7 (2021), pp. 847–857. ISSN: 10991018. DOI: 10.1002/fam.2904.
- [14] J. C. Mindeguia et al. “Thermo-mechanical behaviour of cross-laminated timber slabs under standard and natural fires”. In: *Fire and Materials* 45.7 (2021), pp. 866–884. ISSN: 10991018. DOI: 10.1002/fam.2938.
- [15] J. Su et al. *Fire Safety Challenges of Tall Wood Buildings-Phase 2: Task 2 & 3-Cross Laminated Timber Compartment Fire Tests*. Tech. rep. Quincy, MA: Fire Protection Research Foundation, 2018.
- [16] A. Law and R. M. Hadden. “Burnout means Burnout”. In: *SFPE Europe Magazine* 5 (2017), pp. 1–10.
- [17] F. Wiesner et al. “Structural capacity in fire of laminated timber elements in compartments with exposed timber surfaces”. In: *Engineering Structures* 179 (2019), pp. 284–295. ISSN: 18737323. DOI: 10.1016/j.engstruct.2018.10.084.
- [18] *Concept NTA 6125 Brandveiligheid massieve houtbouw voor consultatie*. Tech. rep. NEN, 2025.

- [19] A. I. Bartlett, R. M. Hadden, and L. A. Bisby. “A Review of Factors Affecting the Burning Behaviour of Wood for Application to Tall Timber Construction”. In: *Fire Technology* 55 (2019), pp. 1–49. ISSN: 15728099. DOI: 10.1007/s10694-018-0787-y.
- [20] *NEN-EN 1991-1-2, Eurocode 1 - Actions on Structures. Part 1-2: Actions on Structures Exposed to Fire*. Tech. rep. Brussels: CEN, 2024.
- [21] J. M. Franssen and T. Gernay. “Modeling structures in fire with SAFIR®: Theoretical background and capabilities”. In: *Journal of Structural Fire Engineering* 8.3 (2017), pp. 300–323. ISSN: 20402325. DOI: 10.1108/JSFE-07-2016-0010.
- [22] J.-M. Franssen. “SAFIR: A Thermal/Structural Program for Modeling Structures Under Fire”. In: *Engineering Journal* (2005), pp. 143–158.
- [23] J.-M. Franssen and T. Gernay. *User’s Manual for SAFIR (Version 2025): A Computer Program for Analysis of Structures Subjected to Fire — Part 1: General Considerations*. Tech. rep. Liège, Belgium: University of Liège, 2024.
- [24] J.-M. Franssen and T. Gernay. *User’s Manual for SAFIR (Version 2025): A Computer Program for Analysis of Structures Subjected to Fire — Part 2: Thermal Calculations*. Tech. rep. Liège, Belgium: University of Liège, 2024.
- [25] J.-M. Franssen and T. Gernay. *User’s Manual for SAFIR (Version 2025): A Computer Program for Analysis of Structures Subjected to Fire — Part 3: Structural Calculations*. Tech. rep. Liège, Belgium: University of Liège, 2024.
- [26] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006. ISBN: 026218253X.
- [27] R. Garnett. *Bayesian Optimization*. Cambridge University Press, 2023. ISBN: 9781108425780. DOI: 10.1017/9781108348973.
- [28] G. Zhao et al. “Efficient Active Learning for Gaussian Process Classification by Error Reduction”. In: *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*. 2021.
- [29] G. Zhao et al. “Efficient Active Learning for Gaussian Process Classification by Error Reduction: Appendix”. In: *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*. 2021.
- [30] F. Robert et al. “Natural fire tests on GLT columns including the cooling down phase”. In: *13th World Conference on Timber Engineering, WCTE 2023*. Vol. 3. 2023, pp. 1848–1854. ISBN: 9781713873297. DOI: 10.52202/069179-0244.
- [31] T. Gernay et al. “Experimental investigation of structural failure during the cooling phase of a fire: Timber columns”. In: *Fire and Materials* 47.4 (2023), pp. 445–460. ISSN: 10991018. DOI: 10.1002/fam.3110.
- [32] S. Renard et al. “Structural behavior of timber columns in wood crib compartment fire tests”. In: *Fire Safety Journal* 155 (2025). ISSN: 03797112. DOI: 10.1016/j.firesaf.2025.104413.
- [33] S. Renard et al. “Parametric experimental study on GLT columns stability during natural fire tests including the cooling phase”. In: *The 13th International Conference on Structures in Fire*. SiF 2024, 2024.
- [34] T. Gernay. “A method for measuring the sensitivity of building structural members to fire decay phases”. In: *Acta Polytechnica* 56.5 (2016), pp. 344–352. ISSN: 18052363. DOI: 10.14311/AP.2016.56.0344.
- [35] J. Stanke, E. Klement, and R. Rudolphi. *Das Brandverhalten von Holzstützen unter Druckbeanspruchung*. Tech. rep. 24. Berlin: Bundesanstalt für Materialforschung und -prüfung (BAM), 1973.
- [36] C. Douglas et al. “Structural Design of Timber Columns in Realistic Fires”. In: *Proceedings from the 14th World Conference on Timber Engineering: Advancing Timber for the Future Built Environment, WCTE 2025*. World Conference on Timber Engineering (WCTE), 2025, pp. 253–262. ISBN: 9798331320898. DOI: 10.52202/080513-0032.
- [37] T. Han and S. Tesfamariam. “Reliability analysis of timber columns under fire load using numerical models with equivalent section temperature”. In: *Engineering Structures* 324 (2025). ISSN: 18737323. DOI: 10.1016/j.engstruct.2024.119345.

- [38] J. D. van der Wulp. *GPC-Glulam-Timber-Columns-Fire-Resistance*. Version v1.0-thesis-submission. GitHub repository. 2026. URL: <https://github.com/jvanderWulp/GPC-Glulam-Timber-Columns-Fire-Resistance/releases/tag/v1.0-thesis-submission>.
- [39] *EN 14080:2013 – Timber Structures: Glued Laminated Timber and Glued Solid Timber – Requirements*. Tech. rep. Brussels: CEN, 2013.
- [40] S. Schilling et al. “Probabilistic description of the mechanical properties of glued laminated timber made from softwood”. In: ETH Zürich Research Collection, Institute of Structural Engineering, 2021. DOI: 10.3929/ethz-b-000505371.
- [41] *Binderholz - Glulam GLT*. Tech. rep. Binderholz GmbH - Glulam factory.
- [42] *Schilliger Holz - Glulam*. Tech. rep. SCHILLIGER HOLZ AG.
- [43] D. J. Hopkin. “Predicting the thermal response of timber structures in natural fires using computational ‘heat of hydration’ principles”. In: *Fire and Materials* 37.4 (2013), pp. 311–327. ISSN: 03080501. DOI: 10.1002/fam.2133.
- [44] D. J. Hopkin et al. “An effective thermal property framework for softwood in parametric design fires: Comparison of the Eurocode 5 parametric charring approach and advanced calculation models”. In: *Construction and Building Materials* 25.5 (2011), pp. 2584–2595. ISSN: 09500618. DOI: 10.1016/j.conbuildmat.2010.12.002.
- [45] *NEN-EN 1995-1-2, Eurocode 5 DRAFT - Design of Timber Structures. Part 1-2: Structural Fire Design*. Tech. rep. Brussels: CEN, 2023.
- [46] E. Garcia-Castillo, T. Gernay, and I. Paya-Zaforteza. “Probabilistic Models for Temperature-Dependent Compressive and Tensile Strengths of Timber”. In: *Journal of Structural Engineering* 149.2 (2023). ISSN: 0733-9445. DOI: 10.1061/jsendh.steng-11369.
- [47] H. J. Blaß. “Tragfähigkeit von Druckstäben aus Brettschichtholz unter Berücksichtigung streuender Einflußgrößen”. PhD thesis. Karlsruhe, Germany: University Fridericana Karlsruhe, 1987.
- [48] S. Šulc, V. Šmilauer, and F. Wald. “Thermal model for timber fire exposure with moving boundary”. In: *Materials* 14.3 (2021), pp. 1–10. ISSN: 19961944. DOI: 10.3390/ma14030574.
- [49] *NEN-EN 1995-1-1, Eurocode 5 - Design of timber structures - Part 1-1: General-Common rules and rules for buildings*. Tech. rep. Brussels: CEN, 2004.



Alternative Research Approaches

Several alternative modelling approaches were considered before selecting the approach adopted in this thesis. These approaches were explored to investigate ways of addressing the challenges in modelling delayed failure and burnout resistance of glued laminated timber columns under realistic fire exposure.

The first two approaches considered aimed to develop a model capable of efficiently predicting the structural response across multiple input dimensions. Such models would enable sensitivity analyses of key parameters and allow for rapid estimation of the duration of the heating phase leading to delayed failure, as the determination of burnout resistance for a given column configuration typically requires iterative numerical simulations.

In addition, a third approach was considered with a slightly different primary objective. This approach sought to exploit the key advantage of a Physics-Informed Neural Network (PINN), which is the ability to embed governing differential equations within the learning process. In this case, the focus was on describing or calibrating parameters governing the heat conduction equation, motivated by limitations identified in their current definition.

The alternative approaches discussed in this appendix were ultimately not adopted, but they provide important context for the modelling choices made in this thesis. Each approach is briefly outlined in the following sections, followed by a short evaluation in relation to the adopted numerical–surrogate modelling framework.

1. Parametric Physics-Informed Neural Network

The first approach considered was the use of a parametric Physics-Informed Neural Network (PINN) to solve the heat conduction equation directly. In this approach, the PINN functions as the solver for the partial differential equation governing the thermal analysis and could generalise across different geometries and fire curves. SAFIR would then be used as a mechanical post-processor and for validation. This approach has the advantage of retaining the physics of the problem within the learning process, thereby reducing the “black-box” nature typically associated with purely data-driven models. Furthermore, unlike conventional neural networks, a PINN does not require a large training dataset, as it is trained against the governing differential equations rather than pre-generated data.

Despite these advantages, the method introduces considerable complexity. The training of parametric PINNs is known to be challenging, as optimisation during the training process becomes increasingly difficult when additional dimensions are introduced. Moreover, because the PINN would only replace the thermal analysis, the mechanical analysis would still depend on traditional numerical methods. As a result, the potential computational gains are limited. Even once the thermal solver is accelerated, the overall model would only reduce the total computational demand by approximately half once trained. A schematic overview of this first approach is illustrated in Figure A.1.

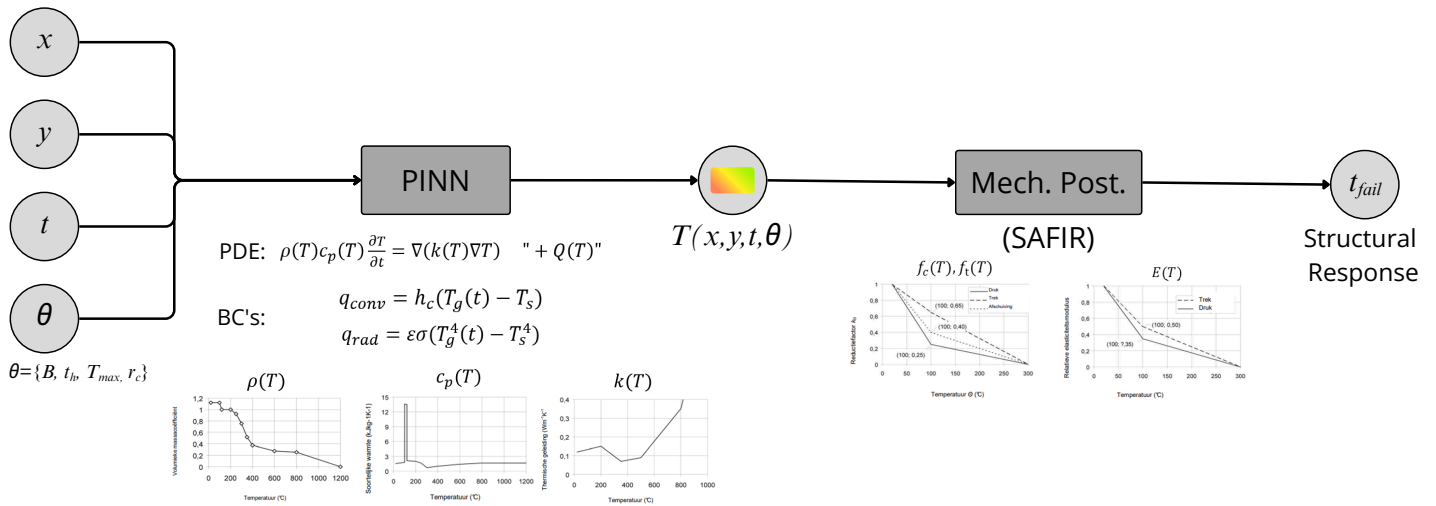


Figure A.1: Parametric PINN approach

2. Surrogate Neural Network on SAFIR Simulations

A second approach considered was the development of a surrogate neural network trained directly on SAFIR simulations. In this approach, the network would learn to map input parameters to the structural response without explicitly embedding physical constraints. The main advantage of this approach lies in its relative simplicity and its ability to deliver very fast predictions once the model is trained. However, neural networks generally require large amounts of training data, which in this case would require an extensive number of numerical simulations with SAFIR to generate a sufficiently broad dataset. In addition, this approach has a more pronounced “black-box” character, as input parameters are directly mapped to output responses. A schematic overview of this second approach is illustrated in Figure A.2.

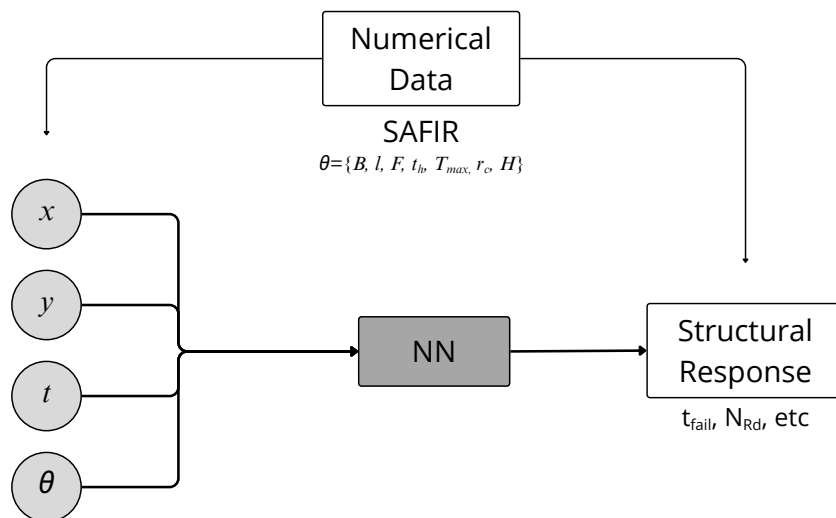


Figure A.2: Surrogate Neural Network approach

3. Inverse Physics-Informed Neural Network

A third alternative considered focused on the definition and calibration of the thermal properties of the governing heat conduction equation using an inverse Physics-Informed Neural Network (PINN). This approach aimed to improve the accuracy of the thermal property definitions for timber during cooling, which are currently calibrated primarily against ISO 834 heating conditions. By calibrating these properties using experimental data, the thermal behaviour of timber could be represented more accurately. Since the mechanical response depends directly on the thermal response, this would also lead to improved predictions of structural performance. In this way, the behaviour of timber could be captured

more realistically over the full duration of a fire, thereby addressing the identified limitation of the defined properties.

To enable such an approach, the governing physical equations would need to describe the behaviour of timber in a more fundamental manner, such that both heating and cooling phases are represented consistently. This could involve the introduction of a material state variable, allowing the thermal response of timber to be characterised as a function of its evolving state, rather than solely through temperature-dependent properties.

While scientifically valuable, this approach presents practical challenges in identifying the origin of potential modelling errors. Such errors could arise from limitations of the PINN formulation, the limited availability of suitable experimental data, or assumptions embedded in the governing defined physical equations. To help distinguish between these sources, one possible strategy would be to develop a separate numerical solver based on the same physical laws as the PINN. Artificial datasets with prescribed “synthetic” thermal properties could then be generated to assess whether the PINN is able to recover these dependencies. This would help clarify whether shortcomings are primarily data-driven or related to limitations of the PINN itself. However, this strategy would require the development of both a PINN and an additional numerical solver, which is not realistic within the scope and time frame of a master’s thesis.

Even with such measures, significant risks would remain. If the available experimental data prove insufficient, the resulting calibration may not lead to reliable improvements. Similarly, deficiencies in the underlying physical formulation would necessitate an iterative refinement of the governing equations, introducing further complexity. Finally, leading researchers, including Thomas Gernay, have indicated that the current Eurocode thermal properties already provide reasonable performance for many applications. As a result, this issue was considered less critical within the broader context of delayed failure and burnout resistance. A schematic overview of this third approach is illustrated in Figure A.3.

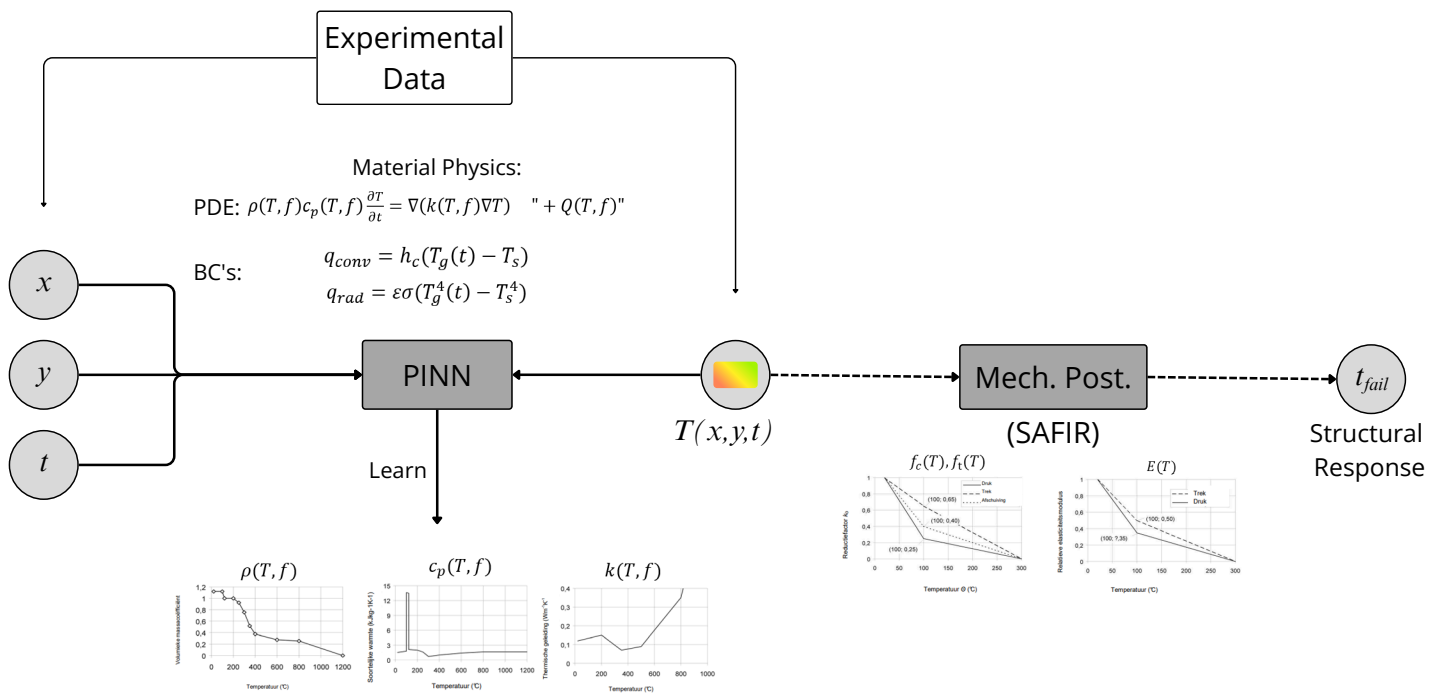


Figure A.3: Inverse PINN approach

Conclusion and Evaluation

The alternative approaches discussed in this appendix illustrate different strategies for addressing the challenges of modelling delayed failure and burnout resistance of glulam timber columns. In particular, the objective underlying the first two approaches to enable efficient exploration of a multi-dimensional parameter space was considered highly relevant, as it directly supports sensitivity analyses and systematic assessment of burnout resistance.

However, the effectiveness of these approaches differs primarily in terms of data requirements and practical feasibility. Fully data-driven neural network surrogates offer fast predictions once trained, but require large amounts of numerical data, making them less suitable when each data point is computationally expensive. Physics-informed neural networks reduce this dependency on data by embedding governing equations, but introduce significant training complexity and provide limited overall efficiency gains when only part of the modelling chain is replaced.

The third approach addresses a more fundamental modelling limitation related to thermal property definitions, but its implementation would require enough experimental data and sufficient validation accuracy, making it less suitable within the scope of this thesis.

Overall, these considerations highlight the need for a modelling framework that combines physical fidelity with computational efficiency, while remaining robust and feasible. This perspective motivated the selection of the adopted numerical–surrogate modelling approach with the use of Gaussian Processes, which offers a balanced compromise between the strengths and limitations of the considered alternatives.

B

Symmetry Mesh Modelling

To assess the computational efficiency gained by exploiting geometrical symmetry, an additional comparison was performed between full and quarter-symmetry models. The objective was to quantify the reduction in simulation time achieved when only a quarter of the cross-section is modelled, while maintaining equivalent accuracy. The analysis was based on the benchmark configuration of the glulam column tests 3 and 6 reported by Gernay et al. [31]. The full column was modelled as a square cross-section, while in the symmetry model only one quarter was analysed with symmetric boundary conditions along the central axes.

Simulation Inputs

Table B.1: Input parameters used for the symmetry modelling.

Category	Parameter	Symbol / Value	Unit
Material	Eccentricity	$e_0 = 0.02$	m
	Density	$\rho = 420$	kg/m ³
	Elastic modulus	$E = 11.5 \times 10^9$	Pa
	Poisson's ratio	$\nu = 0.3$	–
	Compressive strength	$f_c = 24$	MPa
	Tensile strength	$f_t = 19.2$	MPa
	Moisture content	$w = 12$	%
	Ultimate compressive strain	$\varepsilon_{cu} = 1.0$	–
Structural	Section width	$B = 0.28$	m
	Column length	$l = 3.65$	m
	Axial load	$F = 3.22 \times 10^5$	N
Fire	Ambient temperature	$T_0 = 20$	°C
	Heating duration	$t_h = 15$	min
	Maximum temperature	$T_{max} = 754.51$ ($\Gamma = 1$)	°C
	Cooling rate	$r_c = 10.4$	°C/min
Mesh	Structural mesh size	<code>mesh_size = 0.1825</code>	m
	Thermal mesh size	<code>fine_size = 0.005</code>	m

Simulation Results

For both modelling procedures, an identical failure time of 96.71 min was obtained. Both thermal and subsequent structural analyses were performed using identical mesh refinement and material parameters. The comparison shown in Table B.2 clearly demonstrates the computational benefit of the symmetry approach. For this configuration, the column fails during fire exposure, and the total simulation time is reduced from 3 min 57 s to 1 min 46 s, an improvement of more than a factor of two.

It should be noted that when the column does not reach failure, the mechanical simulation time typically


increases even more. Therefore, the symmetry approach offers a highly efficient strategy for thermal–structural simulations without compromising accuracy.

Table B.2: Comparison of total runtimes for full and symmetry-based models.

Model type	Thermal simulation	Structural simulation	Total time
Full cross-section	2 min 42 s	1 min 15 s	3 min 57 s
Quarter-symmetry	1 min 00 s	0 min 46 s	1 min 46 s

Conclusion

Using a symmetry-based model substantially reduces computation time while providing equivalent results. In the tested benchmark configuration, the symmetry model was more than twice as fast as the full model. For non-failing columns, where the structural phase continues for a longer duration, the relative efficiency gain is expected to be even greater. Consequently, symmetry modelling is chosen to be adopted throughout the simulations in this study.



Thermal Mesh Sensitivity Analysis

To determine an appropriate mesh density for the thermal simulations, a mesh sensitivity analysis was carried out. The objective was to identify a mesh size that provides sufficient numerical accuracy while keeping the computational effort reasonable. The analysis was performed using the configuration of the glulam column of tests 3 and 6 reported by Gernay et al. [31], which served as a representative benchmark for the thermo-structural response. Several meshes with varying element sizes were analysed, and the influence of mesh refinement on the predicted failure time and computational cost was evaluated. The results of this short study were subsequently applied in the thermal numerical simulations in the main report. See Subsection 4.2.2.

Simulation Inputs

The same fire, structural, and material input parameters were used for all simulations, and the thermal and structural modules WOODDEC5 were applied. The adopted input values are summarised in Table C.1.

Table C.1: Input parameters used for the thermal mesh sensitivity analyses.

Category	Parameter	Symbol / Value	Unit
Material	Eccentricity	$e_0 = 0.02$	m
	Density	$\rho = 420$	kg/m ³
	Elastic modulus	$E = 11.5 \times 10^9$	Pa
	Poisson's ratio	$\nu = 0.3$	–
	Compressive strength	$f_c = 24$	MPa
	Tensile strength	$f_t = 19.2$	MPa
	Moisture content	$w = 12$	%
	Ultimate compressive strain	$\varepsilon_{cu} = 1.0$	–
Structural	Section width	$B = 0.28$	m
	Column length	$l = 3.65$	m
	Axial load	$F = 3.22 \times 10^5$	N
Fire	Ambient temperature	$T_0 = 20$	°C
	Heating duration	$t_h = 15$	min
	Maximum temperature	$T_{max} = 754.51$ ($\Gamma = 1$)	°C
	Cooling rate	$r_c = 10.4$	°C/min
Mesh	Structural mesh size	$mesh_size = 0.1825$	m

Results of mesh sensitivity analysis

In order to quantify the results of the simulations the relative failure time error even as the runtime ratio are calculated. The corresponding formulas are presented in Equation C.1 and Equation C.2, respectively.

$$\text{Relative error: } \varepsilon(h) = \frac{|Q_{\text{ref}} - Q(h)|}{Q_{\text{ref}}} \times 100\% \quad (\text{C.1})$$

$$\text{Runtime ratio: } R_t(h) = \frac{t_r(h)}{t_{r,\text{ref}}} \quad (\text{C.2})$$

Let h denote the chosen fine mesh size, t_r the runtime, Q the failure time [min], Q_{ref} a chosen reference failure time, and $t_{r,\text{ref}}$ a chosen reference runtime.

Here, the reference failure time is chosen to be the one corresponding to the most refined mesh, 100.1 min. The baseline mesh size for the calculation of the runtime ratio is chosen at 5 mm.

Table C.2: Thermal mesh sensitivity results.
(reference $Q_{\text{ref}} = 100.1$ min; baseline for runtime ratio $h_b = 5$ mm)

fine_size h [m]	Failure time $Q(h)$ [min]	Relative error $\varepsilon(h)$ [%]	Runtime $t_r(h)$ [hh:mm:ss]	Runtime ratio $R_t(h)$ [-]
0.0175	65.4	34.67	00:00:30	0.27
0.00875	89.8	10.29	00:01:08	0.60
0.005	96.7	3.40	00:01:53	1.00
0.001	99.9	0.20	00:30:10	16.0
0.0005	100.1	0.00	02:07:02	67.4

Conclusion

The mesh sensitivity analysis demonstrates a clear convergence in the predicted failure time with progressive mesh refinement. The relative error decreases from 34.7% for the coarsest mesh ($h = 17.5$ mm) to 3.4% for $h = 5$ mm, and further to below 0.2% for finer meshes. However, this improvement in accuracy comes at a substantial computational cost, with the runtime increasing by approximately a factor of 16 when refining from 5 mm to 1 mm, and by more than 60 when refining to 0.5 mm. The mesh with $h = 5$ mm therefore provides a suitable balance between accuracy and efficiency, yielding results within 3.4% of the reference while maintaining practical runtimes for sequential simulations.

Based on these findings, values for the `fine_size` parameter are chosen around 5 mm for subsequent thermal analyses in the report. A formal Grid Convergence Index (GCI) assessment was not performed, as the results already provide sufficient evidence to justify the trade-off between computational time and numerical accuracy.

D

Numerical Restart Algorithm

In SAFIR, the thermal and mechanical analyses are performed sequentially, meaning that the mechanical response directly depends on the preceding thermal history. Since the objective of this study is to capture delayed failure behaviour, the thermal analysis must extend far enough to include the most critical temperature distribution of the cross-section. The end time is therefore defined as the moment when the central node begins to cool down, indicating that the heat wave has fully propagated through the section.

Because of the highly non-linear behaviour of the problem this critical moment cannot be precisely known in advance. Therefore, a restart-based approach was introduced to progressively extend the thermal analysis until cooling of the central node is detected. The restart functionality in SAFIR allows a simulation to continue from a previously saved thermal state without repeating the transient temperature response that occurred prior to that time step. After each run, the temperature field at the final time step is stored and then reloaded as the initial condition for the next run. By repeating this process, the analysis can be extended stepwise until the required thermal condition is reached.

Ideally, the use of restarts should provide results identical to those of a single continuous simulation of equal duration, as the state variables from the previous run are directly transferred to the next. In such an ideal case, restarts serve merely as a computational convenience, allowing long simulations to be divided into manageable intervals while maintaining numerical continuity. This approach therefore provides an efficient means to determine the appropriate end time of the thermal analysis without unnecessary computational effort.

Results

To assess the performance of the restart procedure, thermal simulations were carried out using the configurations of Tests 3 and 6 reported by Gernay et al. [31]. The thermal input parameters adopted for these simulations are summarised in Table D.1.

Table D.1: Input parameters used for the restart thermal simulations

Parameter	Symbol / Value	Unit
Density	$\rho = 420$	kg/m ³
Moisture content	$w = 12$	%
Section width	$B = 0.28$	m
Thermal mesh size	<code>fine_size = 0.005</code>	m

To investigate the behaviour of the restart algorithm, temperature measurements were extracted at three distinct positions within the cross-section, each exhibiting a different type of thermal response. The locations of these measurement points are illustrated in Figure D.1.





Figure D.1: Position of nodes in the cross-section for temperature development measurements. (One-quarter of the cross-section is visualised)


As fire exposure, the initial simulations were carried out using the parametric fire curve reported by Gernay for Tests 3 and 6. The corresponding input parameters are listed in Table D.2, and the resulting fire curve is shown in Figure D.2a.

Table D.2: Fire curve variables corresponding to tests 3 and 6 of Gernay et al. [31]

Parameter	Symbol / Value	Unit
Heating duration	$t_h = 15$	min
Maximum temperature	$T_{max} = 754.51$ ($\Gamma = 1$)	$^{\circ}\text{C}$
Ambient temperature	$T_0 = 20$	$^{\circ}\text{C}$
Cooling rate	$r_c = 10.4$	$^{\circ}\text{C}/\text{min}$



(a) Fire curve tests 3 and 6



(b) Comparison of single run against restart algorithm for central node

Figure D.2: Restart algorithm for central node


As shown in Figure D.2b, the temperature curve obtained using the restart procedure initially behaves as expected when viewed on its own. However, when compared with a continuous single-run simulation, clear deviations appear once the first restart is introduced. These deviations result in an error in

the predicted time at which the central node begins to cool down.


Figure D.3a provides an enlarged view of the error introduced by the restart procedure. From the continuous single-run simulation, it is known that cooling of the central node begins at 391.4 minutes. In the restart-based simulation, the initial end time was set to 250 minutes, and each extension was performed in increments of 60 minutes. Reaching the true cooling point would therefore require three restarts. However, only two restarts were executed in this simulation, and cooling was already detected.

To examine whether the observed error depends on the number of restarts and/or the end time of the initial run, additional simulations were performed. These results are shown in Figure D.3b. The blue curve corresponds to a restart-based simulation in which the end time of the initial run was increased to 290 minutes, while the restart increment was kept at 60 minutes. The purple curve represents a simulation where the initial end time remained at 250 minutes, but the restart increment was reduced to 30 minutes.

The results demonstrate that extending the end time of the initial run (blue curve) noticeably reduces the error in the predicted cooling time. In contrast, an increased amount of restarts leads to minor changes, with the predicted cooling point remaining close to the prior observed value.



(a) Error of restart algorithm



(b) Influence of extended end time of initial run

Figure D.3: Errors of restart algorithm for the central node

To examine the influence of the cooling phase on the restart algorithm, a node within the cross-section was selected that had already begun to cool by the time the first restart occurred. The corresponding results for node 316 are shown in Figure D.4. The figure shows that, once the first restart is introduced, a distinct kink appears in the cooling curve and the node cools more rapidly than in the continuous simulation. When comparing the second restart with the first, no clear repetition of this behaviour is observed.






Figure D.4: Behaviour of restart algorithm for cooling for node 316


To assess whether the deviations observed in the restart algorithm were specifically related to the cooling phase of the fire curve, the procedure was also evaluated using a standard ISO 834 fire curve shown in Figure D.5a. Two locations within the cross-section were examined, one where the temperature began to rise immediately and another where a delayed thermal response occurred. The corresponding temperature histories are presented in Figure D.5b and Figure D.5c. In both cases, the comparison with the continuous single-run simulation reveals negligible differences, which indicates that the restart procedure performs well when no cooling phase is present.



(a) ISO834 gas temperature



(b) Temperature development of node 7



(c) Temperature development of node 316

Figure D.5: Restart algorithm for ISO834 heating

Conclusions

The numerical investigation shows that the deviations between the restart-based simulations and the continuous single-run simulations originate almost entirely from the first restart. This initial restart leads to a noticeable underestimation of the cooling time, whereas subsequent restarts introduce only minor additional differences. Extending the end time of the initial run reduces this first-restart error, while an increased amount of restarts has only little effect on the error made.

The origin of the error lies in the loss of temperature history when a restart is introduced. Since SAFIR stores only the temperature field at the end of a run the irreversible thermal properties associated with the maximum temperature reached by each fibre are not transferred to the next run. As a result, fibres that have already begun to cool down experience a jump in their thermal properties once a restart occurs. This leads to a discontinuity in thermal diffusivity and produces the observed kink in Figure D.4 between the initial run and the first restart.

The effects of this discontinuity are the strongest at the first restart, since this is the moment where the influence of the peak of the fire curve is removed. By the time the first restart occurs, the outer elements have reached very high temperatures and have already cooled down considerably. Losing their temperature history at this moment results in very large changes in their thermal properties.

Once the transition between the first and second restart takes place, the peak temperatures in the history of each element have already been removed. Consequently, the proportional change in thermal properties becomes much smaller, and no comparable behaviour is observed in the subsequent restarts. This can be seen clearly in Figure D.4, where the transition from the first to the second restart no longer produces a noticeable kink.

The faster cooling observed between the initial run and the first restart in Figure D.4 is explained by the response of a more inner ring of fibres. The fibres that have reached maximum temperatures in the range of 100–120 °C carry a specific heat value corresponding to the spike in this temperature region. Once the first restart happens, these elements are assigned a specific heat outside this peak region, which leads to a substantial increase in thermal diffusivity and therefore a faster cooling rate.

Importantly, the error only occurs when the fire curve contains a cooling phase, since only then do mesh elements begin to cool before the restart. When the temperature continuously rises, as in the ISO 834 curve, no such change in thermal properties occurs and the restart procedure produces results consistent with a single continuous simulation. See Figure D.5

It can therefore be concluded that the restart algorithm is not reliable on its own for predicting the exact cooling time of the central node when a fire curve contains a cooling phase. Nevertheless, it remains a useful safety-net when applied iteratively and combined with a sufficiently accurate initial estimate. Starting from this estimate, the thermal analysis is first performed with the restart functionality, and a post-processing routine checks whether the central node has begun to cool. If cooling is not detected, the simulation is extended with restarts in one-hour increments until this criterion is met. An iterative scheme is then applied in which the end time identified from the previous restart-based run becomes the new initial guess for the following thermal simulation. This procedure is repeated until a full thermal analysis can be completed without any restarts. When the initial estimate lies reasonably close to the true cooling time, only a limited number of iterations are required, and the final continuous simulation provides the correct thermal response, ensuring a consistent basis for the subsequent mechanical analysis.

A hybrid approach, combining a reasonable initial guess with the iterative restart algorithm, remains computationally preferable to prescribing a highly conservative end time. Thermal simulations in SAFIR are relatively inexpensive, whereas the mechanical analysis is considerably more demanding. Selecting an overly long end time would therefore increase the total computational effort unnecessarily, particularly in cases where the column does not fail. By contrast, refining the end time through a small number of restart-based thermal iterations ensures that the cooling point is captured accurately while avoiding excessive mechanical simulation time. In short, the limited additional cost of a few extra thermal runs is far outweighed by the computational savings achieved in the mechanical analysis. For this reason, the iterative restart algorithm is accepted as an effective and efficient back-up approach.

E

Simplified Analytical Approach

A reasonable first estimate of the time at which the central node begins to cool is essential for the restart-based procedure used to determine the end time of the thermal analysis. The restart algorithm is only effective when the initial guess is sufficiently accurate; an estimate that is too low triggers unnecessary iterations, whereas an estimate that is too high results in excessive computational effort during the mechanical analysis. To avoid prescribing a highly conservative end time while still ensuring that the full thermal response is captured, a simplified analytical model was developed to provide a grounded initial prediction.

For this analytical approach to remain solvable in closed form, the governing heat-conduction problem must be kept linear. Any source of nonlinearity, such as temperature-dependent thermal properties or the radiation boundary condition term, makes the analytical solution unsolvable.

To include simplicity and maintain linearity the problem was simplified to a one-dimensional heat-conduction model with constant thermal properties. The model represents half of the cross-section, with a symmetry boundary at the centre and a prescribed temperature at the exposed surface. Although radiation could be incorporated through a linearised expression, and cases with convection or convection combined with linearised radiation were examined, these alternatives offered no meaningful improvement compared with a simple Dirichlet boundary condition. For this reason, the applied fire curve is introduced directly as a prescribed temperature at the boundary.

Furthermore, the fire curve was simplified to make the analysis easier. A piecewise linear curve was used, consisting of a linear heating phase, a linear cooling phase, and finally a constant ambient temperature.

With these simplifications, a substantial deviation in the predicted cooling time of the central node compared with the non-linear numerical model is to be expected. Nevertheless, the analytical model employs the same physical inputs as the non-linear SAFIR simulations and therefore still provides a meaningful basis for generating a first estimate of the cooling time.

Derivation of the simplified analytical model

Governing equation and boundary conditions

The one-dimensional heat-conduction problem is formulated on the interval $x \in [0, L]$, where $x = 0$ corresponds to the exposed face and $x = L = B/2$ represents the centre of the cross-section. Assuming constant thermal properties, the transient temperature field $T(x, t)$ satisfies

$$\rho c \frac{\partial T}{\partial t} = k \frac{\partial^2 T}{\partial x^2}, \quad 0 < x < L, \quad t > 0, \quad (\text{E.1})$$

which can be written as

$$\frac{\partial T}{\partial t} = \alpha \frac{\partial^2 T}{\partial x^2}, \quad \alpha = \frac{k}{\rho c}. \quad (\text{E.2})$$

Here, k is the thermal conductivity ($\text{W m}^{-1} \text{K}^{-1}$), ρ the density (kg m^{-3}), and c the specific heat capacity ($\text{J kg}^{-1} \text{K}^{-1}$). The parameter α is the thermal diffusivity ($\text{m}^2 \text{s}^{-1}$), which governs the rate at which

temperature propagates through the material.

The boundary conditions consist of a prescribed gas-side temperature $T_g(t)$ at the exposed face and a symmetry condition at the centre:

$$T(0, t) = T_g(t), \quad t > 0, \quad (\text{E.3})$$

$$\frac{\partial T}{\partial x}(L, t) = 0, \quad t > 0, \quad (\text{E.4})$$

while the initial condition is taken as a uniform temperature

$$T(x, 0) = T_0, \quad 0 \leq x \leq L. \quad (\text{E.5})$$

Piecewise linear fire curve

The applied fire curve is approximated by a piecewise linear function (see Figure E.1), consisting of a linear heating phase up to a maximum gas temperature T_{\max} , followed by a linear cooling phase and finally a return to the ambient temperature T_0 . The duration of the heating phase is denoted by t_h , while the cooling rate is represented by r_c , defined as the negative slope of the cooling branch. The transition time t_c marks the end of the cooling phase and the moment at which the gas temperature reaches T_0 again.

In terms of the time t , the fire curve is written as

$$T_g(t) = \begin{cases} T_0 + m_1 t, & 0 < t \leq t_h, \\ T_{\max} + m_2(t - t_h), & t_h < t \leq t_c, \\ T_0, & t > t_c, \end{cases} \quad (\text{E.6})$$

where

$$m_1 = \frac{T_{\max} - T_0}{t_h}, \quad m_2 = -r_c, \quad t_c = t_h + \frac{T_{\max} - T_0}{r_c}. \quad (\text{E.7})$$

Hence the time derivative of the gas temperature is piecewise constant:

$$\dot{T}_g(t) = \begin{cases} m_1, & 0 < t \leq t_h, \\ m_2, & t_h < t \leq t_c, \\ 0, & t > t_c. \end{cases} \quad (\text{E.8})$$




Figure E.1: Approximated fire curve used as thermal boundary condition in the simplified analytical approach.

Homogenisation of the boundary condition

To handle the non-homogeneous Dirichlet condition at $x = 0$, a transformed temperature field is introduced,

$$\theta(x, t) = T(x, t) - T_g(t). \quad (\text{E.9})$$

Substituting into the governing equation gives

$$\frac{\partial \theta}{\partial t} = \alpha \frac{\partial^2 \theta}{\partial x^2} - \dot{T}_g(t), \quad 0 < x < L, \quad t > 0, \quad (\text{E.10})$$

with boundary and initial conditions

$$\theta(0, t) = 0, \quad t > 0, \quad (\text{E.11})$$

$$\frac{\partial \theta}{\partial x}(L, t) = 0, \quad t > 0, \quad (\text{E.12})$$

$$\theta(x, 0) = T_0 - T_g(0) = 0, \quad 0 \leq x \leq L, \quad (\text{E.13})$$

since $T_g(0) = T_0$ in (E.6).

Eigenfunction expansion

The homogeneous spatial boundary-value problem associated with (E.10) is

$$X''(x) + \mu^2 X(x) = 0, \quad X(0) = 0, \quad X'(L) = 0. \quad (\text{E.14})$$

This yields the eigenvalues and eigenfunctions

$$\mu_n = \frac{(n + \frac{1}{2})\pi}{L}, \quad X_n(x) = \sin(\mu_n x), \quad n = 0, 1, 2, \dots \quad (\text{E.15})$$

which form a complete orthogonal basis on $[0, L]$. The solution $\theta(x, t)$ is expanded as

$$\theta(x, t) = \sum_{n=0}^{\infty} a_n(t) \sin(\mu_n x). \quad (\text{E.16})$$

Substituting (E.16) into (E.10) and using orthogonality gives, for each mode n ,

$$\frac{da_n}{dt} + \lambda_n a_n(t) = -b_n \dot{T}_g(t) \quad (\text{E.17})$$

with

$$b_n = \frac{2}{(n + \frac{1}{2})\pi}, \quad \lambda_n = \alpha \mu_n^2, \quad (\text{E.18})$$

and zero initial condition $a_n(0) = 0$ due to $\theta(x, 0) = 0$.

The solution of (E.17) is obtained by variation of constants,

$$a_n(t) = -b_n \int_0^t \exp[-\lambda_n(t - \tau)] \dot{T}_g(\tau) d\tau. \quad (\text{E.19})$$

Analytical solution

Applying an eigenfunction expansion to the homogenised governing equation and solving the resulting modal evolution equations yields an analytical representation of the transformed temperature field $\theta(x, t)$. Substituting the modal solution $a_n(t)$ from (E.19) into the series expansion (E.16) gives

$$\theta(x, t) = - \sum_{n=0}^{\infty} b_n \sin(\mu_n x) \int_0^t \exp[-\lambda_n(t - \tau)] \dot{T}_g(\tau) d\tau. \quad (\text{E.20})$$

The physical temperature field follows from the transformation $T(x, t) = \theta(x, t) + T_g(t)$, which leads to the compact form

$$T(x, t) = T_g(t) - \sum_{n=0}^{\infty} b_n \sin(\mu_n x) I_n(t), \quad (\text{E.21})$$

where the modal convolution integral is defined as

$$I_n(t) = \int_0^t \exp[-\lambda_n(t - \tau)] \dot{T}_g(\tau) d\tau. \quad (\text{E.22})$$

This expression provides the full analytical temperature response at any location $x \in [0, L]$ and any time $t > 0$ for the simplified model.

Temperature at the centre

The physical temperature at the centre of the cross-section follows from (E.21) evaluated at $x = L$. This gives $\sin(\mu_n L) = (-1)^n$, and one obtains

$$T(L, t) = T_g(t) - \sum_{n=0}^{\infty} b_n (-1)^n I_n(t). \quad (\text{E.23})$$

Because the derivative of the fire curve $\dot{T}_g(t)$ is piecewise constant, the convolution integral $I_n(t)$ from (E.22) takes closed-form expressions for each phase of the fire curve.

Heating phase ($0 < t \leq t_h$):

$$I_n(t) = \frac{m_1}{\lambda_n} (1 - e^{-\lambda_n t}). \quad (\text{E.24})$$

Cooling phase ($t_h < t \leq t_c$):

$$I_n(t) = \frac{m_1}{\lambda_n} (e^{-\lambda_n(t-t_h)} - e^{-\lambda_n t}) + \frac{m_2}{\lambda_n} (1 - e^{-\lambda_n(t-t_h)}). \quad (\text{E.25})$$

Ambient phase ($t > t_c$):

$$I_n(t) = \frac{m_1}{\lambda_n} (e^{-\lambda_n(t-t_h)} - e^{-\lambda_n t}) + \frac{m_2}{\lambda_n} (e^{-\lambda_n(t-t_c)} - e^{-\lambda_n(t-t_h)}). \quad (\text{E.26})$$

Combining the corresponding piecewise definitions of $T_g(t)$ and $I_n(t)$ with (E.23), provides the full analytical temperature response at the centre of the one-dimensional member.

Results

To assess the performance of the simplified analytical model, several comparisons were carried out against the non-linear thermal simulations performed in SAFIR. The input parameters used in these test cases are listed in Table E.1 and Table E.2. Test 1 corresponds to the configuration of tests 3 and 6 reported by Gernay et al. [31].

For each test, the response of the simplified analytical model was compared with the non-linear SAFIR simulation in terms of the time at which the central node begins to cool, and the maximum temperature attained at the centre. These results are summarised in Table E.3 and Table E.4.

All SAFIR simulations were carried out using a convection coefficient of $35 \text{ W m}^{-2} \text{ K}^{-1}$ during heating and $4 \text{ W m}^{-2} \text{ K}^{-1}$ during cooling. The value of the thermal emissivity was taken as 0.8.

For the analytical model, the specific heat capacity and thermal conductivity were taken as their Eurocode ambient-temperature values. A specific heat capacity of $1530 \text{ J kg}^{-1} \text{ K}^{-1}$ and a thermal conductivity of $0.12 \text{ W m}^{-1} \text{ K}^{-1}$ were used throughout the calculations. The analytical solution was obtained using 60 modes.

Table E.1: Input parameters used for the comparison between the simplified analytical model and the non-linear SAFIR thermal analysis (Tests 1–4).

Variable	Symbol [Unit]	Test 1	Test 2	Test 3	Test 4
Inputs					
Ambient temperature	T_0 [°C]	20	20	20	20
Heating duration	t_h [min]	15	25	25	25
Cooling rate	r_c [°C/min]	10.4	10.4	6.0	6
Maximum gas temperature	T_{max} [°C]	754.51 ($\Gamma = 1$)	816.68 ($\Gamma = 1$)	816.68 ($\Gamma = 1$)	1150 ($\Gamma = 9.49$)
Section width	B [m]	0.28	0.40	0.32	0.5
Thermal mesh size	<code>fine_size</code> [m]	0.005	0.005	0.005	0.0039
Density	ρ [kg/m ³]	420	420	420	420

Table E.2: Input parameters used for the comparison between the simplified analytical model and the non-linear SAFIR thermal analysis (Tests 5–8).

Variable	Symbol [Unit]	Test 5	Test 6	Test 7	Test 8
Inputs					
Ambient temperature	T_0 [°C]	20	20	20	20
Heating duration	t_h [min]	5	5	35	35
Cooling rate	r_c [°C/min]	6	15	5	5
Maximum gas temperature	T_{max} [°C]	1150 ($\Gamma = 47.45$)	1150 ($\Gamma = 47.45$)	750 ($\Gamma = 0.41$)	750 ($\Gamma = 0.41$)
Section width	B [m]	0.2	0.32	0.32	0.5
Thermal mesh size	<code>fine_size</code> [m]	0.00625	0.005	0.005	0.0039
Density	ρ [kg/m ³]	420	420	420	420

Table E.3: Comparison of SAFIR and analytical model results (Tests 1–4).

Result	Test 1		Test 2		Test 3		Test 4	
	SAFIR	Analytical	SAFIR	Analytical	SAFIR	Analytical	SAFIR	Analytical
Cooling time of central node [min]	391.4	327.3	702.35	637.6	496.38	446.0	830.06	1017.4
Maximum temperature at centre [°C]	70.18	53.2	55.96	52.5	83.68	70.6	75.99	59.9
Ratio cooling time (SAFIR/analytical) [-]	1.20		1.10		1.11		0.82	

Table E.4: Comparison of SAFIR and analytical model results (Tests 5–8).

Result	Test 5		Test 6		Test 7		Test 8	
	SAFIR	Analytical	SAFIR	Analytical	SAFIR	Analytical	SAFIR	Analytical
Cooling time of central node [min]	89.41	243.4	372.32	409.7	525.38	460.0	1140.23	1007.4
Maximum temperature at centre [°C]	687.76	230.1	73.41	56.7	85.85	73.0	57.64	41.9
Ratio cooling time (SAFIR/analytical) [-]	0.37		0.91		1.14		1.13	

Conclusions

When comparing the results of the simplified analytical model to the non-linear SAFIR simulations, clear deviations are observed. The most likely origin of these deviations is expected to be the influence of the temperature-dependent thermal properties used in the numerical simulations.

For the predicted time at which cooling begins, the largest observed ratio corresponds to a factor of 1.20. When examining the cooling times more closely, a clear trend appears between cases with slower heating phases (low Γ values) and those with faster heating phases (high Γ values). For the slower heating phases, the analytical approach tends to underpredict the SAFIR cooling times, whereas for the fast heating phases it tends to overpredict. When interpreting these differences, it should be kept in mind that the non-linear SAFIR simulations are not an absolute ground truth either, particularly since their temperature-dependent thermal properties are calibrated solely for heating scenarios with a Gamma-factor equal to one (ISO 834).

In the end, to arrive at a sufficiently conservative initial guess for the end time of the thermal analysis, a safety factor of 1.3 is applied to the analytical solution. Although one could consider applying a different factor for low- Γ and high- Γ fires, the current dataset is considered too limited to justify such a distinction. For this reason, a uniform conservative factor of 1.3 is adopted.

Additional analytical cases including convection only, and convection combined with linearised radiation, were also investigated. These variations did not provide any significant improvement in predicting the onset of cooling when compared to the baseline analytical formulation. Therefore, the simpler analytical model, combined with the uniform correction factor, remains the preferred approach for establishing the initial estimate of the end time.

F


Charring Depth Validation

In order to validate the applicability of the cumulative-temperature equation (A.4.3.2(2)) from the draft of Eurocode 5 [45], this appendix examines the charring behaviour observed in a set of natural fire experiments conducted by Gernay et al. [30, 32, 33]. The objective is to compare the experimentally measured charring front, defined as the position of the 300 °C isotherm, with the charring-depth predictions obtained from the natural fire curves recorded during these tests.


The recorded natural fire curves are used as direct input to compute the charring depth over time using the cumulative-temperature formulation given in Subsection 4.2.2. The resulting predictions are then evaluated against the experimental measurements to assess whether the model accurately captures both the progression and the eventual stabilisation of the 300 °C isotherm.

Results

Below the results for a set of natural fire experiments conducted by Gernay et al. are presented. The plots of the charring-depth development are shown only up to the time of structural failure of each test. Beyond this point, the thermocouple measurements become unreliable due to the collapse of the column and the associated disturbance of the instrumentation. For the experimental measurements eight thermocouples positioned at depths ranging from 0 to 80 mm, spaced at 10 mm intervals, were used. These measurements provide the experimentally observed location of the 300 °C isotherm.




(a) Natural fire curve of Test 9




(b) Charring depth development of Test 9


Figure F.1: Fire curve and calculated versus measured charring depths for Test 9.




(a) Natural fire curve of Test 10



(b) Charring depth development of Test 10


Figure F.2: Fire curve and calculated versus measured charring depths for Test 10.

(a) Natural fire curve of Test 12




(b) Charring depth development of Test 12


Figure F.3: Fire curve and calculated versus measured charring depths for Test 12.



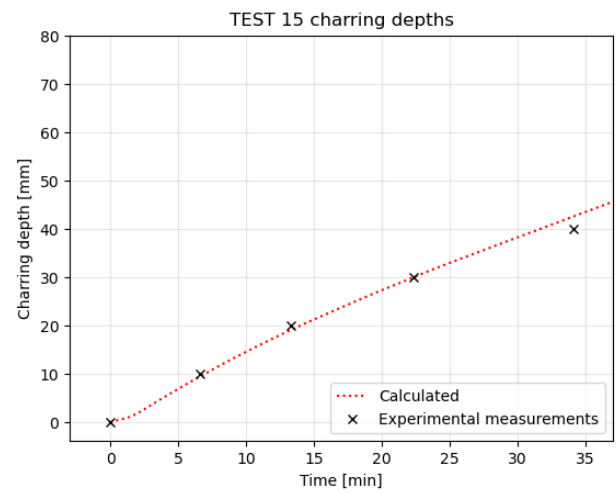
(a) Natural fire curve of Test 14



(b) Charring depth development of Test 14

Figure F.4: Fire curve and calculated versus measured charring depths for Test 14.

(a) Natural fire curve of Test 15



(b) Charring depth development of Test 15

Figure F.5: Fire curve and calculated versus measured charring depths for Test 15.





Figure F.6: Fire curve and calculated versus measured charring depths for Test 17.






Figure F.7: Fire curves and calculated versus measured charring depths for Test 18.

Conclusions

As an overall conclusion, the comparison between the experimentally measured charring depths and the predictions obtained from the cumulative-temperature model shows a generally good agreement. When examining the results in more detail, the cumulative formulation slightly overpredicts the measured experimental data during the cooling phase. However, it should be noted that the 300 °C isotherm is only an approximate indicator of the true charring front. During the decay phase of the fire, limited ongoing pyrolysis may allow the actual charring depth to advance marginally beyond the strict position of this isotherm. Nevertheless, the experimental data suggest that this effect is small and does not significantly change the overall trend.

The behaviour is most clearly illustrated in Test 18. In this test, the 50 mm thermocouple never exceeds 300 °C, but it reaches a maximum temperature of 298.41 °C at around 79.5 minutes, as shown in Figure F.7c. When this value is considered in the reasoning, it becomes evident that the charring front stabilises once the compartment temperature falls below 300 °C, which occurs at approximately 83 minutes. This supports the decision to cap the predicted charring depth at the moment the fire curve drops below the charring threshold.

Overall, the observations from the natural fire experiments indicate that the cumulative-temperature model provides a sufficiently accurate and conservative representation of the charring behaviour for the purposes of this study.

G

Moment Initialisation

This appendix compares two equivalent methods for introducing an initial bending effect associated with load eccentricity in the column analysis. The comparison is included to verify that the chosen implementation of the initial moment does not influence the predicted structural response or failure time. The analysis was performed using the configuration of the glulam column tests 3 and 6 reported by Gernay et al. [31], which served as a representative benchmark for the thermo-structural response.

Results

Two loading strategies were analysed for an identical column configuration. In the first approach, an initial first-order bending moment $M_0 = F \cdot e_0$ was applied at both the top and bottom of the column, resulting in a constant moment distribution along the member. In the second approach, the axial force was applied eccentrically at the top cross-section, allowing the nonlinear mechanical solver to generate the corresponding bending effects internally.

The analyses were performed using identical material, structural, fire, and mesh parameters, which are summarised in Table G.1. The thermal and structural modules WOODDEC5 were used for the analyses.

Table G.1: Input parameters used for the moment initialisation analyses.

Category	Parameter	Symbol / Value	Unit
Material	Eccentricity	$e_0 = 0.02$	m
	Density	$\rho = 420$	kg/m ³
	Elastic modulus	$E = 11.5 \times 10^9$	Pa
	Poisson's ratio	$\nu = 0.3$	–
	Compressive strength	$f_c = 24$	MPa
	Tensile strength	$f_t = 19.2$	MPa
	Moisture content	$w = 12$	%
	Ultimate compressive strain	$\varepsilon_{cu} = 1.0$	–
Structural	Section width	$B = 0.28$	m
	Column length	$l = 3.65$	m
	Axial load	$F = 3.22 \times 10^5$	N
Fire	Ambient temperature	$T_0 = 20$	°C
	Heating duration	$t_h = 15$	min
	Maximum temperature	$T_{max} = 754.51$ ($\Gamma = 1$)	°C
	Cooling rate	$r_c = 10.4$	°C/min
Mesh	Structural mesh size	<code>mesh_size = 0.1825</code>	m
	Thermal mesh size	<code>fine_size = 0.005</code>	m


The resulting failure times for both approaches are summarised in Table G.2.

Table G.2: Comparison of moment initialisation approaches

Loading approach	Nodeline position [m]	Failure time [min]
End moments applied ($M_0 = F \cdot e_0$)	0.14	96.70
Eccentric axial load applied	0.12	96.58

Conclusions

The comparison shows that applying the initial bending moment explicitly at the column ends or introducing it implicitly through eccentric axial loading leads to almost identical structural responses for the considered configuration. No comparison was made with an initially deformed column shape, such as a sinusoidal curve. However, no significant influence is expected, as failure is governed by the critical mid-height cross-section. For simplicity and clarity, the approach using explicit initial end moments is therefore adopted in the main analyses.



Mechanical Mesh Sensitivity Analysis

To assess the influence of the mechanical discretisation on the predicted structural response, a mechanical mesh sensitivity analysis was performed. The objective was to identify a suitable structural element length that provides sufficient numerical accuracy while maintaining reasonable computational efficiency. The analysis was carried out using the configuration of the glulam column tests 3 and 6 reported by Gernay et al. [31], which served as a representative benchmark for the thermo-mechanical response. The results of this study were subsequently used in the mechanical analyses presented in the main report.

Simulation Inputs

All simulations were performed using identical fire, material, and structural input parameters. The thermal and structural modules WOODDEC5 were applied. The adopted input values are summarised in Table H.1.

Table H.1: Input parameters used for the mechanical mesh sensitivity analyses.

Category	Parameter	Symbol / Value	Unit
Material	Eccentricity	$e_0 = 0.02$	m
	Density	$\rho = 420$	kg/m ³
	Elastic modulus	$E = 11.5 \times 10^9$	Pa
	Poisson's ratio	$\nu = 0.3$	–
	Compressive strength	$f_c = 24$	MPa
	Tensile strength	$f_t = 19.2$	MPa
	Moisture content	$w = 12$	%
	Ultimate compressive strain	$\varepsilon_{cu} = 1.0$	–
Structural	Section width	$B = 0.28$	m
	Column length	$l = 3.65$	m
	Axial load	$F = 3.22 \times 10^5$	N
Fire	Ambient temperature	$T_0 = 20$	°C
	Heating duration	$t_h = 15$	min
	Maximum temperature	$T_{max} = 754.51$ ($\Gamma = 1$)	°C
	Cooling rate	$r_c = 10.4$	°C/min
Mesh	Thermal mesh size	<code>fine_size = 0.005</code>	m

Results of Mechanical Mesh Sensitivity Analysis

The mechanical mesh was varied by changing the number of structural elements along the column length, resulting in different element lengths. The influence of mesh refinement on the predicted failure time and computational cost was evaluated. The relative failure time error was computed with respect to the converged failure time obtained for the finest discretisation. For the runtime ratio, a baseline

case with four structural elements was selected as reference, as this mesh lies in the mid-range of the tested discretisations.

The relative failure time error and runtime ratio are evaluated using the same definitions as introduced in the thermal mesh sensitivity analysis (see Equation C.1 and Equation C.2).

Table H.2: Mechanical mesh sensitivity results.
(reference failure time $Q_{\text{ref}} = 96.71$ min; baseline runtime at 4 elements)

Elements	Element length [m]	Failure time [min]	Relative error [%]	Runtime [hh:mm:ss]	Runtime ratio [-]
1	3.6500	119.88	23.9	00:00:50	0.85
2	1.8250	96.95	0.25	00:01:05	1.10
4	0.9125	96.74	0.03	00:00:59	1.00
8	0.4563	96.73	0.02	00:01:11	1.20
16	0.2281	96.71	0.00	00:01:34	1.59
32	0.1141	96.71	0.00	00:02:06	2.14

Conclusion

The mechanical mesh sensitivity analysis shows rapid convergence of the predicted failure time with increasing mesh refinement. A very coarse discretisation with a single element leads to a significant overestimation of the failure time, while meshes with four or more elements yield nearly identical results. Further refinement beyond an element length of approximately 0.25 metres does not result in meaningful improvements in accuracy.

In contrast to the thermal analysis, the additional computational cost associated with refining the mechanical mesh is relatively modest. As a result, adopting a converged structural discretisation does not significantly affect the overall runtime. Based on these observations, a structural element length of approximately 0.25 metres is selected for the mechanical analyses in this thesis, providing sufficient accuracy while maintaining computational efficiency.

Plastic Compression Strain

This appendix investigates the influence of the prescribed compressive plastic strain limit on the predicted structural failure time of glulam columns exposed to fire. In SAFIR, timber is modelled using a temperature-dependent constitutive law that is elastic–brittle in tension and elastic–plastic in compression, as illustrated in Figure 4.9. The extent of plastic deformation in compression is governed by the limiting compressive plastic strain parameter ε_{cu} , which is not uniquely defined in existing standards or experimental calibration.

To assess the importance and sensitivity of this parameter, four experimental fire tests on glulam columns reported by Stanke et al. [35] were selected. These tests were chosen to represent a range of column widths, material properties, and load levels, allowing the influence of compressive plasticity to be evaluated across different structural configurations. All tests are subjected to standard fire exposure.

Results

All simulations were carried out using the thermal module WOODDEC5. The column length of the tests are equal to $l = 3.65$ m for all cases. Boundary conditions correspond to a hinged–fixed configuration, with both translational degrees of freedom restrained at the bottom, and transverse translation and rotation restrained at the top.

An initial eccentricity was introduced to trigger bending, following the approach adopted by Gernay et al. [12]. The eccentricity was taken as $e_0 = l/400$, resulting in $e_0 = 0.009125$ m, consistent with Eurocode recommendations. Material properties and applied loads were taken directly from the experimental test descriptions. For each test, simulations were performed using two values of the compressive plastic strain limit, $\varepsilon_{cu} = 0.0$ and $\varepsilon_{cu} = 1.0$. Two mechanical strength reduction models were considered, WOODDEC5 and WOODPRBWE, for which a Weibull quantile of 0.5 was applied in the case of WOODPRBWE. The thermal and mechanical meshes were defined based on the outcomes of the thermal and mechanical mesh sensitivity analyses.

Test H27A

The test-specific material, geometric, and loading parameters for test H27A are listed in Table I.1. The corresponding simulation results obtained for different strength reduction models and compressive plastic strain limits are summarised in Table I.2.

Table I.1: Test-specific input parameters for test H27A of Stanke et al. [35].

Parameter	Symbol	Value
Section width	B	0.27 m
Density	ρ	420 kg/m ³
Elastic modulus	E	13.1 GPa
Compressive strength	f_c	42.9 MPa
Moisture content	w	12.6 %
Axial load	F	538.366 kN

Table I.2: Simulation results for the H27A test of Stanke et al. [35].
Experimental failure time: 59 min

Simulation	ε_{cu}	Strength reduction curve	Failure time (min)
1	0.0	WOODEC5	52.55
2	1.0	WOODEC5	61.42
3	0.0	WOODPRBWE	64.99
4	1.0	WOODPRBWE	66.18

Test R20A

The test-specific material, geometric, and loading parameters for test R20A are listed in Table I.3. The corresponding simulation results obtained for different strength reduction models and compressive plastic strain limits are summarised in Table I.4.

Table I.3: Test-specific input parameters for test R20A of Stanke et al. [35].

Parameter	Symbol	Value
Section width	B	0.20 m
Density	ρ	400 kg/m ³
Elastic modulus	E	11.032 GPa
Compressive strength	f_c	40.9 MPa
Moisture content	w	8.1 %
Axial load	F	251.056 kN

Table I.4: Simulation results for the R20A test of Stanke et al. [35].
Experimental failure time: 34 min

Simulation	ε_{cu}	Strength reduction curve	Failure time (min)
1	0.0	WOODEC5	25.77
2	1.0	WOODEC5	28.56
3	0.0	WOODPRBWE	30.49
4	1.0	WOODPRBWE	30.91

Test R15A

The test-specific material, geometric, and loading parameters for test R15A are listed in Table I.5. The corresponding simulation results obtained for different strength reduction models and compressive plastic strain limits are summarised in Table I.6.

Table I.5: Test-specific input parameters for test R15A of Stanke et al. [35].

Parameter	Symbol	Value
Section width	B	0.15 m
Density	ρ	380 kg/m ³
Elastic modulus	E	12.411 GPa
Compressive strength	f_c	41.3 MPa
Moisture content	w	10.2 %
Axial load	F	106.89 kN

Table I.6: Simulation results for the R15A test of Stanke et al. [35].
Experimental failure time: 26 min

Simulation	ε_{cu}	Strength reduction curve	Failure time (min)
1	0.0	WOODEC5	16.85
2	1.0	WOODEC5	17.91
3	0.0	WOODPRBWE	18.88
4	1.0	WOODPRBWE	19.12

Test H40

The test-specific material, geometric, and loading parameters for test H40 are listed in Table I.7. The corresponding simulation results obtained for different strength reduction models and compressive plastic strain limits are summarised in Table I.8.

Table I.7: Test-specific input parameters for test H40 of Stanke et al. [35].

Parameter	Symbol	Value
Section width	B	0.40 m
Density	ρ	410 kg/m ³
Elastic modulus	E	9.653 GPa
Compressive strength	f_c	39.0 MPa
Moisture content	w	9.4 %
Axial load	F	1372.937 kN


Table I.8: Simulation results for the H40 test of Stanke et al. [35].
Experimental failure time: 114 min

Simulation	ε_{cu}	Strength reduction curve	Failure time (min)
1	0.0	WOODEC5	83.17
2	1.0	WOODEC5	101.00
3	0.0	WOODPRBWE	105.06
4	1.0	WOODPRBWE	109.22

Conclusions

The sensitivity study demonstrates that both the selected strength reduction curve and the prescribed compressive plastic strain limit influence the predicted failure time of glulam columns under fire exposure. Allowing compressive plastic deformation generally leads to increased failure times, although the magnitude of this effect varies between tests and material models.

Importantly, for all investigated configurations, the predicted failure times obtained with $\varepsilon_{cu} = 0.0$ and $\varepsilon_{cu} = 1.0$ remain within the range of the experimental observations reported by Stanke et al. No pronounced outlier behaviour is observed for either limiting case. These results indicate that the appropriate choice of the compressive plastic strain limit is problem-dependent and cannot be uniquely defined. On this basis, both limiting cases are retained in the analyses presented in this thesis.



GPC AL: 2D Trial Functions

This appendix presents the results of a series of two-dimensional synthetic trial-functions used to develop the GPC surrogate modelling framework and its associated active learning strategy. Two-dimensional problems were selected as the initial test case, as they allow full visualisation of the decision boundary, training data, and model predictions, while providing a known ground truth against which performance can be assessed.

Multiple two-dimensional trial functions were considered to evaluate the robustness of the active learning algorithm under different boundary characteristics. The considered cases include a nonlinear boundary, a parabolic boundary, a square-root boundary, and a sinusoidal boundary. These functions exhibit varying levels of nonlinearity and smoothness, allowing assessment of training-point selection, hyperparameter optimisation, and the quality of the resulting surrogate prediction model. In addition, effective algorithmic settings and convergence criteria for the two-dimensional problems are identified. The results provide insight into the number of model evaluations required and form the basis for extending the framework to higher-dimensional input spaces.

Results

All monitoring metrics described in Subsection 4.3.2 are evaluated for the two-dimensional trial functions. Hyperparameter optimisation and convergence are performed following the standard procedure outlined in the main text, unless stated otherwise. The monitoring metrics are evaluated using a fixed Sobol sample set consisting of 20,000 points.

The final surrogate model is evaluated using a global misclassification error, computed by comparing the model predictions with the known ground truth. This error is estimated using a Sobol sampling set consisting of 200,000 points distributed over the input domain.

The active learning configuration used for these two dimensional cases includes 200 gradient-based optimisation steps per iteration with a learning rate of 0.001 for the stochastic gradient descent procedure used in the SMOCU point optimisation. For each iteration, an initial random Sobol sampling set of 20,000 points is generated, from which a retaining fraction of 0.01 is applied, resulting in 200 candidate points for SMOCU evaluation. The one-dimensional SMOCU integral is approximated using 2,000 truncated Gaussian points.

Nonlinear Boundary

The first two-dimensional trial function considered features a nonlinear decision boundary in the normalised input space $(x_1, x_2) \in [0, 1]^2$. The boundary is defined as

$$f(x_1) = 0.2 + 0.6 x_1^2,$$

with class labels assigned according to

$$y(x_1, x_2) = \begin{cases} 0, & \text{if } x_2 > f(x_1), \\ 1, & \text{if } x_2 \leq f(x_1). \end{cases}$$

The active learning process is considered converged after 69 iterations, at which point the boundary spread metric drops below the threshold value of 0.01. At convergence, the global misclassification error is found to be 0.143%. The final GPC prediction field obtained after convergence is shown in Figure J.1.




Figure J.1: GPC prediction field for the two-dimensional nonlinear boundary trial function.

To assess the convergence behaviour of the surrogate model and the active learning strategy, several monitoring metrics are tracked throughout the learning process. These metrics include the acquisition values, changes in the estimated decision boundary, the boundary spread, the log-likelihood of the model, and the misclassification ratio. The evolution of these quantities for the nonlinear boundary case is presented in Figure J.2.

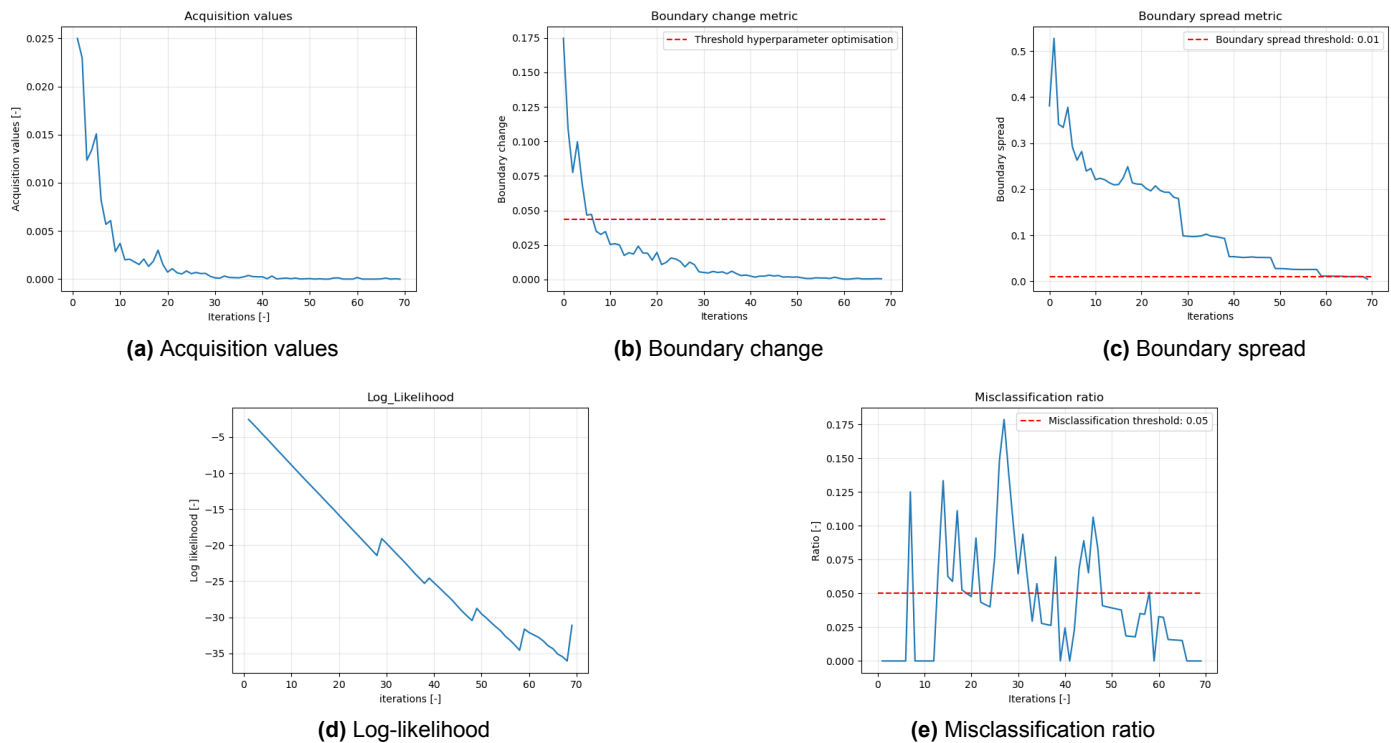


Figure J.2: Monitoring metrics for the two-dimensional nonlinear boundary trial function.

Parabola Boundary

The second two-dimensional trial function considered features a parabolic decision boundary in the normalised input space $(x_1, x_2) \in [0, 1]^2$. The boundary is defined as

$$f(x_1) = 0.2 + 0.6 \cdot 7(x_1 - 0.5)^2.$$

Class labels are assigned according to

$$y(x_1, x_2) = \begin{cases} 0, & \text{if } x_2 > f(x_1), \\ 1, & \text{if } x_2 \leq f(x_1). \end{cases}$$

The active learning process is considered converged after 92 iterations, at which point the boundary spread metric drops below the threshold value of 0.01. At convergence, the global misclassification error is found to be 0.098%. The final GPC prediction field obtained after convergence is shown in Figure J.3.




Figure J.3: GPC prediction field for the two-dimensional parabola boundary trial function.

As in the previous case, several monitoring metrics are tracked throughout the active learning process to assess convergence behaviour and surrogate model performance. These metrics include the acquisition values, changes in the estimated decision boundary, the boundary spread, the log-likelihood, and the misclassification ratio. The evolution of these quantities for the parabola boundary case is presented in Figure J.4.

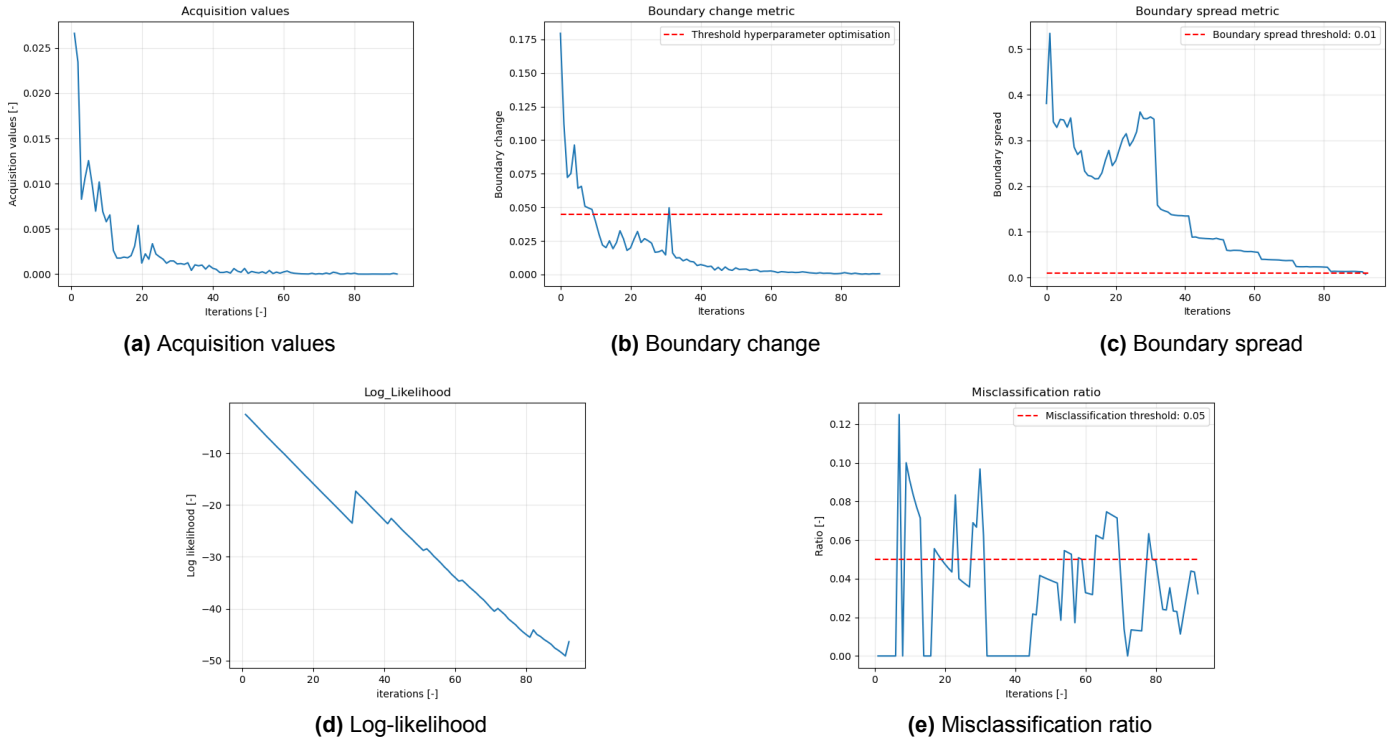


Figure J.4: Monitoring metrics for the two-dimensional parabola boundary trial function.

Square Root Boundary

The third two-dimensional trial function considered features a square-root-shaped decision boundary in the normalised input space $(x_1, x_2) \in [0, 1]^2$. The boundary is defined as

$$f(x_1) = 0.5 \sqrt{\max(x_1 - 0.2, 0)}.$$

Class labels are assigned according to

$$y(x_1, x_2) = \begin{cases} 0, & \text{if } x_2 > f(x_1), \\ 1, & \text{if } x_2 \leq f(x_1). \end{cases}$$

The active learning process is considered converged after 58 iterations, at which point the boundary spread metric drops below the threshold value of 0.01. At convergence, the global misclassification error is found to be 0.365%. The final GPC prediction field obtained after convergence is shown in Figure J.5.

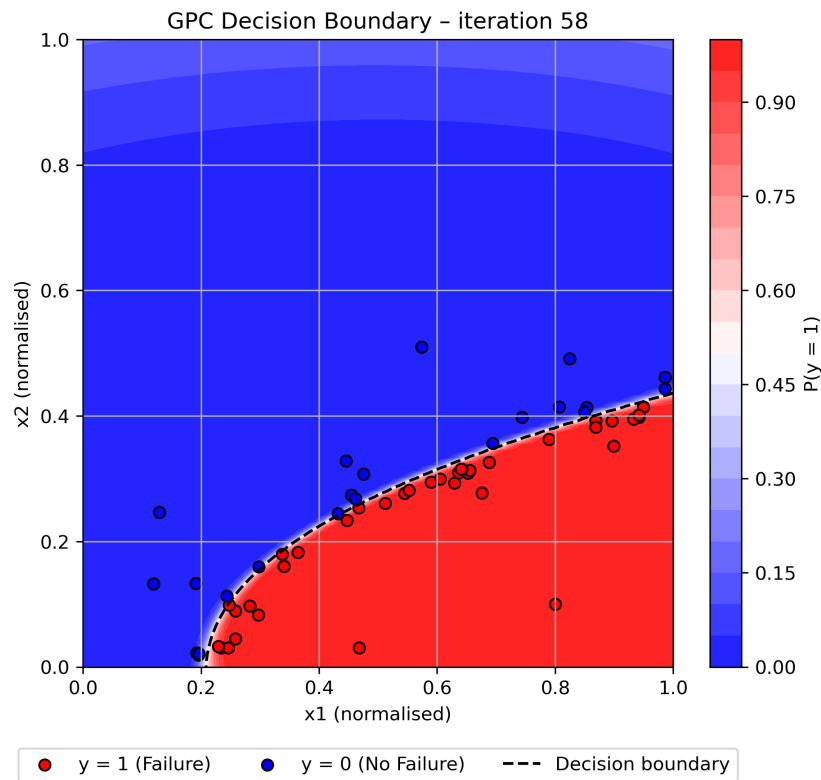


Figure J.5: GPC prediction field for the two-dimensional square-root boundary trial function.

As in the previous cases, several monitoring metrics are tracked throughout the active learning process to assess convergence behaviour and surrogate model performance. These metrics include the acquisition values, changes in the estimated decision boundary, the boundary spread, the log-likelihood, and the misclassification ratio. The evolution of these quantities for the square-root boundary case is presented in Figure J.6.

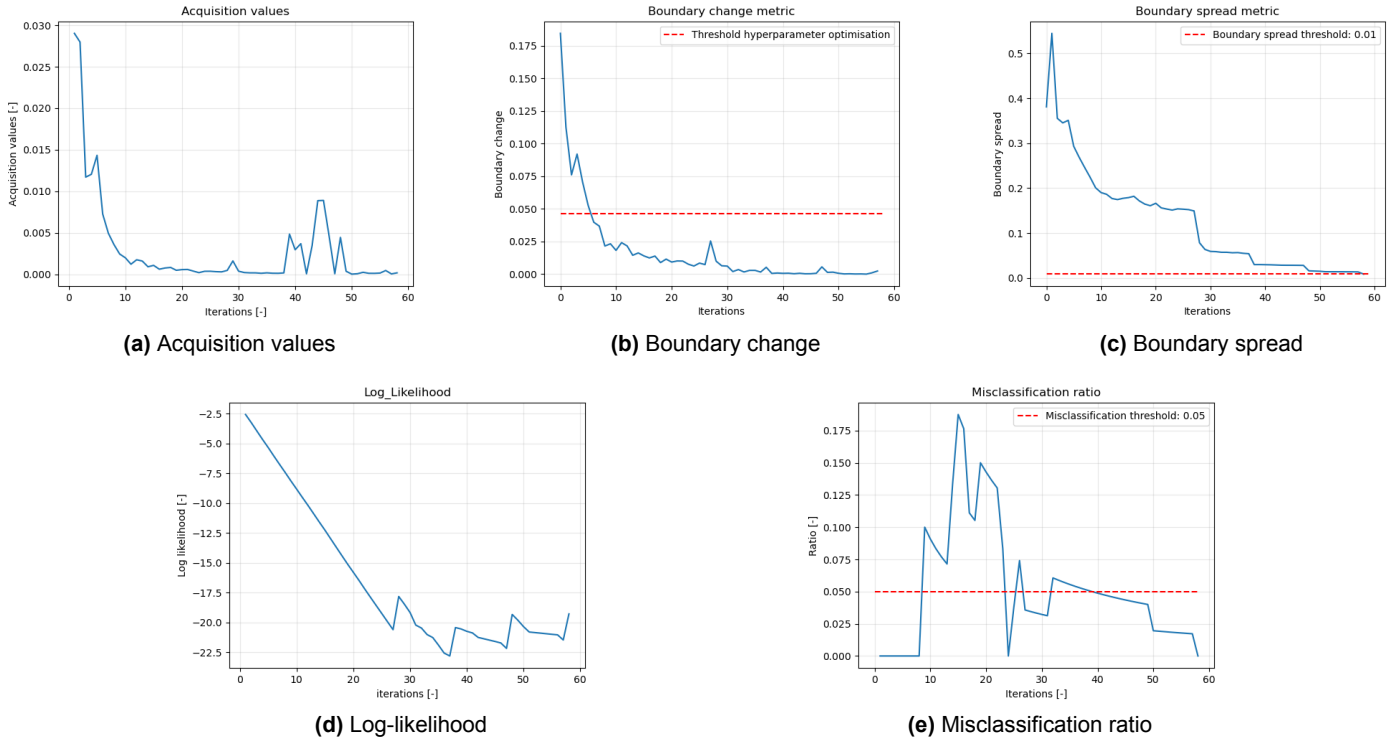


Figure J.6: Monitoring metrics for the two-dimensional square-root boundary trial function.

Sinus Boundary

The fourth two-dimensional trial function considered features a sinusoidal decision boundary in the normalised input space $(x_1, x_2) \in [0, 1]^2$. The boundary is defined as

$$f(x_1) = x_1 + 0.2 \sin(7\pi x_1).$$

Class labels are assigned according to

$$y(x_1, x_2) = \begin{cases} 0, & \text{if } x_2 > f(x_1), \\ 1, & \text{if } x_2 \leq f(x_1). \end{cases}$$

The active learning process is considered converged after 221 iterations, at which point the boundary spread metric drops below the threshold value of 0.01. At convergence, the global misclassification error is found to be 0.327%. The final GPC prediction field obtained after convergence is shown in Figure J.7.




Figure J.7: GPC prediction field for the two-dimensional sinus boundary trial function.

As in the previous cases, several monitoring metrics are tracked throughout the active learning process to assess convergence behaviour and surrogate model performance. These metrics include the acquisition values, changes in the estimated decision boundary, the boundary spread, the log-likelihood, and the misclassification ratio. The evolution of these quantities for the sinus boundary case is presented in Figure J.8.

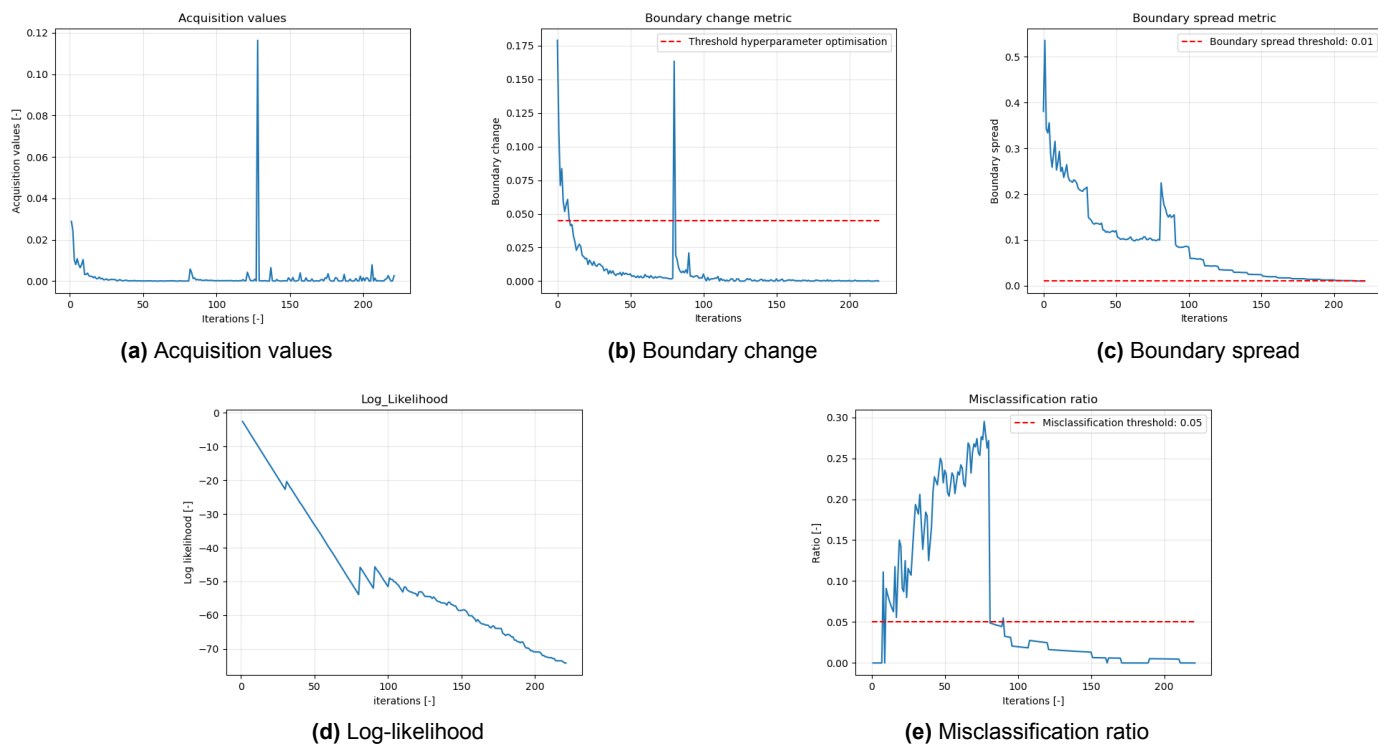



Figure J.8: Monitoring metrics for the two-dimensional sinus boundary trial function.

Conclusions

The two-dimensional trial-functions considered demonstrate that the proposed active learning strategy for GPC effectively identifies and refines the decision boundary across a range of boundary shapes and nonlinearities. In all cases, newly selected training points increasingly concentrate near the decision boundary as the active learning process progresses, while relatively few points are placed deep within homogeneous class regions. This behaviour is consistent with the intended objective and strategy to efficiently resolve the classification between two states. In addition, the applied active learning settings are therefore found to be adequate for the two-dimensional trial functions considered.

For all investigated trial functions, a boundary spread value below 0.01 is found to coincide with a highly accurate surrogate prediction field and a low final global misclassification error. This metric therefore provides a robust convergence criterion for terminating the active learning process. The boundary change metric is found to be useful for guiding hyperparameter optimisation. When the boundary change drops below approximately 25% of its maximum observed value, the decision boundary is found to be sufficiently well enclosed to allow stable optimisation. To ensure robustness, this condition is required to be satisfied for 20 consecutive iterations before hyperparameter optimisation is triggered.

The number of active learning iterations required to achieve convergence is found to depend strongly on the complexity and nonlinearity of the underlying decision boundary. Smooth boundaries converge relatively quickly, while highly oscillatory boundaries, such as the sinusoidal case, require a substantially larger number of model evaluations. Nevertheless, in all cases the final surrogate models achieve low misclassification errors, confirming the robustness of the framework. The insights gained from these controlled trial-function studies provide a sound basis for extending the framework to higher-dimensional input spaces, as explored in subsequent appendices.



GPC AL: 3D Trial Functions

This appendix presents the results of two three-dimensional synthetic trial functions used to further develop the GPC surrogate modelling framework and its associated active learning strategy. The three-dimensional cases extend the two-dimensional trial functions discussed previously and allow the behaviour of the framework to be examined in a higher-dimensional input space.

While full visualisation of the surrogate prediction field is no longer feasible in three dimensions, it remains possible to inspect the distribution of sampled training points and the estimated decision boundary. This still provides valuable qualitative insight into the convergence behaviour of the active learning algorithm and the effectiveness of the training-point selection strategy.

Two three-dimensional trial functions are considered, featuring a nonlinear boundary and a sinusoidal boundary. These cases are analysed alongside the two-dimensional trial functions to identify which convergence criteria and algorithmic settings are effective for three-dimensional problems. The results also provide an indication of the number of model evaluations required as dimensionality increases and serve as a basis for further extension of the framework to higher-dimensional input spaces.

Results

All monitoring metrics described in Subsection 4.3.2 are evaluated for the three-dimensional trial functions. Hyperparameter optimisation and convergence are performed following the standard procedure outlined in the main text, unless stated otherwise. The monitoring metrics are evaluated using a fixed Sobol sampling set consisting of 200,000 points.

The final surrogate models are assessed using a global misclassification error, computed by comparing the model predictions with the known ground truth. This error is estimated using an independent Sobol sampling set of 200,000 points distributed over the input domain.

The active learning configuration used for these three dimensional cases includes 200 gradient-based optimisation steps per iteration with a learning rate of 0.001 for the stochastic gradient descent procedure used in the SMOCU point optimisation. For each iteration, an initial random Sobol sampling set of 200,000 points is generated, from which a retaining fraction of 0.001 is applied, resulting in 200 candidate points for SMOCU evaluation. The one-dimensional SMOCU integral is approximated using 2,000 truncated Gaussian points.

Nonlinear Boundary

The first three-dimensional trial function considered features a nonlinear decision boundary in the normalised input space $(x_1, x_2, x_3) \in [0, 1]^3$. The decision boundary is defined as

$$f(x_1, x_2) = 0.2 + 0.3(x_1^2 + x_2^2),$$

with class labels assigned according to

$$y(x_1, x_2, x_3) = \begin{cases} 0, & \text{if } x_3 > f(x_1, x_2), \\ 1, & \text{if } x_3 \leq f(x_1, x_2). \end{cases}$$

The active learning process is considered converged after 87 iterations, at which point the boundary spread metric drops below the threshold value of 0.01. At convergence, the global misclassification error is found to be 0.241 %. A visualisation of the estimated decision boundary is shown in Figure K.1.




Figure K.1: Estimated decision boundary for the three-dimensional nonlinear boundary trial function.

To assess convergence behaviour and surrogate model performance, several monitoring metrics are tracked throughout the active learning process. These include the acquisition values, changes in the estimated decision boundary, the boundary spread, the log-likelihood, and the misclassification ratio. The evolution of these quantities for the nonlinear boundary case is presented in Figure K.2.

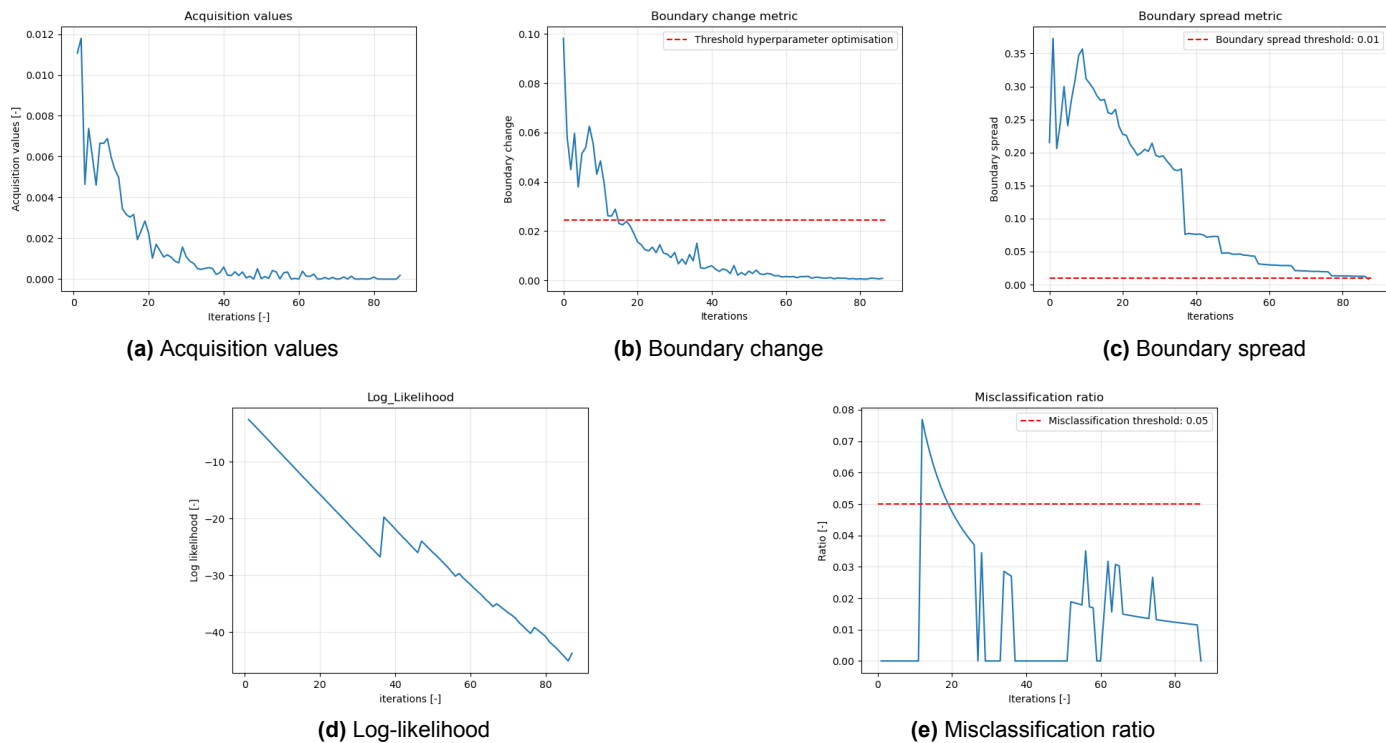


Figure K.2: Monitoring metrics for the three-dimensional nonlinear boundary trial function.

Sinus Boundary

The second three-dimensional trial function considered features a nonlinear decision boundary with an additional sinusoidal component in the normalised input space $(x_1, x_2, x_3) \in [0, 1]^3$. The decision boundary is defined as

$$f(x_1, x_2) = 0.2 + 0.3(x_1^2 + x_2^2) + 0.1 \sin(2\pi x_1),$$

with class labels assigned according to

$$y(x_1, x_2, x_3) = \begin{cases} 0, & \text{if } x_3 > f(x_1, x_2), \\ 1, & \text{if } x_3 \leq f(x_1, x_2). \end{cases}$$

The active learning process is considered converged after 128 iterations, at which point the boundary spread metric drops below the threshold value of 0.01. At convergence, the global misclassification error is found to be 0.224%. A visualisation of the estimated decision boundary is shown in Figure K.3.




Figure K.3: Estimated decision boundary for the three-dimensional sinus boundary trial function.

As in the previous case, monitoring metrics are tracked throughout the active learning process to assess convergence behaviour and surrogate model performance. The evolution of these quantities for the sinus boundary case is presented in Figure K.4.

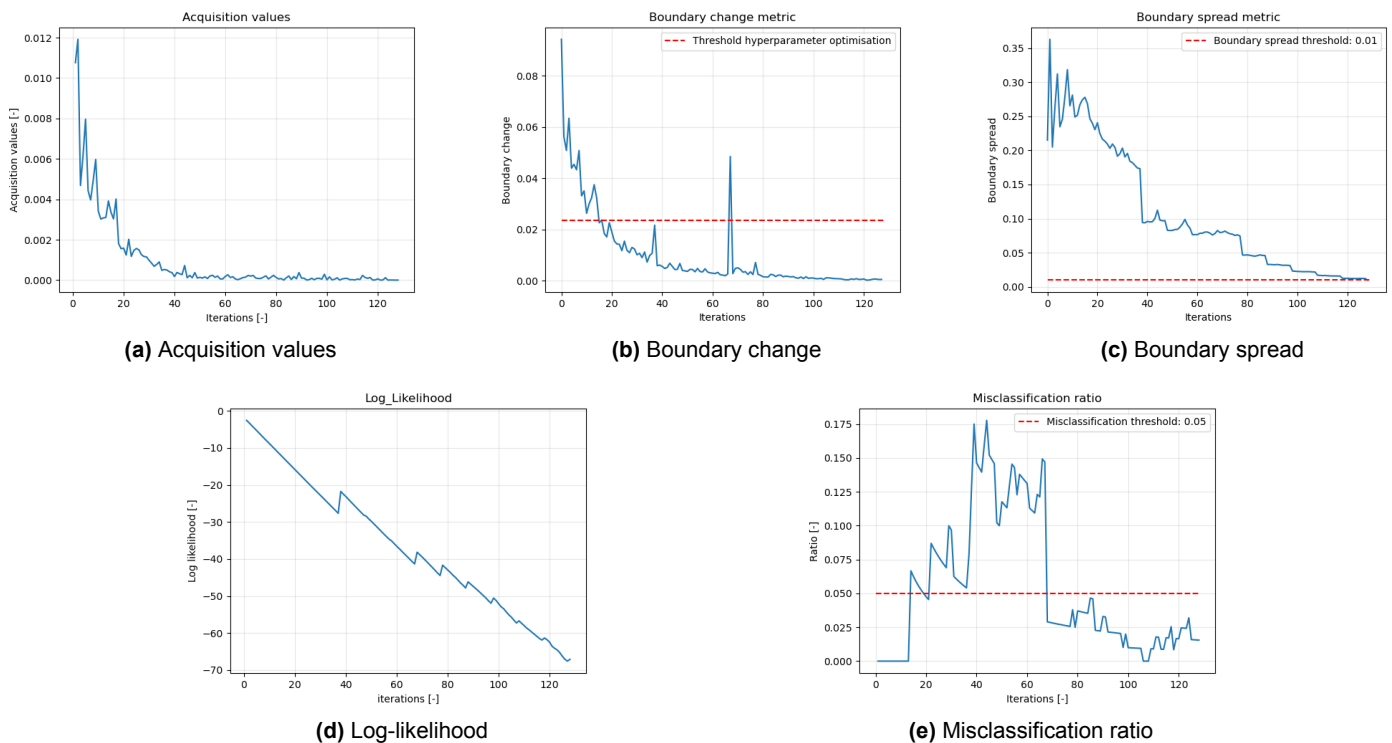



Figure K.4: Monitoring metrics for the three-dimensional sinus boundary trial function.

Conclusions

The three-dimensional trial functions show that the active learning strategy consistently places new training points close to the decision boundary, which aligns with the intended objective of efficiently

refining the classification surface. The applied active learning settings are found to be adequate for the three-dimensional trial functions considered. Similar to the behaviour observed in the two-dimensional trial functions, the number of iterations required for convergence remains problem dependent and increases with the complexity of the decision boundary.

The convergence strategy defined for the two-dimensional trial functions is found to remain applicable in three dimensions. In particular, the boundary spread metric provides a reliable indicator of convergence, while the boundary change metric offers useful guidance for the timing of hyperparameter optimisation. These results indicate that the proposed active learning framework extends naturally to three-dimensional input spaces and provides a sound basis for further expansion to higher-dimensional problems.



GPC AL: 6D Trial Function

This appendix presents the results of a six-dimensional synthetic trial function used to further evaluate the GPC surrogate modelling framework and its associated active learning strategy. The six-dimensional case represents a substantial increase in problem complexity compared to the lower-dimensional trial functions discussed in previous appendices and serves to assess the practical applicability of the framework in higher-dimensional input spaces.

Due to the dimensionality of the problem, direct visualisation of the prediction field, the full decision boundary, and the distribution of sampled training points is no longer feasible. Insight into the performance of the surrogate model is therefore obtained through a combination of monitoring metrics, the final global misclassification error, and two-dimensional slices through the six-dimensional input space. Median and randomly selected slices are subtracted which provide to some extent qualitative information on the local accuracy of the learned decision boundary. The monitoring metrics and global error quantify convergence behaviour and overall model performance.

The six-dimensional trial function includes nonlinear terms and cross-interactions between input variables, providing a representative test case for the active learning strategy. The results are analysed alongside those of the lower-dimensional trial functions to identify which convergence criteria and algorithmic settings are effective in six dimensions and to provide an indication of the number of model evaluations required at this dimensionality.

Results

All monitoring metrics described in Subsection 4.3.2 are evaluated for the six-dimensional trial function. Hyperparameter optimisation and active learning settings follow the standard procedure outlined in the main text.

Monitoring metrics are evaluated using a fixed Sobol sampling set consisting of 1,000,000 points. The final global misclassification error is also estimated using a Sobol sample of 1,000,000 points distributed over the input domain.

The active learning configuration used for this case includes 200 gradient-based optimisation steps per iteration with a learning rate of 0.001 for the stochastic gradient descent procedure used in the SMOCU point optimisation. For each iteration, an initial random Sobol sampling set of 1,000,000 points is generated, from which a retaining fraction of 0.001 is applied, resulting in 1,000 candidate points for SMOCU evaluation. The one-dimensional SMOCU integral is approximated using 4,000 truncated Gaussian points.

The nonlinear decision boundary in the normalised input space $(x_1, x_2, x_3, x_4, x_5, x_6) \in [0, 1]^6$ is defined by the hypersurface

$$f(x_1, x_2, x_3, x_4, x_5) = 0.25 + 0.20 x_1^2 + 0.10 x_2 + 0.15 x_3^2 + 0.10 x_4^2 x_5 \\ + 0.05 x_1 x_3^2 - 0.05 x_2 x_4.$$

Class labels are assigned according to

$$y(x_1, x_2, x_3, x_4, x_5, x_6) = \begin{cases} 0, & \text{if } x_6 > f(x_1, x_2, x_3, x_4, x_5), \\ 1, & \text{if } x_6 \leq f(x_1, x_2, x_3, x_4, x_5). \end{cases}$$

The active learning procedure is terminated after 411 iterations due to computational considerations. At this point, the final global misclassification error is measured to be equal to 0.246 %.

To assess local accuracy of the learned decision boundary, two-dimensional slices through the six-dimensional space are extracted. Median slices, obtained by fixing all non-plotted variables at their median values, are shown in Figure L.1. Random slices, obtained by fixing non-plotted variables at randomly selected values, are shown in Figure L.2.

The evolution of the monitoring metrics throughout the active learning process is shown in Figure L.3.




Figure L.1: Median slices of the six-dimensional considered trial function




Figure L.2: Random slices of the six-dimensional considered trial function

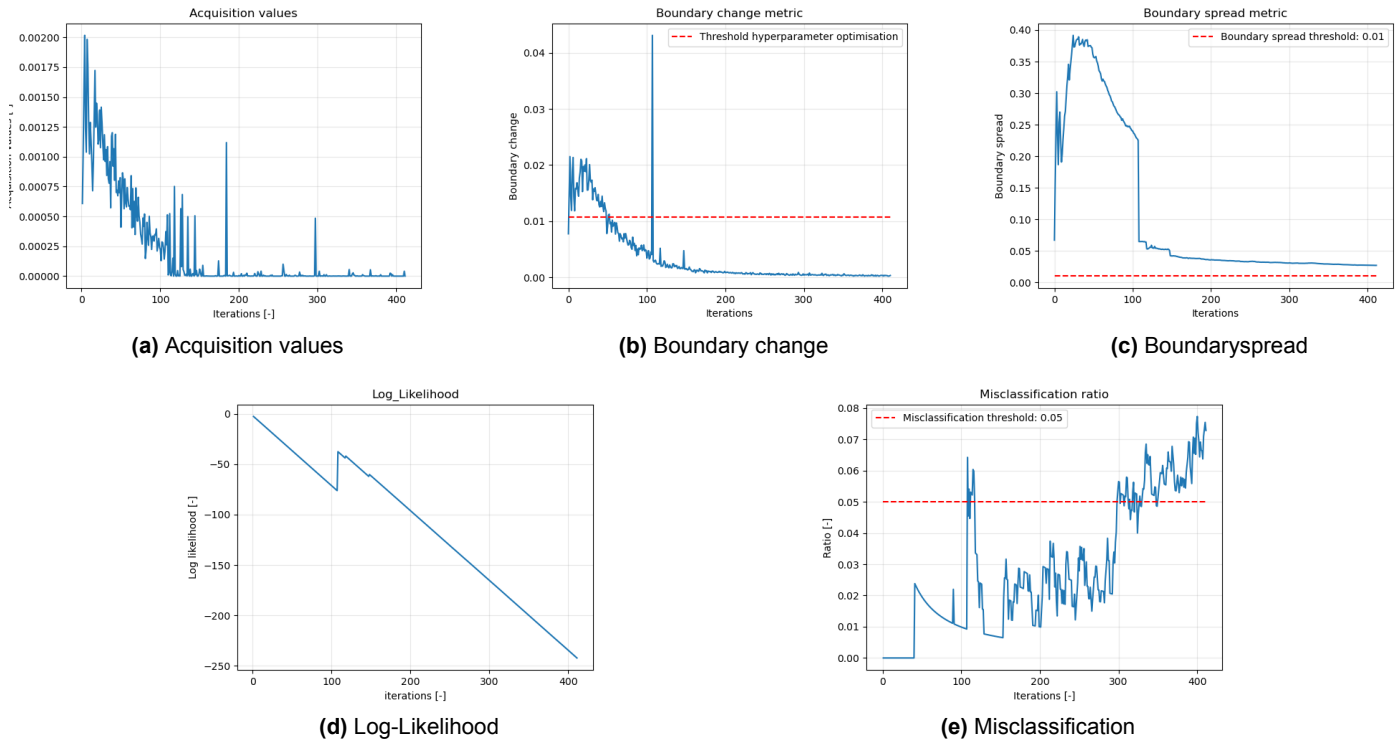


Figure L.3: Monitoring metrics for the six-dimensional considered trial function


Conclusions

When considering the final global misclassification error, the extracted two-dimensional slices, and the evolution of the monitoring metrics, the results indicate that the surrogate modelling framework remains effective for this six-dimensional classification problem. Together, these measures suggest that the learned decision boundary is captured with sufficient accuracy, despite the increased dimensionality of the input space.

To account for the higher dimensionality, the final global misclassification error is evaluated using an increased Sobol sampling set of 1,000,000 points. Although this represents a substantial increase compared to the lower-dimensional cases, it should be noted that the sampling remains relatively sparse in six dimensions. Nevertheless, repeated evaluations using independently generated Sobol sample sets result in consistent global misclassification errors. This indicates that the chosen sampling size provides a sufficiently stable estimate of the overall surrogate performance for this case.

Convergence is not fully achieved according to the boundary spread criterion, which attains a value of 0.0271 after 411 iterations. However, as discussed in Section 4.3, reductions in boundary spread are indicative of improved prediction field fitting. The agreement observed in the extracted median and random slices suggests that the decision boundary is captured with sufficient accuracy at this stage, even though the formal convergence threshold has not yet been reached. These slices provide valuable qualitative insight into local accuracy, which cannot be assessed using global misclassification metrics alone.

The applied active learning settings are found to be adequate for the six-dimensional trial function considered. This is reflected in the overall downward trend of the acquisition values and the decreasing boundary change metric, both of which indicate continued refinement of the decision boundary. The misclassification metric remains above the convergence level during once the active learning process has been terminated. However this is consistent with the fact that full convergence has not yet been achieved. As sampling increasingly concentrates near the decision boundary, temporary increases or slower reductions in the misclassification error are expected. When a final hyperparameter optimisation is performed on the resulting dataset, the misclassification metric is observed to drop below the prescribed threshold.



GPC AL: 12D Trial Function

This appendix presents the application of the GPC surrogate modelling framework to a twelve-dimensional synthetic trial function. The twelve-dimensional case represents a substantial increase in problem complexity compared to the lower-dimensional trial functions considered in previous appendices and is used to explore the practical limitations and challenges of active learning in high-dimensional input spaces.

A single twelve-dimensional trial function is considered, incorporating nonlinear terms and cross-interactions between input variables. Due to the dimensionality of the problem, direct visualisation of the full prediction field and the distribution of sampled training points is no longer feasible. Insight into model behaviour and accuracy is therefore obtained through a combination of monitoring metrics, global misclassification measures, and extracted two-dimensional slices through the high-dimensional input space. These slices provide qualitative information on local agreement between the surrogate prediction and the known ground truth, while the monitoring metrics and final error measures are used to assess convergence and overall model performance.

Due to the increased size of the twelve-dimensional design space, sampling becomes inherently sparse for any practical number of model evaluations. For this reason, the active learning strategy is adapted relative to the approaches used for lower-dimensional problems. These modifications aim to improve sampling efficiency and numerical stability, while retaining the core principles of the proposed framework. As a result, this appendix places greater emphasis on methodological insights, algorithmic behaviour, and observed limitations, rather than on demonstrating full convergence.

The results and discussion presented in this appendix therefore serve both as an evaluation of the current framework in twelve dimensions and as a basis for identifying directions for further improvement and future research.

Results

For the twelve-dimensional trial function, a direct extension of the active learning approach used for the six-dimensional case becomes computationally prohibitive. Although Sobol sampling sets of up to one million points were feasible in six dimensions, such sampling becomes increasingly inefficient in twelve dimensions due to the rapid growth of the design space. To address this, the active learning strategy is adapted to improve sampling efficiency while retaining the core principles of the framework.

At each active learning iteration, an initial global Sobol sampling of 1,000,000 points is performed over the twelve-dimensional input space. Given the sparsity of this sampling, an intermediate gradient-based entropy optimisation step is introduced. A small fraction of the Sobol points, corresponding to the highest entropy values, is retained to identify candidate locations near the estimated decision boundary. Specifically, a retaining fraction of 2×10^{-5} is applied, resulting in 20 candidate points. For each of these points, 80 stochastic gradient-based optimisation steps are performed on the entropy field using a learning rate of 10^{-5} , guiding the points closer to the decision boundary. These parameter values were selected based on iterative trial-and-error experimentation across lower-dimensional cases and the present twelve-dimensional problem.

To ensure that the optimised candidate points are indeed located near the boundary, a probabilistic

boundary filter is applied, retaining only points with predicted class probabilities between 0.4 and 0.6. The SMOCU acquisition function is then evaluated for these filtered points, and the highest-scoring point is selected. This point is subsequently refined using 200 stochastic gradient-based optimisation steps on the SMOCU acquisition field with a learning rate of 0.001. The one-dimensional SMOCU integral is approximated using 5,000 truncated Gaussian points.

Monitoring metrics are evaluated using a fixed Sobol sampling set of 1,000,000 points. To promote exploration in the early stages of learning, the initial kernel variance is set to 5, which reduces the strength of the prior class preference and encourages broader exploration of the design space. Hyperparameter optimisation is enabled once the boundary change metric has dropped below 25% of its maximum value on twenty occasions. Once enabled, optimisation is performed every ten iterations. During hyperparameter optimisation, only the kernel variance is updated if the log-marginal likelihood improves by more than 5%.

The nonlinear decision boundary in the normalised input space $(x_1, x_2, \dots, x_{12}) \in [0, 1]^{12}$ is defined by the hypersurface

$$\begin{aligned} f(x_1, \dots, x_{12}) = & 0.30 + 0.12 x_1^2 + 0.08 x_2 + 0.10 x_3^2 + 0.06 x_4 x_5 + 0.07 x_6^2 \\ & + 0.05 x_7 x_8 + 0.06 x_9^2 + 0.04 x_{10} x_{11} + 0.03 x_1 x_3 \\ & - 0.04 x_2 x_4 + 0.02 x_5 x_6^2. \end{aligned}$$

Class labels are assigned according to

$$y = \begin{cases} 0, & \text{if } x_{12} > f(x_1, \dots, x_{11}), \\ 1, & \text{if } x_{12} \leq f(x_1, \dots, x_{11}). \end{cases}$$

The active learning procedure is terminated after 932 iterations due to computational constraints. At this stage, the estimated global misclassification error is found to be 0.699%, evaluated using 1,000,000 Sobol points distributed over the design space. To gain qualitative insight into local model accuracy, two-dimensional median slices and randomly selected slices of the twelve-dimensional input space are extracted, as shown in Figure M.1 and Figure M.2. These slices illustrate the local agreement between the surrogate prediction and the known ground truth for selected variable pairs. In addition, the evolution of the acquisition values, boundary change, boundary spread, log-likelihood, and misclassification ratio shown in Figure M.3 provides complementary information on convergence behaviour and the effectiveness of the adapted active learning strategy.




Figure M.1: Median slices of the twelve-dimensional considered trial function

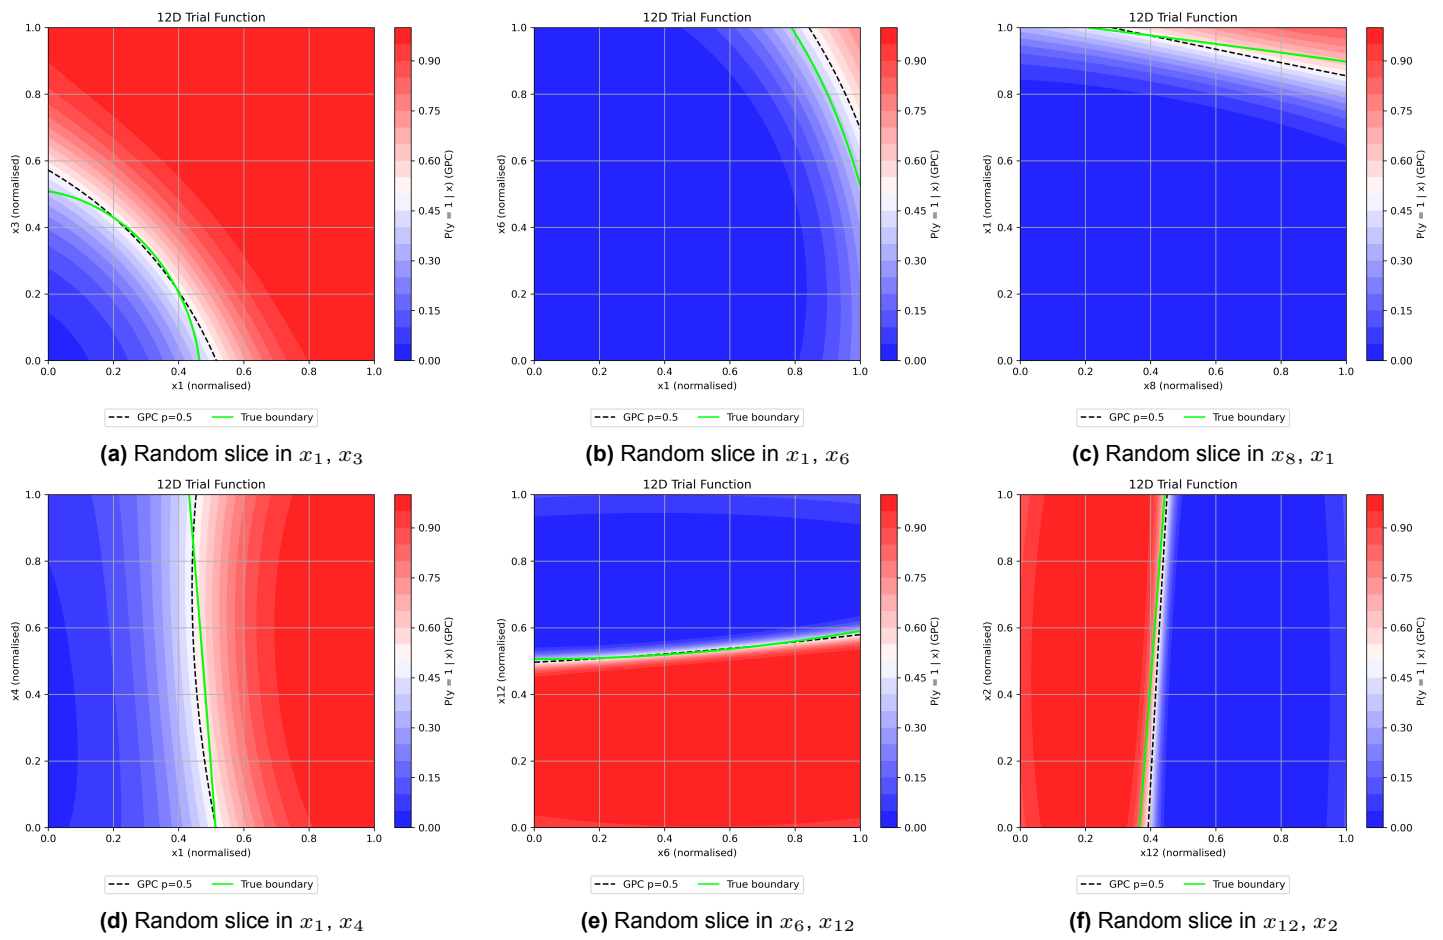


Figure M.2: Random slices of the twelve-dimensional considered trial function

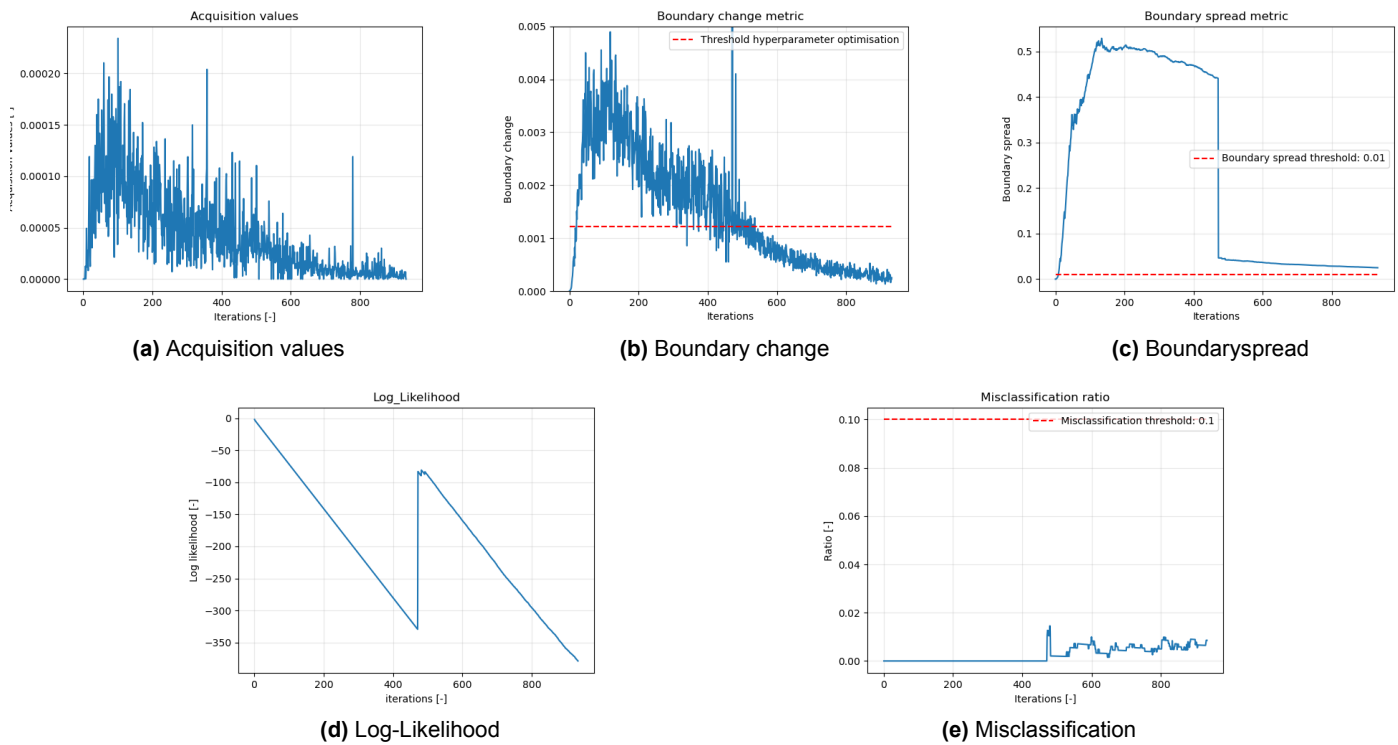


Figure M.3: Monitoring metrics for the twelve-dimensional considered trial function

Discussions and Conclusions

The twelve-dimensional trial function shows that the proposed surrogate modelling framework can be extended to substantially higher-dimensional classification problems when suitable adaptations to the active learning strategy are introduced. Although the decision boundary is not yet fully captured at termination, the monitoring metrics and extracted two-dimensional slices indicate that the active learning procedure is converging and that increasingly informative points are being selected. The acquisition values exhibit a clear downward trend and the boundary change metric stabilises, both of which signal progressive refinement of the classification surface. The extracted median and random slices further suggest that the decision boundary continues to improve as additional samples are acquired. The active learning procedure is terminated due to computational constraints rather than a lack of convergence, and the observed trends indicate that continued sampling would likely lead to further improvements in boundary accuracy.

In addition, it is observed that in the x_{12} direction, where class separation is explicitly defined, the decision boundary is already captured with relatively high accuracy. This behaviour may be related to the formulation of the trial function, in which the separation between classes is directly governed by this variable.

A key methodological change in this active learning procedure is the effective transition to a hybrid entropy–SMOCU strategy. Gradient-based optimisation on the entropy field primarily acts as a boundary-seeking mechanism, while subsequent SMOCU-based optimisation focuses on refining the acquisition along the estimated decision boundary. Initial attempts using solely the standard SMOCU procedure did not exhibit stable or meaningful convergence behaviour in this high-dimensional setting. This remained the case even when the active learning settings were increased within practical computational limits, which motivated the adoption of the hybrid approach.

However, several important limitations remain when combining the hybrid active learning approach with the expansion to a higher-dimensional design space.

First, the increased computational cost associated with entropy-based gradient optimisation limits the number of candidate points that can be evaluated using SMOCU, which in turn reduces the level of exploration along the decision boundary.

Second, optimisation of the kernel variance improves the overall accuracy of the surrogate model but simultaneously leads to a sharpening of the posterior entropy field. As additional training data are acquired near the decision boundary, predictive uncertainty is strongly reduced in the neighbourhood of sampled locations. This causes the entropy field to become relatively flat around existing samples, while increasingly sharp peaks remain concentrated at the decision boundary itself. Consequently, gradient-based optimisation on the entropy field becomes less stable, and reliably locating the decision boundary in very high-dimensional spaces becomes more challenging. For this reason, kernel variance optimisation had to be applied in a restricted manner.

Third, as the number of training points increases in order to effectively learn an eleven-dimensional decision boundary embedded within a twelve-dimensional design space, Gaussian Process Classification becomes progressively more computationally expensive, ultimately limiting the feasible number of active learning iterations.

Finally, convergence monitoring based on global Sobol sampling becomes less informative in high-dimensional input spaces due to the increasing sparsity of the sampling.

Based on these observations, several recommendations can be made for further development. Access to larger computational resources would enable denser initial sampling or optimisation of a larger number of entropy-based candidate points, which could improve the effectiveness of the SMOCU point selection. In addition, transitioning to sparse Gaussian Processes may significantly improve scalability for problems beyond six dimensions, as these approaches scale more efficiently with increasing available data. Finally, an alternative convergence criteria could be explored by locally assessing changes in the prediction field before and after the inclusion of new samples, for example through importance sampling in the region around newly selected points. This concept aligns with the underlying formulation of the current acquisition function, but would evaluate the actual impact of a newly added sample once it has been incorporated into the model.

A further promising direction for improvement lies in replacing gradient-based optimisation on the entropy field with optimisation on the GPC latent mean field. The latent mean field is expected to exhibit more stable behaviour, particularly after kernel variance optimisation. Moreover, the initial Sobol sampling over the design space may be used to estimate the typical magnitude and range of latent mean values, enabling adaptive selection of learning rates, optimisation step counts, and boundary filters defined directly in latent space. Such an approach is expected to improve both the robustness and efficiency of the active learning procedure and represents a valuable direction for future development of the framework for high-dimensional classification problems.

N

Bounds Determination

This appendix provides supporting information on the determination of the bounds adopted for the design space of the surrogate modelling variables considered in this study. The variables are grouped into three categories. First, the structural design variables are treated which describe the column geometry and loading conditions. Second, the fire variables are presented which characterise the applied thermal exposure, including both heating and cooling phases. Finally, the material variables which represent the governing properties of timber used in the resistance assessment are considered.

The appendix is organised according to this categorisation. Each section reports the variables included in the respective group and documents the determination of their bounds for the surrogate modelling framework.

Structural Design Variables

The structural design space is defined by the column width B , the column length l , and the applied axial force F . These variables represent the governing geometric and loading characteristics for the column configurations considered in the surrogate modelling.

The bounds for B and l were selected using first-order design assumptions for typical timber columns in two representative building typologies: a residential mid-rise building and an office high-rise building. Based on conventional storey heights and practical sizing considerations for primary load-bearing members, the following geometric ranges were adopted:

$$B \in [250, 500] \text{ mm}, \quad l \in [3.0, 5.0] \text{ m}.$$

The axial force bounds were subsequently defined using simplified area-based load estimates for the considered building typologies. A first-order estimate of the axial force in a representative interior column was obtained from a defined influential area A , the number of supported floors n , and characteristic floor actions. The variable action was reduced by a factor 0.3 to reflect the reduced representative value adopted for variable loads in the accidental (fire) design situation in correspondence with NEN-EN 1991-1-2 [20]. Accordingly, the axial force was approximated as

$$F \approx n A (p_g + p_Q),$$

where p_g is the permanent floor load and p_Q is the reduced variable floor load.

Residential mid-rise:

For the residential mid-rise case, the following values were adopted. The influence area was taken as $A \approx 10 \text{ m}^2$, corresponding to a column spacing of approximately $3 \times 3 \text{ m}$.

$$p_g = 4.0 \text{ kN/m}^2, \quad p_Q = 2.25 \times 0.3 = 0.675 \text{ kN/m}^2, \quad A = 10 \text{ m}^2, \quad n \in [6, 10].$$

This yields

$$F_{\min} \approx 6 \cdot 10 \cdot (4.0 + 0.675) = 282 \text{ kN}, \quad F_{\max} \approx 10 \cdot 10 \cdot (4.0 + 0.675) = 470 \text{ kN}.$$

Office high-rise:

For the office high-rise case, the following values were adopted. The influence area was taken as $A = 25 \text{ m}^2$, corresponding to a column spacing of approximately $5 \times 5 \text{ m}$.

$$p_g = 3.0 \text{ kN/m}^2, \quad p_Q = 4.0 \times 0.3 = 1.2 \text{ kN/m}^2, \quad A = 25 \text{ m}^2, \quad n \in [10, 18].$$

This yields

$$F_{\min} \approx 10 \cdot 25 \cdot (3.0 + 1.2) = 1050 \text{ kN}, \quad F_{\max} \approx 18 \cdot 25 \cdot (3.0 + 1.2) = 1890 \text{ kN}.$$

Based on these first-order estimates, and on practical rounding, the force bounds were finally set to

$$F \in [300, 1800] \text{ kN}.$$

To verify which fire resistance range is implied by the selected structural bounds, an inverse fire resistance calculation was performed. For a given column configuration (B, l, F) , a bisection method was performed on the fire exposure duration until a unity check of 1.0 was reached. The resulting fire duration was eventually reported as the fire resistance time R . The procedure follows the Eurocode NEN-EN 1995-1-2 reduced cross-section method, in which an effective charring depth is used to define a reduced load-bearing section [1]. For a standard fire exposure, the notional charring depth is taken as

$$d_{\text{char},n}(t) = \beta_n t,$$

and the effective charring depth is computed as

$$d_{\text{ef}}(t) = d_{\text{char},n}(t) + k_0(t) d_0, \quad k_0(t) = \begin{cases} t/20, & t < 20 \text{ min}, \\ 1, & t \geq 20 \text{ min}, \end{cases}$$

where $d_0 = 14 \text{ mm}$ is adopted in accordance with the draft revision of NEN-EN 1995-1-2 [45].

Assuming four-sided exposure and a square cross-section the effective column width is defined as

$$B_{\text{ef}} = B - 2d_{\text{ef}},$$

from which the effective area and second moment of area are determined and used in the subsequent resistance verification. The unity check accounts for axial compression, including stability effects, and for bending due to an assumed initial eccentricity, following NEN-EN 1995-1-1 [49]. The calculations were performed for strength class GL24h. The corresponding Eurocode design material parameters were adopted in accordance with EN 14080 [39].

Finally, the limiting scenarios of the (l, B, F) bounds were evaluated to provide an indication of the fire resistance range covered by the selected structural design space. The results are summarised in Table N.1.

Table N.1: Corresponding fire resistance time R for the limiting scenarios defined by the structural variable bounds.

Scenario	l [m]	B [mm]	F [kN]	R [min]
1	5.0	250	300	35.12
2	5.0	250	1800	0*
3	5.0	500	300	213.68
4	5.0	500	1800	132.76
5	3.0	250	300	61.42
6	3.0	500	300	239.99
7	3.0	250	1800	0*
8	3.0	500	1800	149.90

* Fails already under ambient design conditions, resulting in no meaningful fire resistance time.

Fire Variables

The thermal exposure is defined by the heating duration t_h , the maximum compartment temperature T_{\max} , and the linear cooling rate r_c . These variables were selected to provide a simple but flexible representation of timber compartment fires, explicitly capturing both heating and cooling phases. The adopted bounds were defined based on qualitative comparison with reported temperature–time histories and trends in experimental timber compartment fire studies, covering both ventilation-controlled and fuel-controlled regimes [13–15, 32].

For the surrogate modelling the resulting selected ranges are:

$$t_h \in [1, 60] \text{ min}, \quad r_c \in [6, 20] \text{ }^\circ\text{C/min}, \quad T_{\max} \in [600, 1200] \text{ }^\circ\text{C}.$$

The upper bound for t_h represents comparatively long heating phases observed in ventilation-controlled timber compartment fires, whereas the range for r_c spans slower cooling associated with ventilation-controlled conditions and faster cooling more typical of fuel-controlled decay. The bounds for T_{\max} were chosen to cover moderate to severe peak temperatures reported across the considered experimental studies. The lower bound of $t_h = 1 \text{ min}$ is included to represent very rapid post-flashover heating observed in some fuel-controlled experimental compartment fire tests. While such short heating phases are generally not representative for real building fires, this lower bound is retained because the adopted piecewise fire curve approximation is defined to start at the onset of flashover, rather than at ignition.

Material Variables

The material variables are selected to represent the variability in the governing mechanical properties of GL24h timber, as well as the initial geometric imperfection used in the stability verification. The considered variables are the density ρ , the modulus of elasticity E , the compressive strength f_c , and the initial eccentricity e_0 . Probabilistic descriptions for ρ , E , and f_c were adopted from Schilling et al. [40]. For the initial eccentricity, the length-dependent stochastic distribution reported by Blaß [47] was adopted. The distributions of the variables are summarised in Table N.2.

Since the material variables are treated as stochastic inputs in the surrogate model, their bounds were defined to capture 99.7% of the adopted probability density functions, corresponding to an interval of ± 3 standard deviations around the mean in the underlying normal space.

For ρ , E , and f_c , a lognormal distribution was assumed, such that

$$\ln(X) \sim \mathcal{N}(\lambda, \zeta^2).$$

Subsequently, an approximate 99.7% interval is obtained as

$$X \in [\exp(\lambda - 3\zeta), \exp(\lambda + 3\zeta)].$$

Applying this to the adopted distributions yields, for the density and compressive strength,

$$\rho \in [359, 491] \text{ kg/m}^3, \quad f_c \in [35.7, 45.8] \text{ N/mm}^2.$$

For practical rounding in the surrogate modelling bounds, the density range was finally taken as

$$\rho \in [360, 490] \text{ kg/m}^3.$$

For the modulus of elasticity, a member-size correction on the distribution spread was taken into account through

$$k_{E,h,\zeta} = \left(\frac{0.60}{h} \right)^{0.78},$$

which is applied as $\zeta_{\text{eff}} = k_{E,h,\zeta} \zeta$. To ensure that the bounds remain valid over the full geometric design space, the two limiting member depths within the considered range were evaluated, i.e. $h = B = 250 \text{ mm}$ and $h = B = 500 \text{ mm}$. The final bounds were then taken as the outer most lower bound and upper bound obtained across these two cases, resulting in

$$E \in [9.44, 17.37] \text{ kN/mm}^2.$$

The initial eccentricity was modelled as the sum of a fixed eccentricity of 20 mm and a stochastic out-of-straightness component, taken as normally distributed with length-dependent parameters,

$$e_{\text{dis}} \sim \mathcal{N}(\mu, \sigma^2), \quad \mu = 1.4 \cdot 10^{-5} L, \quad \sigma = 4.53 \cdot 10^{-4} L.$$

To obtain bounds valid for the full length range, the upper length bound $L = 5000$ mm was used, giving


$$e_{\text{dis}} \in [\mu - 3\sigma, \mu + 3\sigma] = [-6.725, 6.865] \text{ mm}.$$

Including the fixed offset, the resulting eccentricity bounds are

$$e_0 = 20 \text{ mm} + e_{\text{dis}} \Rightarrow e_0 \in [13.275, 26.865] \text{ mm}.$$

Table N.2: Probabilistic descriptions of the applied material variables.

Variable	Distribution	Mean	Parameters
Density (ρ)	Lognormal	420 kg/m ³	$\lambda = 6.04, \zeta = 0.0521$
Modulus of elasticity (E)	Lognormal	12.8 kN/mm ²	$\lambda = 2.55, \zeta = 0.0513$
Compressive strength (f_c)	Lognormal	40.4 N/mm ²	$\lambda = 3.70, \zeta = 0.0416$
Eccentricity (e_0)	Normal	0.02 m	$\mu = 1.4 \cdot 10^{-5} \cdot L$ $\sigma = 4.53 \cdot 10^{-4} \cdot L$



Results: Monitoring Metrics

This appendix presents the monitoring metrics recorded during the active learning procedures for the surrogate modelling results presented in Section 5.2. The figures shown here are provided for completeness and are intended to support the interpretation of the convergence behaviour and performance of the active learning framework.

For each considered surrogate model configuration, the evolution of the acquisition values, boundary change, boundary spread, log-likelihood, and global misclassification error is shown as a function of the active learning iterations.

6D Structural Design and Fire Variables

2D Surrogate Model of Structural Design and Fire Variables (t_h, B)

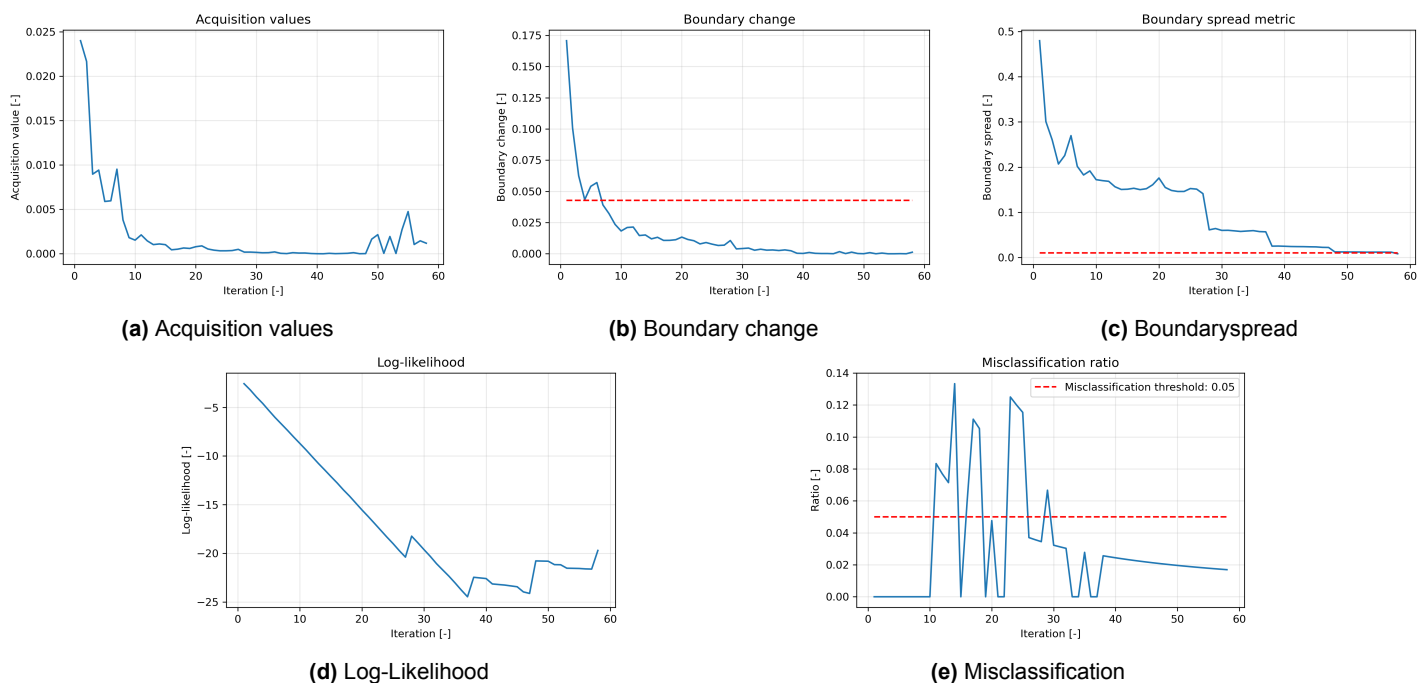


Figure O.1: Monitoring metrics 2D (t_h, B) structural design and fire variables

3D Surrogate Model of Structural Design and Fire Variables (t_h, B, F)

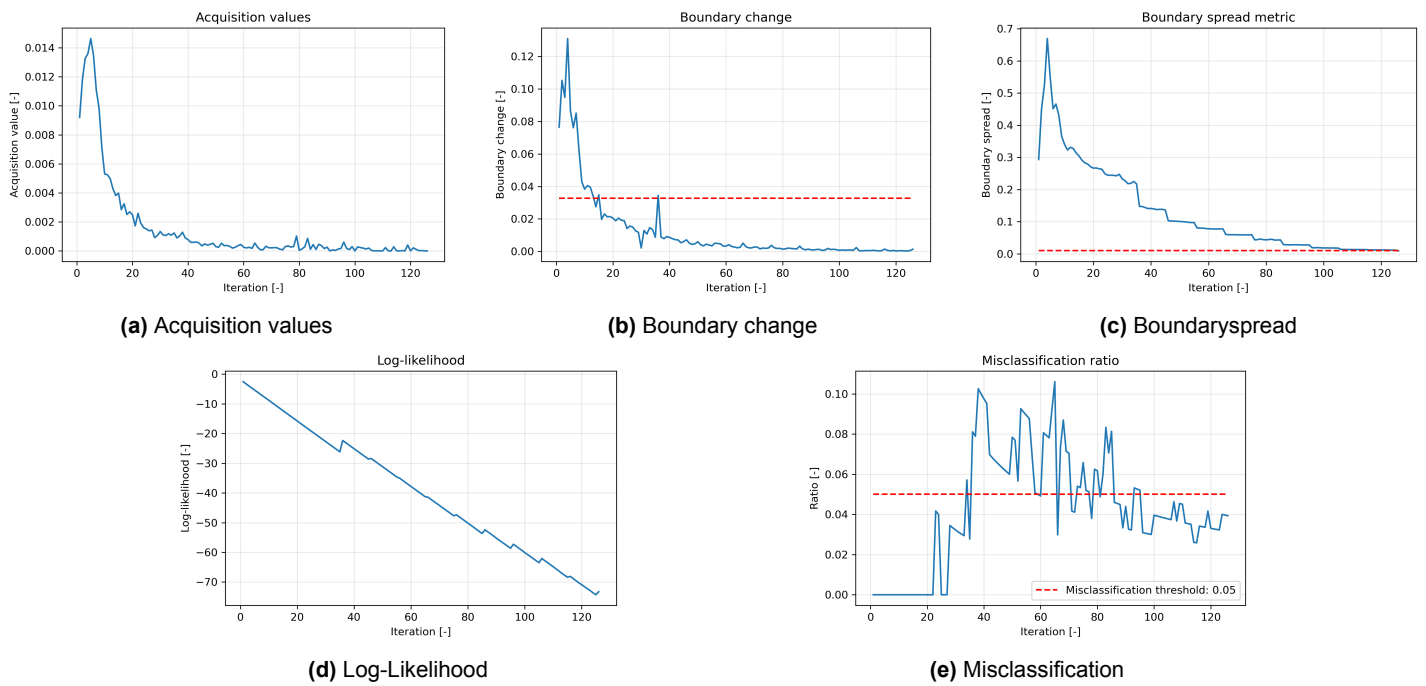


Figure O.2: Monitoring metrics 3D (t_h, B, F) structural design and fire variables

6D Surrogate Model of Structural Design and Fire Variables ($t_h, B, F, l, T_{max}, r_c$)

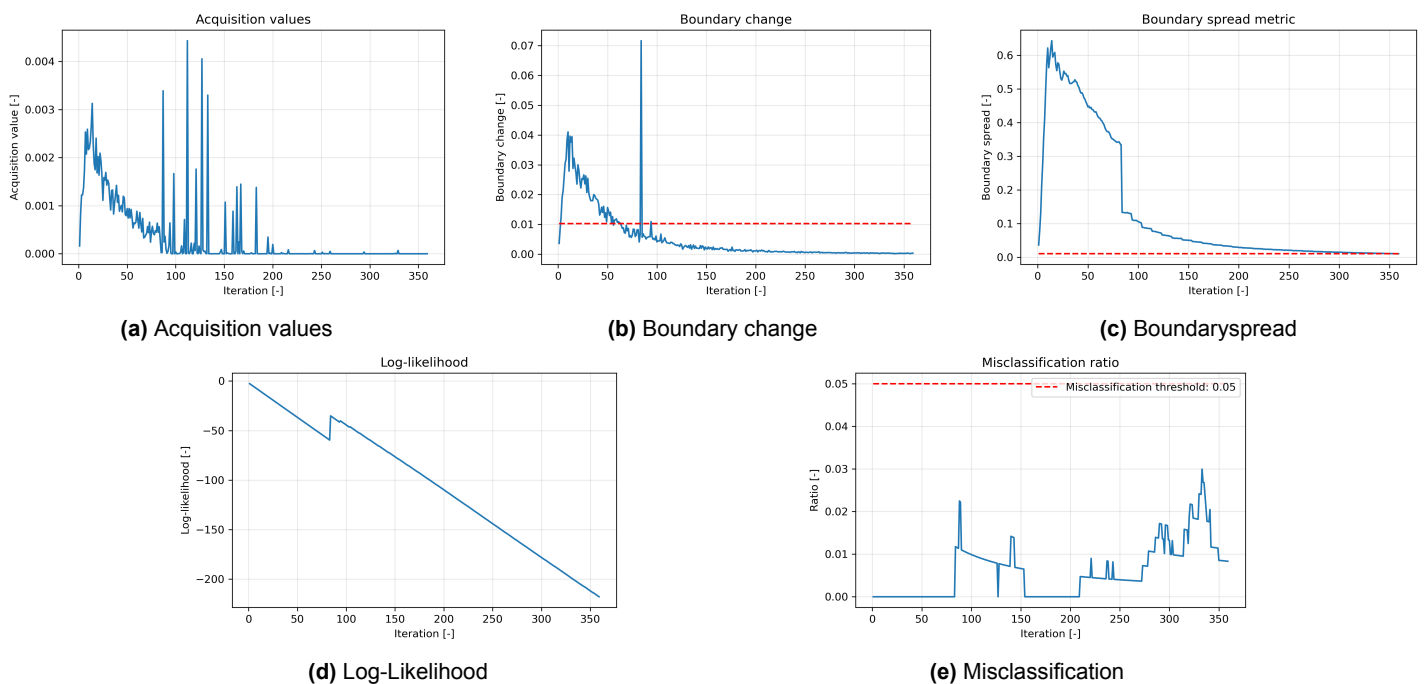


Figure O.3: Monitoring metrics 6D ($t_h, B, F, l, T_{max}, r_c$) structural design and fire variables

6D Material Variables

2D Surrogate Model with Material Variables (t_h, ρ)

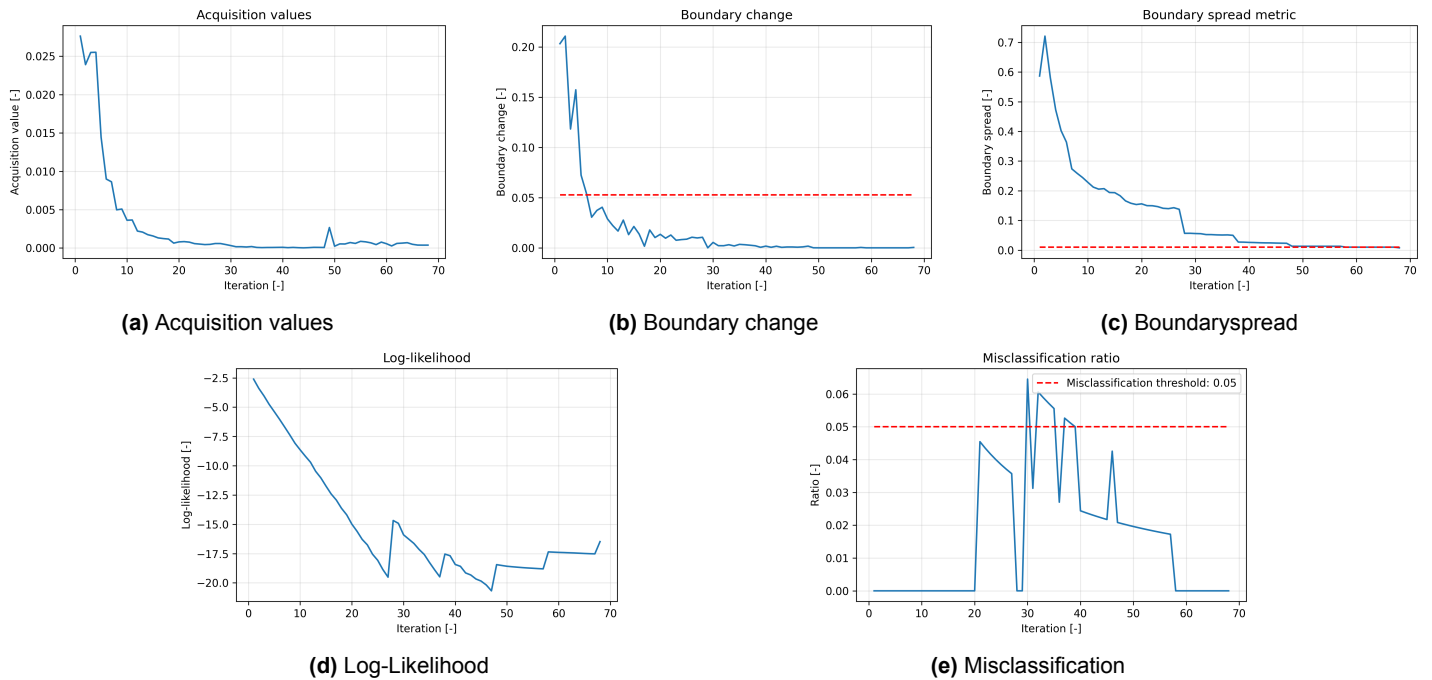


Figure O.4: Monitoring metrics 2D (t_h, ρ) material variables

3D Surrogate Model with Material Variables (t_h, ρ, E)

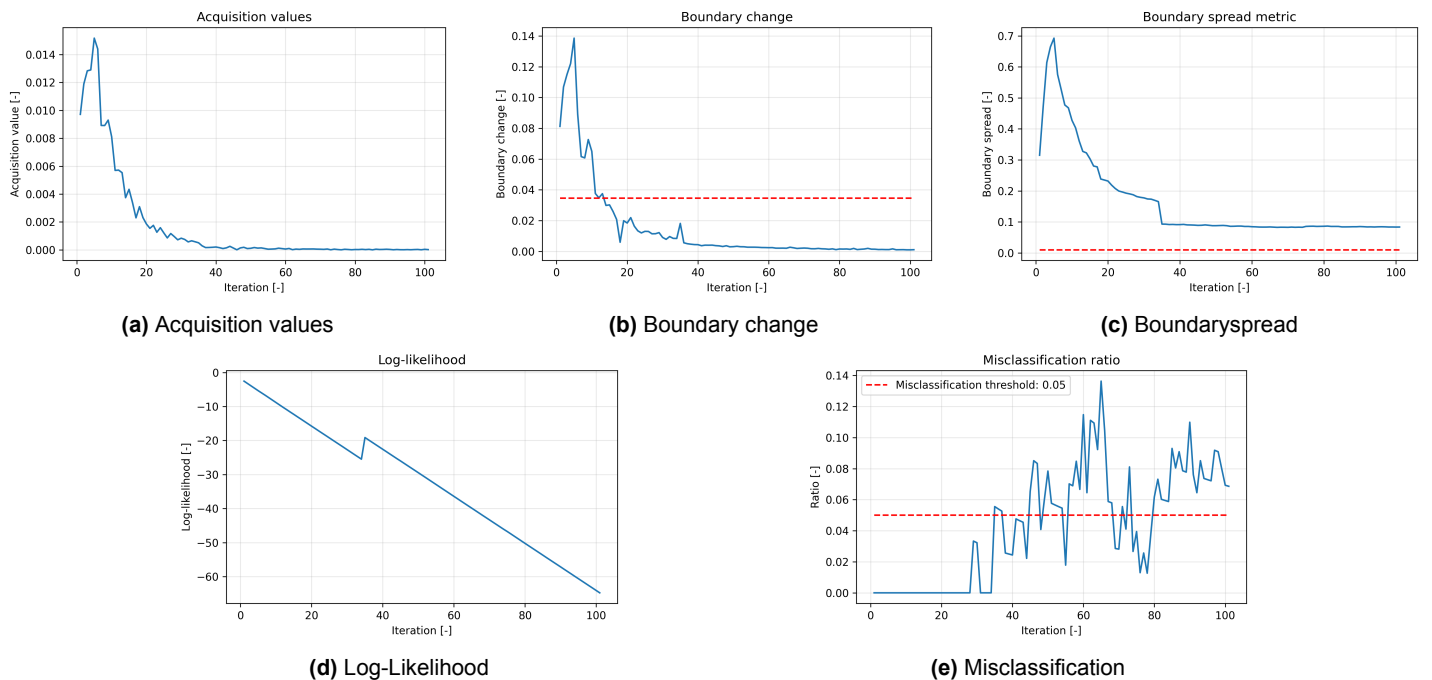


Figure O.5: Monitoring metrics 3D (t_h, ρ, E) material variables

6D Surrogate Model with Material Variables ($t_h, \rho, E, B, f_c, e_0$)

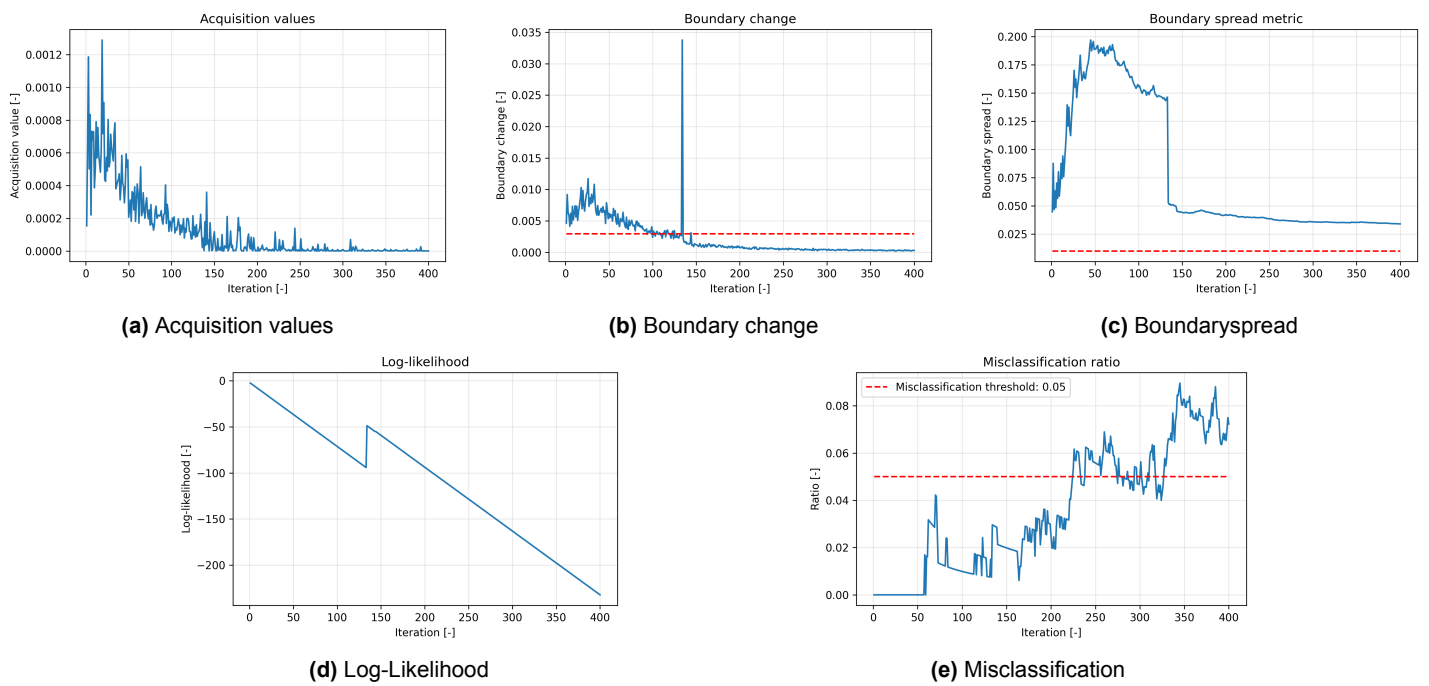


Figure 0.6: Monitoring metrics 6D ($t_h, \rho, E, B, f_c, e_0$) material variables