ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DIPARTIMENTO DI INGEGNERIA CIVILE, CHIMICA,
AMBIENTALE E DEI MATERIALI

Master of Science in Chemical and Process Engineering - Curriculum S.T.E.M.

"OPTIMIZED REACTION MECHANISM FOR HYDROGEN-AMMONIA COMBUSTION"

Candidate: Pietro Malpezzi

Supervisor: Prof. Ernesto Salzano

Co-supervisors: Prof. Alessandro Parente

Dr. Yuki Murakami

Academic Year 2024/2025

Session V

Contents

List of Figures

List of Tables

List of Equations

Chapter 1

Introduction

The transition toward climate–neutral energy systems has intensified the need for clean, reliable, and efficient fuels suitable for both stationary and mobile power applications. Growing environmental concerns related to fossil fuel combustion, combined with increasingly stringent emission regulations, have motivated extensive research into carbon–free or carbon–neutral energy carriers. Hydrogen (H_2) is widely regarded as one of the most promising alternatives due to its high reactivity and the absence of carbon-based pollutants. However, challenges related to storage, transport, and safety hinder its large-scale deployment. Despite these limitations, hydrogen chemistry plays a central role in combustion science: the H/O radical pool governs the reactivity of a broad spectrum of fuels, including hydrocarbons, oxygenated fuels, and nitrogen-containing species [3].

Ammonia (NH_3) has recently emerged as an attractive carbon–free fuel and hydrogen carrier [4]. Its mature production chain, favorable storage characteristics, and high hydrogen density make it particularly suitable for large-scale energy transport and long-term storage. Blends of H_2 and NH_3 combine the advantages of both fuels: hydrogen enhances the low reactivity of ammonia, while ammonia mitigates the handling and safety challenges associated with hydrogen. As a result, H_2–NH_3 mixtures represent a promising solution for zero-carbon combustion technologies.

Nevertheless, ammonia combustion remains complex. Pure NH_3 exhibits low flame speed, high ignition temperature, and narrow flammability limits, hindering its ignition and stabilization. Furthermore, ammonia oxidation is strongly coupled to NO_x formation, requiring accurate modelling of N-chemistry to meet environmental targets. Despite the recent years progress, existing detailed kinetic mechanisms show significant discrepancies when predicting key combustion characteristics of H_2–NH_3 systems, such as ignition delay times, laminar flame speeds, and species profiles. These differences arise from uncertainties in elementary reaction rates and from incomplete knowledge of coupled H/N/O pathways.

As modern combustion simulations increasingly rely on high-fidelity kinetic models, quantifying and reducing uncertainties in Arrhenius parameters has become essential. Even small deviations in selected rate constants can propagate into large discrepancies in predicted ignition behaviour or pollutant formation. For this reason, uncertainty analysis and mechanism optimisation have become central components in kinetic model evaluation. Global optimisation approaches based on genetic algorithms (GAs) have proven particularly effective in this context [5, 6]. Their ability to explore non-linear, high-dimensional, and highly coupled parameter spaces without relying on gradient information makes them well suited to the challenge of kinetic parameter refinement.

In this work, a computational framework is developed to optimise selected kinetic parameters of a detailed H_2–NH_3 mechanism using the PyGAD genetic algorithm library coupled with Cantera. The workflow integrates sensitivity analysis, uncertainty quantification, gene-space formulation, feasibility constraints, and validation against reference datasets. The methodology aims to improve the predictive accuracy of existing mechanisms while maintaining physical consistency and computational tractability.

1.1 Scope and Objectives

The objective of this thesis is to optimise selected elementary reaction parameters within a detailed H_2–NH_3 kinetic mechanism using a genetic-algorithm-based framework. The aim is to improve the mechanism's capability to reproduce key combustion characteristics of hydrogen–ammonia mixtures over a broad range of conditions.

Accurate kinetic models are essential tools for predicting the effect of modifying fuels, operating conditions, or reactor configurations. In nitrogen-containing fuel systems, ammonia represents the simplest molecular structure and provides the foundation for understanding the broader N-chemistry. Because hydrogen chemistry forms the radical backbone for nearly all combustion systems, achieving robust predictions for H_2–NH_3 mixtures requires consistent treatment of both H/O and H/N/O reaction pathways.

The specific objectives of the present work are as follows:

- To perform a detailed sensitivity analysis to identify the most influential reactions affecting combustion targets for H_2–NH_3 mixtures.
- To define uncertainty bounds for selected Arrhenius parameters and incorporate them into a physically consistent search space.

- To develop an efficient gene-space formulation inspired by the REPRICA methodology, ensuring both flexibility and physical feasibility during optimisation.
- To apply a Genetic Algorithm implemented in PyGAD to optimise pre-exponential factors, temperature exponents, and activation energies of selected reactions.
- To validate the optimised mechanism against independent datasets, evaluating improvements in prediction accuracy and assessing limitations.

The overall goal is to establish a reproducible and computationally efficient workflow for kinetic model optimisation, contributing to a more reliable understanding of H_2–NH_3 combustion behaviour.

1.2 Thesis Outline

The thesis is structured as follows:

- **Chapter 2** reviews the state of the art on H_2–NH_3 combustion chemistry, existing kinetic mechanisms, reaction rate formulations, uncertainty quantification, and optimisation techniques relevant to mechanism refinement.
- **Chapter 3** describes the methodology adopted in this work, including the mechanism structure, sensitivity analysis criteria, uncertainty treatment, gene-space formulation, feasibility checks, and implementation details of the GA–Cantera workflow.
- **Chapter 4** presents the results obtained from sensitivity studies, uncertainty evaluation, optimisation, gene-space analysis, and model validation.
- **Chapter 5** summarises the main findings of the thesis, discusses limitations of the present approach, and outlines perspectives for future research.

Chapter 2

State of the Art

Ammonia has gained renewed interest as a carbon-free energy carrier because it enables hydrogen storage, transport, and combustion without introducing CO_2. However, its practical use as a fuel is limited by intrinsically slow reactivity, low laminar burning velocity, and complex nitrogen chemistry that promotes NO and N_2O formation. Blending ammonia with hydrogen has emerged as an effective strategy to overcome these kinetic limitations, but the resulting coupled H_2/NH_3 chemistry introduces new reaction pathways and sensitivities not captured by classical nitrogen or hydrogen submechanisms alone. As a result, the development, validation, and reduction of kinetic mechanisms for H_2–NH_3 mixtures remains an active research area, with numerous studies providing partially overlapping but sometimes contradictory findings [7].

2.1 H_2-NH_3 combustion chemistry background

This section examines how operating conditions—pressure, temperature, equivalence ratio, hydrogen fraction, and oxygen concentration—influence NH_3/H_2 kinetics, with particular emphasis on the formation pathways and diagnostic significance of NO and N_2O intermediates.

Ammonia oxidation proceeds primarily through successive hydrogen-abstraction reactions forming NH_x radicals (NH_2, NH, N), followed by branching through N_2H_2 and NNH intermediates. These pathways are highly sensitive to temperature and pressure and several key reactions involve third-body stabilization (e.g., hydrazine formation, $NH_2 + NH_2 + M \rightarrow N_2H_4 + M$) [2]. Hydrogen oxidation is significantly faster and dominated by the classical H/O system, with chain carriers H, O, and OH controlling ignition, radical buildup, and heat release [8]. When hydrogen is added to ammonia, the fast H/O subsystem strongly perturbs the slower nitrogen chemistry [9], resulting in:

- Accelerated ignition
- Increased radical pool (OH, O, H)
- Enhanced flame speed and reactivity
- Altered NO and N_2O formation rates.

The coupling of these subsystems is non-linear and condition dependent, which explains why no single kinetic mechanism performs robustly across all temperatures, pressures, and mixture compositions.

NH_x and N_2H_2 pathways

Under fuel-rich conditions, or in regimes where the H/O system is not yet fully activated (e.g. low temperature, moderate pressure, or weakly igniting mixtures), the nitrogen sub-mechanism becomes dominated by NH_x and N_2H_2-forming pathways. The conversion of NH_3 proceeds mainly through successive hydrogen abstractions:

$$NH_3 \rightarrow NH_2 \rightarrow NH \rightarrow N$$

followed by recombination and isomerization reactions that produce hydrazine (N_2H_4), diazene (N_2H_2), and the key intermediate NNH. These reactions are well documented as major branching points in ammonia oxidation [2]:

$$\begin{aligned} NH_2 + NH &\rightarrow N_2H_2 + H \\ NH_2 + NH_2 &\rightarrow N_2H_2 + H_2 \\ N_2H_2 + H &\rightarrow NNH + H_2 \end{aligned}$$

These pathways are important for two reasons.

First, they regulate the partitioning of fuel nitrogen toward either molecular nitrogen (the benign, desired route) or oxidized nitrogen species. The NNH intermediate in particular has a dual role. It can decompose thermally,

$$NNH \rightarrow N_2 + H$$

which increases radical concentrations and contributes to the hydrogen-recovery cycle. Alternatively, it can be oxidized:

$$NNH + O_2 \rightarrow N_2 + HO_2$$

which competes with chain branching and redirects flux away from reactive radicals toward HO_2, a less reactive species.

Second, hydrogen addition strongly modifies these NNH/N_2H_2 pathways.
A larger H-radical pool promotes the thermal decomposition of NNH, accelerating the formation of H atoms and therefore increasing flame reactivity [10]. Conversely, an increase in HO_2 via hydrogen peroxide chemistry may enhance the oxidative destruction of NNH, thereby suppressing its contribution to chain propagation. The net effect depends sensitively on the interplay among temperature, equivalence ratio, and pressure.

Because N_2H_2 and NNH are precursors not only to benign N_2 but also to NO and N_2O through competing NH + NO and HNO pathways, accurate representation of these intermediates is essential for predicting both ignition behavior and pollutant formation in H_2-NH_3 systems [11, 10]. In modern mechanisms, errors in the rate constants of NNH formation, decomposition, and oxidation are among the dominant sources of uncertainty in NO and N_2O predictions.

Overall, the NH_x–N_2H_2 subsystem forms the backbone of the slow nitrogen chemistry of ammonia flames. Hydrogen enrichment does not eliminate this chemistry; instead, it reshapes it by altering radical availability, changing relative branching ratios, and shifting the flux distribution between pathways that either promote or suppress reactivity. This complex coupling explains why the predictive accuracy of kinetic models in rich H_2/NH_3 mixtures depends critically on how they represent the interlinked NH_x, NNH, HNO, and N_2H_2 reaction families.

Hydrogen-dominated chain branching

The addition of hydrogen profoundly alters the early-stage ignition and flame-propagation kinetics of H_2-NH_3 mixtures because it introduces the classical H/O chain-branching subsystem. At sufficiently high temperatures, or whenever the mixture contains a non-negligible fraction of H_2, the kinetic response becomes dominated by reactions that rapidly generate and recycle the key radicals H, O, and OH [8]. The most influential steps include:

$$\begin{aligned} H + O_2 &\rightarrow O + OH \\ H + O_2 + M &\rightarrow HO_2 + M \\ H_2O_2 &\rightarrow 2OH \end{aligned}$$

The first reaction is the classical chain-branching step that controls ignition delays in hydrogen flames. The HO_2 and H_2O_2 pathways dominate under lower-temperature or

higher-pressure conditions, providing a secondary source of OH through peroxide decomposition. Collectively, these reactions generate a radical pool that is several orders of magnitude larger than that produced by ammonia oxidation alone [2].

This elevated concentration of OH is particularly important because OH is the principal abstracting radical responsible for initiating ammonia oxidation through:

$$NH_3 + OH \rightarrow NH_2 + H_2O.$$

Since ammonia decomposition is rate-limited by hydrogen abstraction, increasing the OH pool directly accelerates the formation of NH_2 and downstream fuel-nitrogen intermediates (NH, HNO, NNH). As several studies have shown [12, 7], the sensitivity of NO formation to hydrogen enrichment arises predominantly from this acceleration of NH_x conversion.

A larger OH and H radical reservoir shifts the balance between NO-forming and NO-reducing channels, because:

- reactions that produce NO (e.g. $NH_2 + O \rightarrow HNO + H$) scale favorably with O and OH, while
- reactions that consume NO (e.g. $NH_2 + NO \rightarrow N_2 + H_2O$) may become relatively less competitive.

Therefore, the hydrogen-induced enhancement of chain branching is not only a driver of global flame reactivity, but also the primary mechanism by which the nitrogen subnetwork is perturbed. This effect explains why NO formation increases in many conditions with rising H_2 fraction, but decreases in others: the hydrogen-driven radical pool modifies the detailed flux balance among NH_x, HNO, NO, and N_2O in a strongly temperature- and pressure-dependent manner [2].

Relevance of NO and N_2O in NH_3/H_2 Combustion

Ammonia combustion chemistry is fundamentally distinct from hydrocarbon oxidation due to the presence of nitrogen-containing intermediates, most notably nitric oxide (NO) and nitrous oxide (N_2O). These species are not merely pollutants but active participants in the reaction network, governing the branching between chain-propagating and chain-terminating pathways. Four key factors underscore their central role:

1. **Reaction network complexity and multi-pathway branching.**

Unlike hydrocarbons, where oxidation proceeds mainly through well-defined H-abstraction and radical addition sequences, ammonia oxidation involves multiple competing pathways

for nitrogen atom redistribution. The conversion of NH_3 to molecular nitrogen (N_2) proceeds through a cascade of NHx radicals (NH_2, NH, N), with NO and N_2O acting as critical branch points. For instance, the reaction

$$NH_2 + NO \rightarrow N_2 + H_2O$$

serves as a major chain-terminating step under certain conditions, significantly quenching radical activity [2]. This pathway competes with radical-driven chain-branching routes, making the balance between NO formation and consumption a key determinant of overall reactivity.

2. NO and N_2O as reactive intermediates, not inert products.

In NH_3/H_2 flames, NO is not only a combustion product but a reactive intermediate that strongly influences the branching of nitrogen chemistry. Reaction paths such as

$$NH + NO \rightarrow N_2O + H$$

route nitrogen toward N_2O rather than direct N_2 formation, altering the overall oxidation balance [11]. Because NO and N_2O formation competes directly with radical-driven chain-branching and termination pathways, their concentrations govern the global reactivity, flame structure, and pollutant output.

3. Strong sensitivity to hydrogen addition and radical pool dynamics.

The addition of hydrogen dramatically perturbs the H/O radical subsystem, leading to elevated concentrations of H, O, OH, and HO_2. These radicals accelerate NH_3 conversion (e.g., via $NH_3 + OH \rightarrow NH_2 + H_2O$) and shift the relative flux between NO-forming and NO-reducing branches. As a consequence, NO concentrations may increase or decrease depending on equivalence ratio ϕ, pressure, and temperature. N_2O formation, on the other hand, tends to decrease at high temperatures and high radical densities because the destruction reaction

$$N_2O + H \rightarrow N_2 + OH$$

becomes highly favored under those conditions [11], despite active formation pathways remaining operational. This differential sensitivity means that NO and N_2O behave quite differently when hydrogen fraction changes, making N_2O a critical diagnostic of the underlying nitrogen–radical coupling.

4. Diagnostic power for mechanism validation and development.

Because NO and N_2O concentrations integrate over many elementary reactions (involving NHx, HNO, NNH, N_2H_2, etc.), they serve as rigorous markers for the overall correctness of a kinetic mechanism. Many of the most comprehensive mechanism assessments use NO and N_2O as primary validation targets against shock-tube, flame, or reactor experiments [13, 11]. A mechanism that reproduces flame speed and ignition delay times well but fails to predict NO/N_2O profiles typically exhibits incorrect branching ratios or radical flux imbalances, rendering it unreliable for real-world applications.

In summary, a realistic and credible kinetic model for NH_3/H_2 combustion must explicitly include a well-validated NO/N_2O sub-mechanism, because these species are not only environmental pollutants but also central players in the radical and nitrogen reaction network. Their behavior under varying conditions dictates whether ammonia combustion can be considered clean, stable, and efficient.

Influence of Operating Conditions on NH_3/H_2 Kinetics

The performance of ammonia–hydrogen combustion depends strongly on the thermodynamic and compositional environment. Because nitrogen chemistry proceeds via relatively slow hydrogen-abstraction and recombination steps, any factor that significantly modifies radical concentrations, third-body collision efficiencies, or radical pool lifetimes may dramatically alter ignition delays, laminar flame speed, and NOx/N_2O formation. Below, the most critical operating parameters are discussed, supported by representative findings from recent literature.

Pressure

Pressure exerts a major influence on three-body recombination and radical termination reactions. In high-pressure experiments (e.g., rapid compression machines), recombination processes such as

$$NH_2 + NH_2 + M \rightarrow N_2H_4 + M$$

and other third-body-mediated stabilization steps become significantly more efficient as pressure increases, accelerating radical recombination and altering the chain-branching balance [11].

For example, data from RCM experiments on $NH_3/H_2/O_2$ mixtures at pressures between 20 and 60 bar show that increasing H_2 mole fraction leads to substantial reductions in ignition delay times under rich conditions. This suggests that mechanisms which include accurate pressure-dependent reaction rates (including third-body efficiencies and pressure-

falloff behavior) are better suited for high-pressure applications, whereas low-pressure-tuned mechanisms can over-predict ignition delays by more than 30% [13].

Temperature

Temperature has a dual role: it accelerates both chain-branching and radical termination while also shifting the relative importance of distinct reaction pathways. Shock-tube and laminar-flame measurements indicate a transition behavior around a temperature threshold (typically ∼1200 K, though case-dependent): below this threshold, NH_x-dominated pathways (hydrogen abstraction, radical recombination) govern; above it, classical H/O chain branching dominates (e.g., $H + O_2 \rightarrow O + OH$), substantially increasing radical production and accelerating oxidation [7].

In this context, recent measurements of laminar flame speeds for NH_3/H_2/air mixtures with hydrogen volume fractions up to 80% showed that the flame propagation speed increases significantly with increasing preheating temperature and H_2 fraction [14]. For instance, the flame propagation speed increased by a factor of 7.8 when the hydrogen volume fraction was raised from 20% to 80% at 673 K under stoichiometric conditions.

This proves that kinetic mechanisms must include robust, temperature-dependent Arrhenius parameters, especially for radical chain-branching, termination, H-abstraction, and nitrogen-radical decomposition reactions.

Equivalence Ratio (ϕ)

The equivalence ratio impacts reactivity and emissions in a more complex manner than for hydrocarbon fuels. In premixed NH_3/H_2/air flames, experiments with a 70/30 vol% NH_3/H_2 blend over a wide ϕ range (0.55–1.4) showed that laminar burning velocities and NO emissions peak around $\phi \approx 0.8$, decreasing at richer conditions [9]. This behavior contrasts with typical hydrocarbon fuels and reflects the competing effects of radical availability, radical termination (via recombination or third-body reactions), and nitrogen-radical chemistry.

At rich equivalence ratios, radical consumption and radical-pool depletion (via recombination and termination) are more pronounced, while formation of nitrogen intermediates (e.g., N_2H_2, NNH) can act as radical sinks and slow down reactivity. The result is a non-monotonic dependence of flame speed and NOx/N_2O emissions on ϕ, a trait unique to NH_3/H_2 combustion compared to hydrocarbon systems [9].

Hydrogen Fraction

The fraction of hydrogen in the fuel blend strongly controls the radical pool and thus overall reactivity. A recent systematic study observed that adding H_2 markedly enhanced laminar flame speed under engine-like conditions for NH_3/H_2/air mixtures, illustrating how hydrogen injection mitigates NH_3's poor flame propagation characteristics [1].

Figure 2.1 illustrates the effect of hydrogen fraction on (a) ignition delay time and (b) combustion duration for NH_3/H_2/air mixtures. As shown in panel (a), ignition delay decreases exponentially with increasing hydrogen fraction, consistent with enhanced radical production via $H + O_2$ chain-branching. Combustion duration (panel b) exhibits a similar trend, confirming that hydrogen addition accelerates both ignition and flame propagation.

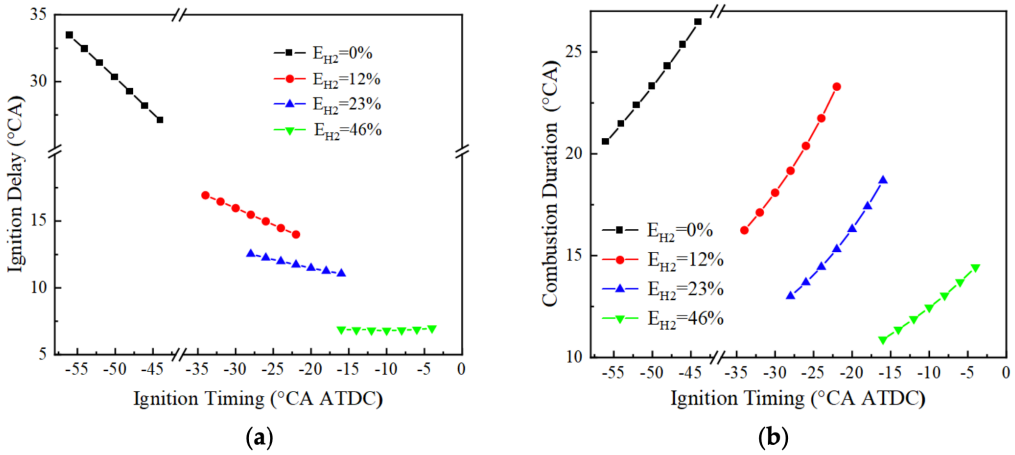

Figure 2.1. Effect of hydrogen fraction on (a) ignition delay time and (b) combustion duration for NH_3/H_2/air mixtures. Data from [1].

The H_2 fraction also affects pollutant formation: increased radical production initially favors NOx formation (via HNO and NH-radical pathways), but may also facilitate radical termination and NOx reduction (e.g., via recombination or N_2O pathways) at very high H_2 fractions. Recent reviews indicate that NOx emissions as a function of H_2 fraction show a non-monotonic trend: emissions increase up to moderate H_2 fractions, then decrease when the radical pool becomes large enough to favor termination routes or when NH_3 is sufficiently diluted [15].

Oxygen Enrichment

Lean vs. Rich Effects on NOx and N_2O. The effects of oxygen concentration, often explored through variations in the equivalence ratio ϕ, significantly influence NH_3/H_2 flame chemistry and emissions. Lean conditions, which imply excess oxygen, generally

lead to higher NOx formation compared to rich conditions [7, 9]. Enhancement of chain-branching reactions involving oxygen, such as

$$H + O_2 \rightarrow O + OH,$$

is consistent with studies showing that H-atom chemistry dominates the laminar burning velocity (LBV) in lean combustion. The addition of hydrogen, which increases overall reactivity, strongly enhances flame speed, particularly under lean conditions where O-atom availability is high [15, 1].

With respect to NOx yields, experimental observations indicate that NO production peaks near stoichiometric conditions ($\phi \approx 0.8$–0.9) and decreases as the mixture becomes richer, due to depletion of the radical pool by termination reactions. Furthermore, N_2O emissions are often observed under very lean conditions ($\phi < 0.7$) [13].

Modeling Challenges and Data Gaps. Although numerical studies show that oxygen enrichment increases reactivity, this elevated activity requires careful modeling of radical-termination pathways to avoid over-prediction of NOx formation. Experimental data at high oxygen concentrations for NH_3/H_2 blends remain limited [11], highlighting a need for further validation studies under oxygen-enriched conditions.

Having established the sensitivity of NH_3/H_2 kinetics to operating conditions, we now review the evolution of detailed kinetic mechanisms developed to capture these dependencies across a wide range of combustion regimes.

2.2 Kinetic mechanisms for H_2-NH_3 systems

A chemical-kinetic mechanism is a mathematically closed description of all elementary reactions that interconvert the molecular species participating in a reacting flow. Each reaction is expressed by a modified Arrhenius rate law:

$$k \;=\; A\,T^{n}\exp\!\left(-\frac{E_a}{RT}\right) \tag{2.1}$$

where:

- A is the pre-exponential factor [$cm^3/(mol{\cdot}s$, for bimolecular reactions)],
- n is the temperature exponent (dimensionless),
- E_a is the activation energy [cal/mol],

- R is the universal gas constant [cal/(mol·K)],
- T is temperature [K].

These three constitute the kinetic parameters. By integrating the coupled ordinary-differential equations for species mass fractions and temperature, the mechanism predicts macroscopic combustion observables such as laminar-burning velocity (LBV), ignition-delay time (IDT), and pollutant formation (NO, NO_2, N_2O). Because the combustion of ammonia–hydrogen blends involves both H/OH-radical chemistry and nitrogen-specific pathways (e.g., HNO, NNH, N_2O), a successful mechanism must contain a balanced sub-mechanism for each domain and must be validated against a broad set of experiments covering lean, stoichiometric, and rich regimes.

Three main strategies characterize current H_2/NH_3 mechanism development:

Updated detailed mechanisms: it consists in the refinement of existing models through updated rate constants and expanded validation datasets.

Systematic reduction of detailed mechanisms: by identifying and removing chemical species and reactions that have only a minor influence on the system's behavior. The result is a significantly smaller and faster chemical mechanism, suitable for large-scale simulations such as computational fluid dynamics or real-time modeling, while still preserving the essential physical and chemical characteristics of the original detailed mechanism.

New mechanisms built from first principles: mechanisms formulated specifically for high-pressure, gas-turbine, or autoignition conditions, typically validated against extensive datasets including laminar flame speeds, shock tube data, and NO_x emissions.

Despite substantial progress, comparative studies conclude that no mechanism exhibits universal accuracy [7]. Performance varies depending on: pressure–temperature regime, hydrogen fraction, equivalence ratio, targeted property (ignition delay, flame speed, NO_x prediction), degree of mechanism reduction. These limitations form the primary motivation for mechanism optimization methods such as genetic algorithms, which can tune reaction parameters to improve agreement with selected validation targets while maintaining computational tractability.

Historical Development of NH_3–H_2 Kinetic Mechanisms

To contextualize the development of modern NH_3/H_2 mechanisms, a brief review is proposed on the historical evolution from early nitrogen chemistry models to contemporary

frameworks incorporating high-level theoretical rate constants.

The development of NH_3–H_2 kinetic mechanisms has progressed through several generations of models. The earliest systematic effort dates to the 1980s, when Miller and Bowman introduced a 98-reaction, 22-species mechanism establishing the canonical $NH_3 \rightarrow NH_2 \rightarrow NH \rightarrow N$ oxidation sequence and the foundational DeNOx routes (e.g., $NH_2 + NO \rightarrow NNH + OH$) [7].
Subsequent refinements by Miller & Bowman (1989) and Miller & Glarborg (1999) expanded the nitrogen sub-mechanism to include prompt-NO chemistry, HNO and N_2O intermediates, and additional NHx pathways [16, 17].

During the 2000s, large hydrocarbon oxidation mechanisms such as those by Konnov (2009) and Dagaut (2008) incorporated nitrogen chemistry as an add-on, producing highly detailed mechanisms containing 1200–1400 reactions and 120–150 species [18, 19]. Although not specifically optimized for NH_3/H_2 flames, these mechanisms formed the basis for many subsequent comparisons.

A major advance came with the comprehensive nitrogen mechanism of Glarborg et al. (2018) [2], which includes approximately 1400 reactions and 151 species. This mechanism integrates updated rate coefficients from high-level theoretical studies (notably by Klippenstein and collaborators) into a unified framework.

Pathway	**Representative Reactions**
NH oxidation	$NH + O \rightarrow NO + H$
	$NH + OH \rightarrow N + H_2O$
	$NH + O_2 \rightarrow HNO + O$
HNO oxidation	$HNO + H \rightarrow NO + H_2$
	$HNO + O \rightarrow NO + OH$
	$HNO + OH \rightarrow NO + H_2O$
Zeldovich thermal-NO	$O + N_2 \rightarrow NO + N$
	$N + O_2 \rightarrow NO + O$
	$N + OH \rightarrow NO + H$[1]
NNH route	$NNH \rightarrow N_2 + H$
	$NNH + O \rightarrow NO + NH$
	$NNH + H \rightarrow NH + NH$
	$NNH + OH \rightarrow HNO + NH$

Table 2.1. Key NO-forming pathways in NH_3/H_2 combustion [2]

It consistently reproduces the dominant NO-forming pathways, which are summarized in Table 2.1.

[1]This reaction has minor contribution at moderate temperatures but is included for completeness.

More recent developments have focused on evaluating and tailoring detailed mechanisms for H_2-rich ammonia flames. Alnasif et al. (2023) assessed 67 published *complete kinetic mechanisms* against burner-stabilized stagnation-flame measurements of NO formation for 70% NH_3/30% H_2 mixtures at $\phi = 0.6$–1.4 [13]. Among the tested mechanisms, the Glarborg model provided the closest overall agreement with the measured NO profiles across lean and stoichiometric conditions, accurately capturing the interplay between NH_2, HNO, and NH pathways.

A parallel study by Alnasif et al. (2024) [11] examined N_2O formation and again identified the Glarborg mechanism, augmented with Klippenstein's updated rate coefficients, as the most reliable across $\phi = 0.6$–1.4. Its performance deteriorates only in the very lean limit ($\phi \approx 0.57$), where kinetic inhibition prevents ignition.

Taken together, these evaluations show that while the Glarborg (2018) framework remains the most consistently validated detailed nitrogen mechanism, several recent optimizations, particularly those incorporating state-of-the-art rate-constant calculations, offer improved predictive robustness under specific flame conditions. In this regard, the mechanism of Stagni et al. (2023) provides an effective balance between completeness, numerical tractability, and accuracy for NH_3/H_2 flames relevant to the present study [20], and will therefore be adopted as the baseline mechanism for optimization in Chapter 3.

2.3 Reaction rate expression and parametrisation

Before presenting the optimization methodology (Chapter 3), we formalize the mathematical representation of reaction rates, as this directly determines the parametrization strategy used in genetic algorithm-based tuning.

Chemical kinetic mechanisms for the NH –H system rely on a set of temperature- and pressure-dependent rate expressions that encode the behaviour of elementary reactions, association reactions, and dissociation channels. These expressions are built around the Arrhenius law and its pressure-dependent extensions (falloff and logarithmic interpolation). Because the optimization strategy adopted in this thesis directly modifies selected kinetic parameters, it is essential to describe in detail the mathematical form of the reaction rate expression and the rationale behind their parametrisation.

In chemical kinetics, the rate of a reaction is expressed as the *rate coefficient* $k(T)$, which typically follows the modified Arrhenius formulation, described in Eq. 2.1:

$$k(T) = A\,T^{n}\,\exp\left(-\frac{E_a}{RT}\right)$$

The pre-exponential factor A represents the frequency and geometric probability of successful molecular collisions, encompassing both the collision rate (proportional to temperature via kinetic theory) and the steric factor that accounts for the relative orientation of reactants. For bimolecular gas-phase reactions, A typically ranges from 10^{10} to 10^{14} $\mathrm{cm^3/(mol{\cdot}s)}$, with lower values indicating sterically hindered reactions (e.g., those requiring specific molecular alignment) and higher values corresponding to barrierless or nearly barrierless processes where almost every collision leads to reaction (e.g., radical recombinations). For unimolecular dissociations, A is expressed in $\mathrm{s^{-1}}$ and reflects the vibrational frequency of the breaking bond, typically ranging from 10^{13} to 10^{16} $\mathrm{s^{-1}}$.

The exponent n accounts for the temperature dependence of the collision frequency and activation entropy, reflecting deviations from the classical Arrhenius form ($n = 0$). For elementary reactions, n typically ranges from -3 to $+3$, with negative values indicating that the reaction rate decreases with increasing temperature (e.g., barrierless radical recombinations) and positive values indicating enhanced temperature sensitivity beyond simple exponential activation.

The activation energy E_a quantifies the minimum energy barrier that reactant molecules must overcome for the reaction to proceed, arising from the need to break or weaken existing bonds before new bonds can form. For elementary reactions, E_a typically ranges from near-zero (for exothermic radical recombinations or barrierless addition reactions) to over 70 000 cal/mol (for highly endothermic bond-breaking processes). A low activation energy ($E_a \lesssim 5\,000$ cal/mol) indicates a kinetically fast reaction with weak temperature dependence, while a high activation energy ($E_a \gtrsim 25\,000$ cal/mol) results in strong temperature sensitivity, where small increases in temperature lead to exponential increases in reaction rate. For NH_3/H_2 combustion, the dominant hydrogen chain-branching reaction $H + O_2 \rightarrow OH + O$ exhibits $E_a \approx 16\,800$ cal/mol, reflecting its moderate barrier and strong temperature dependence in ignition phenomena.

Direct optimisation of (A, n, E_a) is notoriously challenging due to the disparity in magnitudes (e.g., A may vary over 20–30 orders of magnitude), the strong coupling between parameters, and the necessity to enforce physical constraints such as $A > 0$. To mitigate these issues, the parameters are transformed into a more numerically stable set through a logarithmic transformation of the modified Arrhenius rate Law:

$$\log k(T) = \log A + n \log T - \frac{E_a}{RT} = \alpha + n \log T - \frac{\epsilon}{T} \tag{2.2}$$

Here, the optimisable parameters are redefined as:

- $\alpha = \log A$
- n remains the temperature exponent
- $\epsilon = E_a/R$

This transformation linearizes the Arrhenius equation with respect to $\ln k(T)$ and facilitates optimization by:

Compressing the dynamic range: The logarithmic scale maps $A \in (10^{-30}, 10^{30})$ to $\alpha \in \mathbb{R}$, making numerical optimization more stable.

Decoupling parameters: The transformation separates the temperature exponent n from the activation energy term ϵ/T, reducing parameter correlation during optimization.

Enabling bounded search spaces: Gene space bounds (Section 2.5) can be defined symmetrically around the nominal values in (α, n, ϵ) space, which is more natural for evolutionary algorithms than exponentially scaled (A, n, E_a) bounds.

In Chapter 3, we describe how these transformed parameters (α, n, ϵ) are systematically varied within physically constrained bounds, derived using the methodologies reviewed in Section 2.5 (Bertolino and Fürst), to minimize prediction errors against experimental targets. The genetic algorithm searches this transformed parameter space to identify optimal rate constant combinations that improve mechanism fidelity while respecting thermochemical consistency.

2.4 Sensitivity and Uncertainty Analysis

Sensitivity analysis is a fundamental tool in chemical kinetics, allowing us to quantify how variations in reaction parameters, such as rate constants, influence key outcomes of a reaction system, including ignition delay times, species concentrations, and flame speed. By systematically perturbing parameters and observing the resulting changes, sensitivity analysis provides insights into which reactions dominate the overall behavior of the system. In the context of kinetic model development, particularly for complex combustion systems such as H_2/NH_3 mixtures, sensitivity analysis serves several purposes:

- **Model validation and refinement:** Identifying reactions with the greatest impact helps focus experimental and theoretical efforts on accurately determining their rates.

- **Mechanism reduction:** Less sensitive reactions may be approximated or removed to simplify the model without sacrificing predictive accuracy.
- **Optimization and control:** In combustion optimization, identifying the most sensitive reactions guides the adjustment of conditions or catalyst design to achieve desired performance, such as minimal ignition delay or efficient NH_3 conversion.

By normalizing the effect of each reaction on an output of interest, sensitivity analysis establishes a hierarchy that prioritizes reactions for experimental validation or mechanistic refinement, a critical step in the development of robust, predictive kinetic models [21].

Sensitivity Analysis Approaches

Sensitivity analysis encompasses a variety of methodological frameworks designed to evaluate how sensitive a combustion model is to changes in its kinetic parameters. These approaches differ in computational cost, the regions of parameter space they explore, and the type of insight they provide. For detailed combustion kinetic mechanisms, such as those governing H_2/NH_3 mixtures, the most relevant classes are *local*, *global*, and *data-driven* methods.

Local (Differential) Sensitivity Analysis

Local sensitivity methods examine how a model output Y (e.g., ignition delay, flame speed, species concentration) responds to an infinitesimal change in a kinetic parameter k_i.
The normalized first-order sensitivity coefficient is commonly defined as:

$$S_{Y,k_i} = \frac{k_i}{Y}\frac{\partial Y}{\partial k_i} \tag{2.3}$$

This dimensionless coefficient quantifies the relative change in Y resulting from a relative change in k_i. For instance, consider the hydrogen chain-branching reaction:

$$H + O_2 \rightarrow OH + O \quad \text{(R5)}$$

If doubling the rate constant ($k_i \rightarrow 2k_i$, i.e., a +100% change) reduces the ignition delay time from 1000 μs to 500 μs (i.e., $Y \rightarrow 0.5Y$, a -50% change), the sensitivity coefficient is approximately:

$$S_{Y,k_i} \approx \frac{\Delta Y/Y}{\Delta k_i/k_i} = \frac{-0.5}{1.0} = -0.5$$

This negative value indicates that increasing the rate constant *decreases* ignition delay, confirming that R5 is a **chain-branching** reaction that accelerates ignition. Conversely, a positive sensitivity coefficient would indicate a **chain-terminating** reaction.

In practice, most combustion solvers approximate this derivative using finite differences. A rate constant is perturbed by a small factor (typically from 1% to 5%), and the resulting change in the observable is evaluated.

Local sensitivity analysis is computationally efficient and straightforward to perform. However, it probes only a narrow region around the baseline mechanism and therefore does not capture nonlinear parameter interactions or large deviations within the uncertainty range of the kinetic parameters.

Global Sensitivity Analysis

Global approaches investigate the effects of varying parameters over their full uncertainty intervals. Rather than applying small perturbations, reaction rate constants are sampled across broad ranges, often expressed through log-normal uncertainty distributions, and a large number of modified mechanisms are evaluated.

One widely adopted global technique is *Sobol variance decomposition* [22]. For an output Y, the total variance $\mathrm{Var}(Y)$ is decomposed into contributions from each parameter and their interactions. The resulting Sobol indices quantify:

- **First-order effects,** representing the direct influence of individual reactions;
- **Higher-order effects,** representing coupled interactions between reactions.

A computationally cheaper global alternative is the *Morris screening method* [23]. This technique approximates each parameter's influence by computing *elementary effects*, finite-difference estimates of the local gradient, along multiple randomized one-at-a-time (OAT) trajectories in parameter space. The distribution of these elementary effects provides a measure of both the mean effect (sensitivity) and the variance (indicating nonlinearity or interaction strength) for each parameter.

Global sensitivity methods thus capture nonlinearities, parameter couplings, and large-scale uncertainties, providing insight that local analysis alone cannot offer.

Data-Driven Methods: Principal Component Analysis (PCA)

When sensitivity analysis is performed across multiple conditions or for many outputs, the resulting datasets become high-dimensional and difficult to interpret directly.

Statistical tools such as **Principal Component Analysis (PCA)** facilitate a structured interpretation. Consider a sensitivity matrix:

$$\mathbf{S} \in \mathbb{R}^{M \times N},$$

where M represents the number of experimental conditions or observables (e.g., ignition delay times at different temperatures and pressures), and N represents the number of reactions or kinetic parameters under consideration. Each element S_{ij} quantifies the sensitivity of observable i to parameter j:

$$S_{ij} = \frac{k_j}{Y_i} \frac{\partial Y_i}{\partial k_j}$$

Applying PCA to $\mathbf{S}$ yields *principal components*, which represent orthogonal directions in parameter space along which the sensitivities vary most strongly. These components often reveal:

- **Correlated groups of reactions** that act together as functional subsystems;
- **Dominant sensitivity modes** that explain the majority of the variability;
- **Redundancies among reactions,** enabling mechanism reduction.

For instance, in H_2/NH_3 combustion, PCA may identify that reactions involving NH_2 oxidation (e.g., NH_2 + O, NH_2 + OH, NH_2 + HO_2) consistently co-vary across different conditions. This indicates that these pathways form a tightly coupled subsystem governing radical production and consumption. Meanwhile, hydrogen chain-branching reactions may appear in separate principal components, reflecting their distinct functional role.

PCA-based sensitivity analysis becomes indispensable when analyzing large-scale mechanisms with numerous reactions across many operating conditions. Under these circumstances, the resulting sensitivity matrix grows too large for manual interpretation, necessitating automated pattern recognition via statistical decomposition. For smaller mechanisms or when only a few conditions are studied, direct examination of individual sensitivities may suffice. While PCA was not employed in the present study due to time constraints, it represents a powerful tool for future mechanism development efforts, particularly when extending the analysis to broader temperature and pressure ranges or additional fuel blends.

Taken together, local, global, and data-driven sensitivity methodologies offer complementary perspectives. Local sensitivity analysis identifies critical reactions under specific operating conditions; global methods integrate nonlinearities and uncertainty effects across

the full parameter space; and PCA reveals deeper structural relationships among reactions. Employing these techniques yields a comprehensive understanding of parameter influence in H_2/NH_3 combustion kinetics and forms a robust foundation for the uncertainty quantification and optimization tasks addressed in the subsequent sections.

Uncertainty Factors for Chemical Kinetic Parameters

Uncertainty quantification of chemical kinetics is an established component of reaction mechanism development, evaluation, and reduction. A compact and widely used representation of kinetic uncertainty is the *uncertainty factor* associated with a rate coefficient. This multiplicative metric is widely adopted in major kinetic databases and mechanism-evaluation frameworks [24, 25]. The concept assumes that the true rate coefficient lies within a symmetric multiplicative interval around a nominal value k_0. This section summarizes the standard definition, the rationale for the logarithmic representation, methods used to derive uncertainty factors, and recommendations for their use in optimization workflows.

Definition and Log-Space Representation

Let k_0 be the nominal rate coefficient (typically taken from a reference mechanism or a recommended database evaluation). The uncertainty factor f_r for reaction r defines a multiplicative interval:

$$k_{\min} = \frac{k_0}{10^{f_r}}, \qquad k_{\max} = 10^{f_r}\, k_0, \tag{2.4}$$

as commonly applied in mechanism reviews and uncertainty evaluations [26, 27]. The interval $[k_0/10^{f_r},\ 10^{f_r} k_0]$ expresses the range within which the true rate coefficient is expected to lie, based on experimental scatter, theoretical predictions, or expert judgment.

Because chemical kinetic uncertainties span orders of magnitude, it is standard to work in logarithmic space [28, 29]. Defining $\kappa = \log_{10} k$, the uncertainty interval becomes:

$$\kappa \in [\kappa_0 - f_r,\ \kappa_0 + f_r]$$

with f_r representing the half-width in $\log_{10}$-space.
In this thesis, the uncertainty quantity f_r denotes the half-width in $\log_{10}$-space:

$$\log_{10} k \in [\log_{10} k_0 \pm f_r]$$

consistent with practices in recent mechanism optimization studies [30, 31].

For reactions represented by the Arrhenius expression (Eq. 2.1), uncertainty may be placed either on $k(T)$ directly or on individual parameters (A, n, E_a). The former approach is more common in mechanism assessments [26, 32] because experimental data typically constrain $k(T)$ more reliably than they constrain the individual Arrhenius parameters independently.

Methods and Sources to Determine Uncertainty Factors

The literature identifies several methods to derive kinetic uncertainty factors:

1. **Experimental scatter.** Combine reported measurements of $k(T)$ across studies (e.g., NIST Chemical Kinetics Database [26]) and derive f_r such that Eq. (2.4) encloses the experimentally observed spread.

2. **Statistical regression and confidence bands.** Meta-analyses of experimental datasets, as used in the evaluations by Baulch et al. [24], yield statistical uncertainty intervals convertible to f_r.

3. **Theoretical error assessment.** High-level quantum chemical studies may report uncertainty estimates arising from method, basis-set, or tunneling corrections; these can be translated to f_r following approaches described by Klippenstein et al. [25].

4. **Analogy.** For poorly studied reactions, uncertainty factors can be inferred using chemically analogous families or known within-family scatter, following practices outlined in [24].

5. **Expert elicitation.** When data are limited, formal expert judgment methodologies are used to assign conservative bounds [29].

Empirical guidelines found in the literature include:

- $f_r \approx 0.1{-}0.3$ for well-characterized elementary reactions with consistent datasets [24].
- $f_r \approx 0.3{-}0.5$ for reactions with sparse or conflicting measurements [26].
- Larger uncertainty intervals for reactions constrained mainly by theoretical calculations or indirect inference [25].

Regardless of the approach, the choice of f_r should always be documented and justified relative to the data or theoretical sources used.

Use of f_r in Optimization Workflows

Uncertainty factors provide a principled way to define bounded search domains for parameter optimization, including evolutionary strategies and genetic algorithms [30, 31]. Representing the uncertainty directly in $\log_{10} A$-space enables the optimization to explore a physically interpretable and statistically justified domain:

$$\log_{10} A \in [\log_{10} A_0 - f_r,\ \log_{10} A_0 + f_r].$$

When uncertainties on n or E_a are also incorporated, analogous bounded intervals must be specified, with care taken to preserve thermochemical consistency where required.

Given these uncertainty intervals on $\log_{10} A$, the question arises: *how should analogous bounds be defined for the temperature exponent n and activation energy E_a, which are not directly constrained by the uncertainty factor f_r?* Several approaches have been proposed in literature to address this challenge, each with distinct assumptions regarding the interdependence of Arrhenius parameters. These methodologies, ranging from simplified linear approximations to nonlinear constraint formulations, are reviewed in the following section, with their comparative performance evaluated experimentally in Chapter 4.

2.5 Gene Space Definition Approaches: Comparative Review

In kinetic mechanism optimization, the Arrhenius parameters of selected reactions must be varied within physically meaningful bounds. The challenge lies in defining a *gene space*, the allowable parameter range for optimization algorithms, that captures experimental uncertainty while maintaining physical consistency.

The following section reviews three approaches to this challenge: (1) linear approximations (Bertolino et al., 2021), (2) nonlinear bounds (Fürst and Senosiain, 2018), and (3) the feasibility problem common to both methods, which motivates the ICA-based approach presented in Chapter 3.

Experimental Uncertainty and Rate Constant Bounds

Experimental measurements of rate constants are inherently uncertain. This uncertainty is typically quantified by an **uncertainty factor** f_r, which defines multiplicative bounds (see Eq. 2.4) on the nominal rate constant $k_0(T)$.

These bounds represent an **uncertainty envelope** within which the true rate constant is assumed to lie across all temperatures in the range $[T_{\min}, T_{\max}]$. Any set of optimized parameters must produce rate constants that remain within this envelope to be *feasible*.

Linear Approximation: Bertolino et al. (2021)

Bertolino et al. [33] proposed a simplified, linear approach to define parameter bounds. By assuming that the uncertainty factor f_r can be linearly distributed among the three Arrhenius parameters, they derived rectangular bounds in the parameter space (α, n, ϵ), where $\alpha = \ln(A)$ and $\epsilon = E_a/R$:

$$\alpha_{\min} = \ln\left(\frac{A_{\text{nom}}}{10^{f_r}}\right), \quad \alpha_{\max} = \ln\left(A_{\text{nom}} \cdot 10^{f_r}\right) \tag{2.5}$$

$$n_{\min} = n_{\text{nom}} - \frac{\ln(10) \cdot f_r}{\ln(T_{\max})}, \quad n_{\max} = n_{\text{nom}} + \frac{\ln(10) \cdot f_r}{\ln(T_{\max})} \tag{2.6}$$

$$\epsilon_{\min} = \frac{E_{a,\text{nom}} - \ln(10) \cdot f_r \cdot R \cdot T_{\min}}{R}, \quad \epsilon_{\max} = \frac{E_{a,\text{nom}} + \ln(10) \cdot f_r \cdot R \cdot T_{\min}}{R} \tag{2.7}$$

These bounds form a *rectangular gene space* that is computationally efficient but overly conservative. The linear approximation assumes that the uncertainty in $k(T)$ can be *independently partitioned* among A, n, and E_a, neglecting the nonlinear coupling inherent in Eq. 2.1. As a result, many parameter combinations near the boundaries of this space may violate the rate constant envelope at intermediate temperatures, even though they satisfy the bounds at $T_{\min}$ and $T_{\max}$.

Nonlinear Bounds: Fürst and Senosiain (2018)

Recognizing the limitations of linear approximations, Fürst and Senosiain [34] derived *nonlinear bounds* that account for the interdependence of Arrhenius parameters. By analyzing the extreme rate constant values at the temperature boundaries $(T_{\min}, T_{\max})$, they obtained analytical expressions for parameter bounds that guarantee the existence of at least one rate constant curve within the uncertainty envelope.

The bounds for A remain identical to Bertolino's approach (Eq. 2.5). However, the bounds for n are derived from the constraint that rate constants at $T_{\min}$ and $T_{\max}$ must simultaneously satisfy the uncertainty envelope:

$$n_1 = \frac{\ln\left(\frac{k_{\min}(T_{\max})}{k_{\max}(T_{\min})}\right) + \frac{E_{a,\text{nom}}}{R}\left(\frac{1}{T_{\max}} - \frac{1}{T_{\min}}\right)}{\ln\left(\frac{T_{\max}}{T_{\min}}\right)}$$

$$n_2 = \frac{\ln\left(\frac{k_{\min}(T_{\min})}{k_{\max}(T_{\max})}\right) - \frac{E_{a,\text{nom}}}{R}\left(\frac{1}{T_{\max}} - \frac{1}{T_{\min}}\right)}{\ln\left(\frac{T_{\min}}{T_{\max}}\right)}$$

$$n_{\min} = \min(n_1, n_2), \quad n_{\max} = \max(n_1, n_2) \tag{2.8}$$

Similarly, bounds for ϵ are obtained from:

$$\epsilon_1 = \frac{\ln\left(\frac{k_{\min}(T_{\max})}{k_{\max}(T_{\min})}\right) - n_{\text{nom}} \ln\left(\frac{T_{\max}}{T_{\min}}\right)}{\frac{1}{T_{\min}} - \frac{1}{T_{\max}}}$$

$$\epsilon_2 = \frac{\ln\left(\frac{k_{\min}(T_{\min})}{k_{\max}(T_{\max})}\right) - n_{\text{nom}} \ln\left(\frac{T_{\min}}{T_{\max}}\right)}{\frac{1}{T_{\max}} - \frac{1}{T_{\min}}}$$

$$\epsilon_{\min} = \min\left[\epsilon_1, \epsilon_2\right], \quad \epsilon_{\max} = \max\left[\epsilon_1, \epsilon_2\right] \tag{2.9}$$

While Fürst's bounds are more physically grounded than Bertolino's, they guarantee only that **at least one feasible parameter combination exists** within the rectangular space, not that **all** combinations are feasible. In practice, the bounds still define a rectangular gene space that may contain a large fraction of infeasible points, particularly for reactions with low activation energies or large uncertainty factors.

The Feasibility Problem

A critical limitation shared by both the Fürst and Bertolino approaches is that *defining a gene space does not guarantee parameter feasibility.* The rectangular bounds $(A_{\min}, A_{\max}) \times (n_{\min}, n_{\max}) \times (E_{a,\min}, E_{a,\max})$ ensure mathematical consistency at the temperature extremes $(T_{\min}, T_{\max})$ but do not enforce the constraint:

$$k_{\min}(T) \leq k(T; A, n, E_a) \leq k_{\max}(T) \quad \forall T \in [T_{\min}, T_{\max}] \tag{2.10}$$

In practice, many parameter combinations (A, n, E_a) within the gene space produce rate constants that violate the uncertainty envelope at intermediate temperatures, despite satisfying the boundary conditions. The envelope bounds themselves are calculated using only the nominal parameters and uncertainty factor:

$$k_{\min}(T) = \frac{A_{\text{nom}}}{10^{f_r}} \cdot T^{n_{\text{nom}}} \cdot \exp\left(-\frac{E_{a,\text{nom}}}{RT}\right)$$

$$k_{\max}(T) = A_{\text{nom}} \cdot 10^{f_r} \cdot T^{n_{\text{nom}}} \cdot \exp\left(-\frac{E_{a,\text{nom}}}{RT}\right)$$

Crucially, the envelope bounds depend only on the nominal parameters ($A_{\text{nom}}, n_{\text{nom}}, E_{a,\text{nom}}$) and uncertainty factor f_r, not on the candidate parameters being tested. This ensures that all optimization candidates are evaluated against a consistent, experimentally justified reference. To enforce feasibility during optimization, an **explicit rate constant check** must be performed:

$$\text{if } \exists T_i |\ k(T_i; A, n, E_a) < k_{\min}(T_i) \text{ or } k(T_i; A, n, E_a) > k_{\max}(T_i) \quad \Rightarrow \quad \text{reject params}$$

To quantify the severity of this problem, a Monte Carlo sampling (N = 10,000 samples per reaction) was performed for both the Fürst and Bertolino gene spaces for five representative reactions in the H_2-NH_3 mechanism across a temperature range of 300–3000 K. The selected reactions (Table 2.2) span a range of uncertainty factors against ignition delay times from $f_r = 0.1$ to 0.5.

Reaction	**Equation**	$\boldsymbol{f_r}$	**Fürst (%)**	**Bertolino (%)**
R26	$NH_3 + OH \Leftrightarrow H_2O + NH_2$	0.3	6.73	61.24
R25	$NH_3 + H \Leftrightarrow H_2 + NH_2$	0.5	7.04	60.88
R39	$NH_2 + NH \Leftrightarrow N_2H_2 + H$	0.2	6.88	61.17
R77	$NH_2 + NO \Leftrightarrow NNH + OH$	0.1	7.27	61.08
R34	$NH_2 + H \Leftrightarrow NH + H_2$	0.4	6.97	60.73
Joint Feasibility (5 reactions):			**0.0002%**	**8.46%**

Table 2.2. Feasibility rates for the Fürst and Bertolino gene space definitions across selected reactions in the H_2-NH_3 mechanism. Monte Carlo sampling (N = 10,000) was performed over the temperature range 300–3000 K. The joint feasibility probability assumes statistical independence between reactions.

The results (Table 2.2) reveal *severe feasibility gaps in both methods*:

- **Fürst method**: Only 6.7–7.3% of randomly sampled parameter combinations remained within the envelope at all temperatures, despite satisfying the nonlinear boundary constraints.
- **Bertolino method**: The linear approximation yielded higher feasibility rates (60.7–61.2%), but still rejected approximately 40% of candidates.

Notably, the Fürst feasibility rates are nearly independent of the uncertainty factor f_r, remaining consistently around 7%. This suggests that the infeasibility is fundamentally caused by the *shape* of the gene space—the nonlinear boundaries create large regions of parameter space that are mathematically valid at temperature extremes but violate the envelope at intermediate temperatures. The Bertolino method, while achieving higher single-reaction feasibility, suffers from the opposite problem: its linear bounds are overly conservative, excluding potentially valid parameter combinations.

Implications for Multi-Reaction Optimization

For a multi-reaction optimization problem involving N reactions, assuming independent feasibility rates, the probability of finding a *joint-feasible* solution, for which all reactions are simultaneously feasible) can be estimated as:

$$P_{\text{joint}} = \prod_{i=1}^{N} p_i$$

where p_i is the feasibility rate for reaction i. For the five-reaction system analyzed here:

- **Fürst method**: $P_{\text{joint}} = 1.65 \times 10^{-6}$ (0.0002%)
- **Bertolino method**: $P_{\text{joint}} = 0.0846$ (8.46%)

These joint probabilities demonstrate that even for a modest five-reaction subsystem, conventional genetic algorithm optimization becomes **highly inefficient**. For the Fürst method, one would need to evaluate approximately *600,000 candidate individuals* on average to find a single joint-feasible solution. Even for the Bertolino method, roughly 12 candidates must be tested per feasible solution, with significant computational effort wasted on infeasible evaluations.
This check is computationally expensive and introduces a high rejection rate during genetic algorithm optimization, significantly reducing convergence efficiency. The problem scales exponentially with the number of reactions in the mechanism, making traditional analytical gene space definitions impractical for realistic multi-reaction systems.

Motivation for Data-Driven Approaches

Neither the Bertolino [33] linear bounds nor the Fürst [34] nonlinear bounds are sufficiently precise to guarantee parameter feasibility across the entire temperature range of interest. Both methods define gene spaces that are mathematically sound, ensuring correct behavior at temperature boundaries, but *physically incomplete*, permitting large fractions of infeasible interior points.

The compounded infeasibility problem becomes particularly severe in multi-reaction optimization. For the five-reaction H_2-NH_3 subsystem analyzed here, joint feasibility rates range from 0.0002% (Fürst) to 8.46% (Bertolino), making GA optimization highly inefficient, with most computational effort wasted on evaluating infeasible candidates.

This motivates a paradigm shift: rather than defining gene spaces *a priori* through analytical bounds, *what if we could learn the feasible parameter space directly from data?* By sampling the feasible region and identifying its intrinsic structure, we can construct a gene space that naturally captures nonlinear parameter correlations, achieving high feasibility rates by construction, while reducing effective dimensionality through statistical independence, and accelerating genetic algorithm convergence. This data-driven approach, formalized as **REPRICA** (REaction Parameter Reconstruction via Independent Component Analysis), is later presented in Section 3.6.

2.6 Optimisation techniques

Recent efforts in chemical kinetics mechanism development have extended beyond sensitivity and uncertainty analysis to include *parameter optimisation*, whereby uncertain kinetic parameters are adjusted to improve agreement with experimental or high-fidelity reference data. Detailed mechanisms often exhibit highly nonlinear and multimodal objective landscapes, making optimisation a critical step in mechanism calibration.

A widely adopted general-purpose framework is **Dakota**, an open-source suite supporting optimisation, uncertainty quantification, sensitivity analysis, and parameter estimation for complex simulation models. Dakota's modular architecture allows coupling with reacting flow and chemical kinetics solvers, and it supports both gradient-based and derivative-free algorithms, including global and local optimisation strategies [35].

Among derivative-free optimisation strategies, **Genetic Algorithms (GAs)** have been frequently employed in kinetic mechanism calibration. GAs are population-based, stochastic search algorithms that explore the parameter space without requiring derivative information, making them well suited for highly nonlinear or discontinuous objective functions. Recent studies have demonstrated that GAs can efficiently optimise reaction rate coefficients, leading to reduced errors in predicted ignition delays and species profiles [36, 37].

In this work, kinetic parameter optimisation is implemented using **PyGAD** (see Sec. 3.5), a Python-based library for Genetic Algorithms. PyGAD allows flexible definition of populations, fitness functions, and genetic operators, and integrates easily with Python-based simulation tools.

Overall, frameworks such as Dakota and algorithmic paradigms such as Genetic Algorithms provide robust strategies for exploring high-dimensional, nonlinear parameter spaces in chemical kinetics. Their adoption in mechanism calibration reflects the ongoing need for optimisation methods that can handle complex, multimodal objective landscapes inherent in combustion systems.

2.7 Curve Matching Score for Model Validation

The validation of chemical kinetic mechanisms against experimental data is fundamental in optimisation. Traditional validation approaches typically rely on point-wise comparison metrics such as Mean Squared Error (MSE) or Mean Relative Error (MRE). However, these conventional metrics present significant limitations when applied to combustion data, where experimental data may be sparse, irregularly sampled, or affected by uncertainties.

To address these limitations, Bernardi et al. [38] introduced the *Curve Matching (CM) score*, a generalized framework for comparing numerical models with experimental observations. The CM methodology was developed at Politecnico di Milano specifically to provide a more robust and comprehensive assessment of kinetic mechanism performance.

The fundamental innovation of the Curve Matching approach lies in its treatment of experimental data and model predictions as continuous curves rather than discrete point sets. This is achieved through functional data analysis techniques, where both experimental measurements and simulation results are represented as smooth functions using spline interpolation. By adopting this functional perspective, the CM score captures not only the point-wise differences between curves but also their overall shape (through first derivatives) and potential horizontal displacements.

Mathematical Formulation

The CM method evaluates the similarity between two functions f (experimental) and g (simulated), returning a normalized score in the range $(0, 1]$, where 1 indicates perfect agreement. The formulation varies depending on the experimental data type.

For global combustion properties such as *ignition delay times* or *laminar flame speeds*, the CM score is defined as:

$$\text{CM Score}(f, g) = \frac{d_0^{L_2} + d_1^{L_2} + d_0^{Pe} + d_1^{Pe}}{4} \tag{2.11}$$

where:

- $d_0^{L_2}$ is the squared deviation (L_2 norm) between f and g
- $d_1^{L_2}$ is the squared deviation between their first derivatives
- d_0^{Pe} is an index related to the Pearson correlation coefficient
- d_1^{Pe} is an index related to the Pearson correlation of the derivatives

For *speciation data*, where species concentration profiles may exhibit horizontal shifts due to kinetic rate differences, the CM score includes an additional term to account for horizontal misalignment:

$$\text{CM Score}(f, g) = \frac{d_0^{L_2} + d_1^{L_2} + d_0^{Pe} + d_1^{Pe} + 2S}{6} \tag{2.12}$$

where S is a measure of horizontal similarity defined by:

$$S = \max\left(1 - \frac{\delta}{D}\right)$$

$$\delta = \arg\max\left(d_0^{L_2} + d_1^{L_2} + d_0^{Pe} + d_1^{Pe}\right)$$

with D denoting the common domain length of the two curves, and δ the optimal horizontal shift for alignment.

Application in Mechanism Optimisation

One of the key advantages of the CM methodology is its ability to handle multiple experimental datasets simultaneously. When validating a mechanism against multiple targets, an aggregated fitness function can be constructed by averaging the CM scores across all datasets, providing a single metric that reflects the overall mechanism performance. This normalized formulation (ranging from 0 to 1) makes the CM score particularly well suited for optimisation algorithms, such as genetic algorithms, where a reliable fitness function is essential for guiding the search process toward improved model parameters.

Since its introduction, the Curve Matching score has gained recognition in the combustion community as a powerful tool for kinetic mechanism validation and optimisation. Its robustness to data irregularities, scale-invariant formulation, and ability to capture both vertical deviations and horizontal shifts make it superior to traditional point-wise metrics for combustion applications.

Chapter 3

Materials and Methods

3.1 Workflow overview

The optimization workflow (Fig 3.1) implements a REPRICA-based approach integrated with a PyGAD genetic algorithm to systematically improve the kinetic mechanism's predictive accuracy. The procedure begins with the definition of parameter boundaries for the Arrhenius parameters (α, n, ϵ) of selected reactions, establishing physically meaningful constraints that prevent the exploration of unrealistic parameter space regions.

At the core of the workflow lies the REPRICA methodology (sec 3.6), which consists of three interconnected components operating in a closed loop. First, the Forward Independent Component Analysis (ICA) transforms the physical Arrhenius parameters (α, n, ϵ) into statistically independent latent variables (s_1, s_2, s_3), enabling more efficient exploration of the parameter space by the genetic algorithm. The PyGAD optimization engine then manipulates these latent variables to search for improved parameter sets within the established gene space. Subsequently, the Backward ICA transforms the optimized latent variables back into physical Arrhenius parameters, ensuring that the optimized values can be directly substituted into the chemical kinetic mechanism.

A critical feature of the workflow is the boundary validation step, which verifies that the proposed Arrhenius parameters satisfy the predefined physical constraints: this is the feasibility check (Eq. 2.10). If the parameters violate these boundaries, which would lead to chemically unrealistic reaction rates, the algorithm immediately returns a penalty CM score (set to 0.0) to PyGAD, preventing the evaluation of invalid parameter sets and significantly reducing computational cost by avoiding unnecessary Cantera simulations. This pre-screening mechanism ensures that only physically plausible candidates proceed to the computationally expensive evaluation phase.

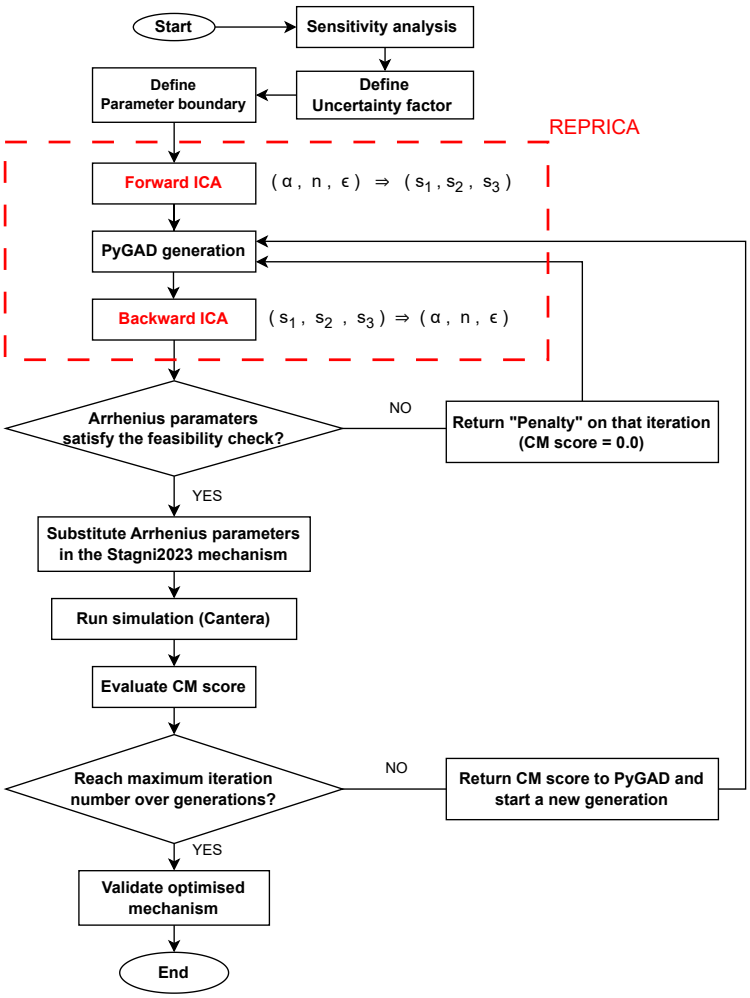

Figure 3.1. Workflow overview, from sensitivity analysis to optimised mechanism validation

For valid parameter sets, the workflow proceeds to substitute the optimized Arrhenius parameters into the chemical kinetic mechanism and performs full Cantera simulations across all experimental targets. The Curve Matching (CM) score is then evaluated to quantify the mechanism's predictive performance. This objective function value is returned to PyGAD, which uses it to guide the evolution of the population towards improved solutions, by prioritizing the best fits of each generation

The iterative process continues until the maximum number of generations is reached, with each generation producing increasingly accurate parameter sets through genetic operations (selection, crossover, and mutation). Upon reaching convergence, the workflow terminates, yielding an optimized mechanism with improved predictive capability across the comprehensive validation dataset.

3.2 Baseline Mechanism and Experimental Datasets

Stagni et al. (2023) Mechanism

The baseline kinetic mechanism employed in this study was developed by Stagni et al. [20] and represents a state-of-the-art detailed model for H_2/NH_3 combustion. The mechanism comprises: *31 chemical species*: H_2, NH_3, N_2, H_2O, NO, N_2O, and intermediate radicals (NH_2, NH, N, OH, H, O, HO_2, etc.); *203 reactions*, including H/O/N radical chemistry, ammonia oxidation pathways, and NO_x formation routes; *Thermodynamic database*, NASA polynomial coefficients for species enthalpy and entropy; *Transport properties*: species-specific collision diameters and Lennard-Jones parameters.

Reaction Types and Optimization Scope This mechanism includes four classes of reactions: *elementary reactions* (simple Arrhenius form), *three-body reactions* (pressure-dependent with collision efficiency factors), *falloff reactions* (pressure-dependent with Troe formalism), and *pressure-dependent Arrhenius reactions* (PLOG format with tabulated rate constants at discrete pressures). Table 3.1 summarizes their prevalence and optimization approach.

Reaction Type	**Count**	**Optimization Scope**
Elementary (simple Arrhenius)	170	(A, n, E_a) optimized
Three-body reactions	7	(A, n, E_a) optimized
Falloff (Troe formalism)	8	Low-P limit $(A_0, n_0, E_{a,0})$ optimized
PLOG (pressure-dependent)	18	Uniform shifts (X_1, X_2, X_3) at all P
Total	**203**	

Table 3.1. Reaction types in the Stagni et al. (2023) mechanism and their treatment in the optimization framework.

For **elementary reactions**, all three modified Arrhenius parameters (A, n, E_a) are optimized within uncertainty bounds. For **three-body reactions**, the rate constant parameters are optimized while third-body collision efficiencies are held constant, as these are typically derived from molecular dynamics simulations or experimental measurements and exhibit low sensitivity to global combustion metrics [2].

For **falloff reactions**, only the low-pressure-limit parameters $(A_0, n_0, E_{a,0})$ are included in the optimization, as the high-pressure limit and Troe broadening factors $(a, T^*, T^{**}, T^{***})$ are typically constrained by master-equation calculations or experimental fits and exhibit low sensitivity to ignition delay and species profiles [39].

For **PLOG reactions**, the present study adopts the approach proposed by Bertolino et al. (2021) [33], which optimizes three uniform shift parameters (X_1, X_2, X_3) applied

to all pressure-dependent rate constant entries rather than treating each pressure point independently. This methodology preserves the shape of the pressure-dependent rate surface derived from master-equation calculations while allowing systematic adjustment of the overall kinetic behavior. The modified rate constant at each pressure is given by:

$$\begin{aligned} \alpha'(P) &= \alpha(P) + X_1, \quad \text{where } \alpha = \ln(A), \\ \beta'(P) &= \beta(P) + X_2, \quad \text{where } \beta = n, \\ \varepsilon'(P) &= \varepsilon(P) + X_3, \quad \text{where } \varepsilon = E_a/R. \end{aligned} \tag{3.1}$$

The shift parameters are bounded by the uncertainty factor f_r for the reaction:

$$\begin{aligned} X_1 &\in \left[-\ln(10^{f_r}), +\ln(10^{f_r})\right], \\ X_2 &\in \left[-\frac{f_r}{\log_{10}(T_{\max})}, +\frac{f_r}{\log_{10}(T_{\max})}\right], \\ X_3 &\in \left[-f_r \cdot T_{\min} \cdot \ln(10), +f_r \cdot T_{\min} \cdot \ln(10)\right], \end{aligned} \tag{3.2}$$

where $T_{\min} = 300$ K and $T_{\max} = 3000$ K define the temperature range of interest. This approach reduces the optimization search space from $3 \times N_{\text{pressures}}$ to 3 variables per PLOG reaction while maintaining thermodynamic consistency across the pressure-dependent rate surface. For example, a reaction that includes rate constants at 4 discrete pressures, would require optimization of only 3 parameters (X_1, X_2, X_3) instead of 4 independent Arrhenius triplets, thereby significantly improving computational efficiency and avoiding physically unrealistic discontinuities in the $k(T, P)$ surface.

Excluded parameters: The following parameters were held constant during optimization to maintain chemical consistency and avoid over-parameterization:

- **Third-body collision efficiencies** (e.g., enhanced collision efficiency of H_2O in recombination reactions): experimentally determined and mechanism-specific
- **Troe and High-pressure-limit falloff parameters** (a, T_1, T_2, T_3): fixed correlation parameters derived from quantum-chemical master equations, and not particularly sensitive to Ignition delay times and species profiles [2, 39]
- **Pressure-dependent shape for PLOG reactions**: only uniform shifts are applied, preserving the relative pressure dependence from the baseline mechanism

This restriction reduces the optimization search space significantly while preserving chemical fidelity. For the present study, 20 target reactions were selected for optimization based on sensitivity analysis (Section 4.1), resulting in 60 degrees of freedom (20 reactions $\times$ 3 Arrhenius parameters/uniform shifts).

Justification for Selection The Stagni et al. (2023) mechanism was selected as the baseline for several reasons:

1. **Recent development**: Published in 2023, it incorporates the latest rate constant developments from Glarborg et al. (2018) [2] and Klippenstein et al. (2017) [25].
2. **Comprehensive validation scope**: The original mechanism was validated against:
   - Shock tube ignition delay times (STIDT) for H_2/NH_3 mixtures [40]
   - Laminar burning velocities (LBV) for lean-to-rich flames [41]
   - Jet-stirred reactor (JSR) species profiles for NH_3 oxidation [42]
   - Flow reactor measurements of NO_x formation [43]
3. **Systematic errors identified**: Despite its broad validation, the Stagni mechanism exhibits: overall *underprediction of ignition delay times*, indicating the mechanism predicts faster ignition than observed experimentally [20] and *overprediction of nitrogen oxides*, particularly problematic in high-temperature oxidation conditions.
4. **Detailed pressure-dependent chemistry**: The inclusion of 18 PLOG reactions (e.g., NH_3 decomposition, $NH_2 + O_2$ reactions) enables accurate modeling across a wide pressure range (0.1–100 atm), which is critical for both fundamental combustion studies and practical applications (gas turbines, internal combustion engines).

Baseline Performance: Motivating the Need for Optimization

Table 3.2 quantifies the baseline mechanism's error metrics, which serve as benchmarks for evaluating optimization performance.

Target Type	Cases	CM Score[1]	MSE[2]	MRE (%)[3]
STIDT	21	0.926	0.123	24.6
JSRSP	85	0.892	1.63×10^{-6}	∞[4]
STSP	54	0.878	2.14×10^{-7}	89.4
Overall	**160**	**0.893**	—	—

Table 3.2. Baseline mechanism performance metrics (Stagni et al. 2023). MSE and MRE values demonstrate fundamental limitations of these metrics for multi-observable mechanism validation (see critical analysis below).

[1] Curve Matching score (0-1 scale): dimensionless similarity metric robust to multi-scale data.
[2] Mean Squared Error exhibits severe multi-scale issues (see discussion below).
[3] Mean Relative Error suffers from numerical instability when experimental values approach zero.
[4] Undefined due to division by near-zero experimental mole fractions.

Critical Analysis: Why MSE and MRE Are Unsuitable

Table 3.2 reveals fundamental flaws in traditional error metrics (MSE, MRE) that render them unsuitable for multi-observable combustion mechanism optimization:

MSE Multi-Scale Problem: MSE exhibits catastrophic scale dependence across observable types. STIDT contributes 99.998% of total MSE ($\text{MSE}_{\text{STIDT}} = 0.123$) despite representing only 13% of validation cases, completely masking errors in JSRSP ($\text{MSE}_{\text{JSRSP}} = 1.63 \times 10^{-6}$) and STSP ($\text{MSE}_{\text{STSP}} = 2.14 \times 10^{-7}$). This six-order-of-magnitude difference arises from inherent observable scales: STIDT measures ignition delays in milliseconds (typically 1–100 ms), while JSRSP/STSP measure species mole fractions (typically 10^{-6}–10^{-2}). Consequently, MSE-based optimization would exclusively target STIDT accuracy while ignoring species concentration errors entirely, unacceptable for mechanism development requiring balanced multi-observable performance.

MRE Numerical Instability: MRE becomes mathematically undefined (∞) for some JSRSP due to division by near-zero experimental mole fractions. When experimental concentrations approach detection limits (e.g., $N_2O \sim 10^{-7}$ mole fraction), even tiny absolute errors produce infinite relative errors:

$$\text{MRE} = \left| \frac{y_{\text{sim}} - y_{\text{exp}}}{y_{\text{exp}}} \right| \times 100\% \to \infty \quad \text{as} \quad y_{\text{exp}} \to 0$$

This numerical pathology renders MRE computationally inadequate for automated optimization algorithms, which require smooth, well-behaved objective functions. In addition, MRE is undefined when $y_{\text{exp}} = 0$ exactly (e.g., intermediate radicals below detection limits), requiring arbitrary exclusion of valid experimental data points.

Curve Matching (CM) as Robust Alternative: In contrast, the CM score (Sec. 2.7) exhibits none of these pathological behaviors. CM is scale-independent by design, treating all observable types with equal weight regardless of absolute magnitude. CM involves no division by experimental values, eliminating numerical instability. The overall CM score (0.893) represents a balanced assessment across 160 validation cases spanning six orders of magnitude in observable scales, precisely the metric needed for robust multi-objective mechanism optimization.

Type	Reference	T Range (K)	P (atm)	ϕ	H_2 %[5]	Pts
Shock Tube Ignition Delay Times (STIDT)						
STIDT	Mathieu & Petersen (2015) [40]	1564–2489	1.4–28.9	0.5–2.0	0	103
STIDT	Chen et al. (2021) [44]	1023–1957	1.2–10.0	1.0	0–70	66
STIDT	Baker et al. (2023) [45]	1438–1694	2.2	1.0	50	5
					STIDT Subtotal:	**204**
Jet-Stirred Reactor Species Profiles (JSRSP)						
JSRSP	Rota et al. (2001) [46]	1022–1387	1.0	0.1–2.0	0	36[6]
JSRSP	Stagni et al. (2020) [47]	500–1200	1.1	0.009–0.019	0	50[7]
JSRSP	Zhang et al. (2021) [48]	800–1281	1.0	0.25–1.0	0–70	376[8]
JSRSP	Osipova et al. (2022) [49]	800–1300	1.0	0.6–1.5	0–30	366[9]
JSRSP	Tang et al. (2022) [50]	700–1190	1.0	0.1–1.0	0	248[10]
					JSRSP Subtotal:	**1,076**
Shock Tube Species Profiles (STSP)						
STSP	Davidson et al. (1990) [51]	2294–2781	0.88–1.03	—	pyrolysis[11]	152
STSP	Halat-Augier et al. (1994) [52]	1100–1190	2.17–2.92	—	pyrolysis	56
STSP	Alturaifi et al. (2022) [53]	2096–3007	0.88–1.26	varied	0–85	56,871
STSP	Mathieu et al. (2024) [54]	1458–2222	1.26–1.50	varied	0	55,521
					STSP Subtotal:	**112,600**
					TOTAL:	**113,880**

Table 3.3. Summary of optimization target datasets (verified from original publications)

Optimization Target Datasets

The optimization was performed against *12 experimental datasets* spanning three complementary observable types, selected to capture different aspects of H_2/NH_3 combustion chemistry across a wide range of temperatures, pressures, and equivalence ratios. Table 3.3 provides a comprehensive overview of all datasets used in this work.

(a) Shock Tube Ignition Delay Times (STIDT) Shock tube ignition delay time (STIDT) measurements constrain the high-temperature radical chemistry that governs autoignition behavior in practical combustion devices. Ignition delay is defined as the time interval between reflected shock compression and the onset of ignition, typically identified by the maximum rate of pressure rise (dP/dt) or OH^* chemiluminescence peak.

The three STIDT datasets employed in this study are: `Mathieu & Petersen (2015)` [40], `Chen et al. (2021)` [44], and `Baker et al. (2023)` [45].

[5] H_2 percentage calculated as: $100 \times \frac{[H_2]}{[H_2]+[NH_3]}$ in fuel mixture.

[6] 22 NO measurements + 14 NO measurements (2 experimental series).

[7] 25 temp points × 2 species (NH_3, NO) across 2 cases.

[8] 94 temp points × 4 species (H_2O, NH_3, NO, N_2O) across 10 cases.

[9] 78 temp points, 4–5 species per case (NH_3, N_2, O_2, NO, N_2O), 4 cases.

[10] 37 temp points, 6–8 species per case, 3 cases.

[11] *Pyrolysis experiment (no oxidizer).* Measures NH and NH_2 radicals in ppm. Validates NH_x radical kinetics ($NH_3 \rightarrow NH_2 \rightarrow NH \rightarrow N$) shared between pyrolysis and fuel-rich combustion zones. High-temperature range ensures rate constant accuracy beyond typical flame conditions.

Chemical relevance: STIDT data primarily constrain *chain-branching pathways* (e.g., $H + O_2 \rightarrow OH + O$, $NH_2 + O \rightarrow$ products) and *radical initiation steps* (e.g., $NH_3 + OH \rightarrow NH_2 + H_2O$), which dominate ignition behavior at high temperatures ($T > 1200$ K).

(b) Jet-Stirred Reactor Species Profiles (JSRSP) Jet-stirred reactor (JSR) species profile measurements probe *intermediate-temperature oxidation pathways* and *NO_x formation chemistry* under near-homogeneous conditions. Unlike shock tubes, JSRs achieve steady-state operation, allowing precise quantification of major and minor species concentrations as functions of temperature, residence time, and equivalence ratio.

The five JSRSP datasets are: `Rota et al. (2001)` [46], `Stagni et al. (2020)` [47], `Zhang et al. (2021)` [48], `Osipova et al. (2022)` [49], and `Tang et al. (2022)` [50].

Special note on Rota et al. (2001): This pioneering study investigated *pure NH_3 oxidation* (no H_2 blending) across an extended equivalence ratio range ($\phi = 0.1$–2.0), providing critical baseline constraints for ammonia-only combustion pathways. The fuel-lean conditions ($\phi < 0.7$) are particularly valuable for validating NO_x formation routes under oxygen-rich environments typical of gas turbine operation.

Chemical relevance: JSR data constrain *radical-intermediate reactions* (e.g., $NH_2 + HO_2 \rightarrow$ products, $NH + NO \rightarrow N_2O + H$) and *$NO_x$ formation pathways* (e.g., $NH + O_2 \rightarrow NO + OH$, N_2O decomposition routes). These pathways dominate fuel conversion and emissions at intermediate temperatures ($T = 700$–1300 K).

(c) Shock Tube Species Profiles (STSP) Shock tube species profile (STSP) measurements provide *time-resolved* validation of NH_x consumption kinetics during the early stages of ignition. Unlike global STIDT measurements (single scalar output), STSP data track the temporal evolution of key species (e.g., NH_3, NO, H_2O, and radical intermediates) behind reflected shocks, revealing the interplay between fuel decomposition, radical production, and heat release.

The four STSP datasets employed in this study are: `Davidson & Hanson (1990)` [51], `Halat-Augier et al. (1994)` [52], `Alturaifi et al. (2022)` [53], and `Mathieu et al. (2024)` [54].

Special note on Davidson et al. (1990): This foundational study investigated *ammonia pyrolysis*, measuring time-resolved NH and NH_2 concentration profiles. The experiments covered temperatures from 2200–2800 K, providing critical validation data for NH_x radical kinetics under extreme high-temperature conditions.

Chemical relevance: STSP data constrain *early-time fuel decomposition pathways* (e.g., $NH_3 \rightarrow NH_2 + H$, $NH_2 + M \rightarrow NH + H + M$), *radical equilibration kinetics* (e.g., $H +$

$O_2 \leftrightarrow OH + O$, $NH_2 + O \leftrightarrow HNO + H$), and *N/O exchange reactions* (e.g., $N + O_2 \leftrightarrow NO + O$). These reactions govern the induction period before thermal runaway and the transition from fuel-bound nitrogen to NO_x products.

Rationale for multi-target optimization: The three observable types constrain different aspects of H_2/NH_3 combustion chemistry:

- **STIDT** (global ignition behavior) → chain-branching rates (H/O radicals), high-T regime ($T > 1200$ K)
- **JSRSP** (steady-state species profiles) → intermediate oxidation + NO_x formation, moderate-T regime ($T = 700$–1300 K)
- **STSP** (transient species evolution) → fuel decomposition + radical equilibration + N/O chemistry, high-T regime ($T > 1350$ K)

Optimizing against all three simultaneously ensures that parameter adjustments improve predictive accuracy across the entire combustion process, from initial fuel decomposition through intermediate oxidation to final autoignition, rather than overfitting to a single experimental observable. The inclusion of the Davidson & Hanson N-atom dataset further ensures that the fundamental N/O radical chemistry remains consistent at extreme temperatures beyond typical NH_3 combustion conditions.

Validation Dataset: Laminar Burning Velocity and Shock-Tube Ignition Delay

To assess the **transferability** of the optimized mechanism to conditions *not included in the optimization*, laminar burning velocity (LBV) measurements and an additional shock-tube ignition delay time (STIDT) dataset were used as independent validation targets.

Type	Reference	T_0 (K)	P_0 (atm)	ϕ	H_2 %	Pts
LBV	Mei et al. (2019) [55]	298	1–5	0.5–1.6	0	30
LBV	Mei et al. (2021) [56]	298	1	1.0–2.0	0	9
STIDT	Shu et al. (2018) [57]	1181–1581	19.7–39.5	0.5–2.0	0	30
					TOTAL VALIDATION TARGETS:	**110**

Table 3.4. Summary of validation datasets

Mei et al. (2019) [55] measured LBV for pure NH_3 using a constant-volume cylindrical bomb with high-speed schlieren imaging. The authors reported an overall measurement uncertainty of about 5%, primarily arising from the flame-radius determination and the linear-fit extraction of the propagation speed.

Mei et al. (2021) [56] examined the laminar flame propagation of NH_3, using as oxidizer a $50\%NO - 50\%N_2$) mixture, with the outwardly propagating spherical-flame technique

at an initial pressure of 1 atm and temperature of 298 K, spanning equivalence ratios $\phi = 1.0 - 2.0$. The measurements were carried out in a constant-volume cylindrical combustion vessel equipped with a schlieren imaging system.

Shu et al. (2018) [57] measured ignition delay times for pure NH_3/air mixtures behind reflected shock waves at high pressures (P_0 = 19.7–39.5 atm) and intermediate temperatures (T_0 = 1100–1600 K), covering fuel-lean ($\phi = 0.5$), stoichiometric ($\phi = 1.0$), and fuel-rich ($\phi = 2.0$) conditions. Ignition was detected via sidewall pressure monitoring, with IDT defined as the time of maximum rate of change of pressure. Reported overall uncertainties in the IDT values were estimated to be $\leq 20\%$ (primarily from temperature and pressure measurement errors and non-ideal facility effects such as boundary-layer growth).

LBV differs fundamentally from the optimization targets in three key aspects:

- **Temperature regime:** LBV probes low-temperature flame zones ($T_{\text{flame}} \approx$ 1800–2200 K from T_0 = 298–398 K), whereas STIDT/STSP probe reflected-shock temperatures ($T >$ 1200 K) and JSRSP probe intermediate oxidation (T = 700–1300 K).
- **Transport effects:** LBV is sensitive to species diffusion and thermal conductivity (Lewis number effects), which do not influence zero-dimensional STIDT or perfectly mixed JSRSP.
- **Observables:** LBV integrates the entire flame structure (preheat, reaction, and post-flame zones), whereas STIDT measures global induction time and JSRSP measures steady-state concentrations.

If the optimized mechanism improves STIDT/JSRSP/STSP predictions without degrading LBV accuracy, it demonstrates that rate constant adjustments are **chemically consistent** across temperature regimes, the mechanism has not **overfitted** to specific conditions, and optimized parameters maintain **predictive power** for transport-sensitive observables. Conversely, worsened LBV predictions indicate overfitting, a critical limitation requiring future refinement (Section 4.6).

The high-pressure STIDT dataset of Shu et al. (2018) was excluded from optimization and retained exclusively for validation. Accurate simulation of these experiments requires explicit incorporation of measured time-dependent volume profiles (VPRO files) to account for non-ideal shock-tube effects (boundary layer growth, post-reflected-shock pressure decay).

Preliminary optimization trials including Shu VPRO cases revealed that these targets disproportionately dominated runtime, biasing the optimizer toward a small subset of

expensive cases at the expense of broader experimental coverage. To maintain balanced representation of diverse observables and tractable optimization times, the Shu dataset was withheld and reserved as an independent, high-pressure validation benchmark not used for parameter tuning.

3.3 Sensitivity analysis criteria

The sensitivity analysis was conducted to identify the reactions exerting the strongest influence on key combustion observables across multiple experimental configurations. In particular, this work adopts a comprehensive approach by considering three distinct types of experimental datasets: shock tube ignition delay times (STIDT), jet-stirred reactor species profiles (JSRSP), and shock tube species profiles (STSP). This multi-observable strategy ensures that the selected reactions for optimization are those that govern not only ignition behavior but also the detailed chemical evolution of $NH_3 - H_2$ combustion systems. This integration stems from fundamental considerations regarding mechanism validation and optimization:

- **Complementary Information Content**: Each experimental configuration probes different aspects of the combustion chemistry. Shock tube ignition delay times (STIDT) are sensitive to the early chain-branching pathways that control ignition onset. Jet-stirred reactor species profiles (JSRSP) capture steady-state intermediate species concentrations under well-mixed conditions, revealing rate-limiting steps in oxidation pathways. Shock tube species profiles (STSP) provide time-resolved measurements of transient species during high-temperature oxidation, offering insight into both initiation and consumption mechanisms.

- **Extended Parameter Space Coverage**: The three datasets collectively span a broader range of thermodynamic conditions than any single experiment type. STIDT data typically cover temperatures from 900 to 1600 K at pressures of 1–30 atm. JSRSP measurements extend to lower temperatures (600–1200 K) and longer residence times (0.5–2 s), accessing near-equilibrium chemistry. STSP experiments reach higher temperatures (1400–2400 K), capturing kinetics often inaccessible to ignition delay measurements alone.

- **Mechanism Robustness**: A kinetic mechanism optimized against multiple observables is inherently more robust than one tuned to a single target. Reactions identified as influential across STIDT, JSRSP, and STSP datasets represent fundamental pathways whose accurate representation is essential for predictive capability

under diverse conditions, including those not explicitly included in the optimization process (e.g., flame speeds, emissions modeling).

- **Avoidance of Overfitting**: Optimizing against a single observable can lead to parameter compensation effects, where errors in less-sensitive reactions are masked by artificial adjustments to highly sensitive pathways. By requiring consistency across multiple observables, the multi-dataset approach constrains the optimization to coherent parameter space regions.

Mathematical Framework

For each experimental condition (T, P, X) in a given dataset, a sensitivity coefficient $S_{i,k}$ quantifies the relative change in the target observable resulting from a perturbation in the rate constant of reaction i at timestep k. The specific form of $S_{i,k}$ depends on the observable type:

For *shock tube ignition delay* measurements, the sensitivity coefficient is defined as:

$$S_{i,k}^{\mathrm{STIDT}} = \left.\frac{\partial \ln[\mathrm{OH}]}{\partial \ln k_i}\right|_{t=t_k} \tag{3.3}$$

where [OH] denotes the hydroxyl radical molar concentration and k_i is the rate constant of reaction i. The OH radical was selected as the target species due to its role as a key indicator of ignition onset, exhibiting a sharp concentration rise at the ignition point that correlates with pressure rise diagnostics used in shock tube experiments.

For *jet-stirred reactor species profile* measurements, sensitivities are computed for each measured species X_j:

$$S_{i,j}^{\mathrm{JSRSP}} = \left.\frac{\partial \ln[X_j]}{\partial \ln k_i}\right|_{t=\tau_{\mathrm{res}}}$$

where τ_{res} is the reactor residence time (typically 0.5–2 s) and $[X_j]$ represents the molar concentration of species j (e.g., NH_3, H_2, N_2O, NO, NO_2, H_2O). The global sensitivity for reaction i in a JSRSP dataset is obtained by averaging over all measured species:

$$S_i^{\mathrm{JSRSP}} = \frac{1}{N_{\mathrm{species}}} \sum_{j=1}^{N_{\mathrm{species}}} \left|S_{i,j}^{\mathrm{JSRSP}}\right| \tag{3.4}$$

For *shock tube species profile* measurements, sensitivities are computed for time-resolved concentration histories:

$$S_{i,j,k}^{\mathrm{STSP}} = \left.\frac{\partial \ln[X_j]}{\partial \ln k_i}\right|_{t=t_k}$$

Time-integration is performed to capture the cumulative influence over the observation window:

$$\mathcal{S}_i^{\mathrm{STSP}} = \int_0^{t_{\mathrm{obs}}} |S_{i,j}(t)| \; dt \tag{3.5}$$

Computational procedure

For each experimental condition, an appropriate reactor model was initialized using the Cantera library (version 3.0.0):

STIDT: Constant-pressure homogeneous reactor with adiabatic boundary conditions;

JSRSP: Perfectly-stirred reactor (PSR) with specified residence time and temperature;

STSP: Constant-pressure homogeneous reactor with time-resolved species tracking.

All reactors were configured with stringent numerical tolerances:

Tolerance	Time Integration	Sensitivity
Relative	$\mathrm{rtol} = 10^{-6}$	$\mathrm{rtol}_{\mathrm{sens}} = 10^{-6}$
Absolute	$\mathrm{atol} = 10^{-15}$	$\mathrm{atol}_{\mathrm{sens}} = 10^{-6}$

All the 203 reactions in the mechanism were then added to the sensitivity tracking list.

STIDT Processing The ignition delay time was defined as the time corresponding to the maximum rate of pressure rise:

$$t_{\mathrm{ign}} = \arg\max_t \left(\frac{dP}{dt}\right) \tag{3.6}$$

Sensitivity coefficients were integrated from $t = 0$ to $t = t_{\mathrm{ign}}$:

$$\mathcal{S}_i^{\mathrm{STIDT}} = \int_0^{t_{\mathrm{ign}}} S_i^{\mathrm{STIDT}}(t) \, dt$$

JSRSP Processing Species sensitivities were extracted at the steady-state residence time and averaged over all measured species as per Equation 3.4.

STSP Processing Time-resolved sensitivities were integrated over the experimental observation window (typically 1–5 ms) according to Equation 3.5.

Normalization within Each Dataset Type

For each dataset type ($d \in \{\text{STIDT}, \text{JSRSP}, \text{STSP}\}$), the integrated sensitivities were normalized by the maximum absolute value within that dataset category:

$$S_{i,d}^{\text{norm}} = \frac{\mathcal{S}_{i,d}}{\max_j |\mathcal{S}_{j,d}|}$$

This intra-dataset normalization is crucial to prevent dominance by any single experimental type, ensuring equal weighting in the subsequent global ranking.

Multi-Level Aggregation Strategy

A hierarchical aggregation approach was implemented to derive a global sensitivity ranking that accounts for variability across experimental conditions, datasets, and observable.

Level 1 – Condition Averaging within Each Dataset

For each individual dataset D_k (e.g., Mathieu STIDT mixture 1) containing N_{cond} experimental conditions, the mean normalized sensitivity for reaction i was computed as:

$$\bar{S}_i(D_k) = \frac{1}{N_{\text{cond}}} \sum_{m=1}^{N_{\text{cond}}} \left| S_{i,m}^{\text{norm}} \right|$$

Absolute values prevent cancellation effects between positive and negative sensitivities at different conditions.

Level 2 – Dataset-Type Averaging

For each experimental type $d \in \{\text{STIDT}, \text{JSRSP}, \text{STSP}\}$ containing N_d individual datasets, the type-level sensitivity was computed as:

$$\bar{S}_i^{(d)} = \frac{1}{N_d} \sum_{k=1}^{N_d} \bar{S}_i(D_k)$$

For example in this work: $N_{\text{STIDT}} = 12$ for the Mathieu datasets on shock tube ignition delay times.

Level 3 – Global Three-Way Combination

The final global sensitivity for reaction i was obtained by averaging across the three experimental types with equal weighting:

$$\bar{S}_i^{(\text{global})} = \frac{1}{3}\left(\bar{S}_i^{(\text{STIDT})} + \bar{S}_i^{(\text{JSRSP})} + \bar{S}_i^{(\text{STSP})}\right) \tag{3.7}$$

This equal weighting ensures that no single observable type dominates the ranking. The factor of 1/3 assigns 33.33% weight to each experimental category, reflecting the principle that ignition behavior, steady-state species distributions, and transient species evolution are equally important for comprehensive mechanism validation.

Level 4 – Filtering and Re-normalization

Reactions R1–R23, comprising the foundational H /O sub-mechanism, were excluded from the global ranking for the following reasons:

- **Mechanism Validation Status**: The hydrogen oxidation chemistry is among the most extensively validated sub-mechanisms in combustion kinetics, with rate constants constrained by thousands of experimental measurements spanning multiple decades of research [58].
- **Optimization Focus**: This work primary objective is to refine the predictive capability of NH -specific pathways, whose kinetic uncertainties is significantly larger.
- **Chemical Coupling Considerations**: While H_2 reactions exhibit high sensitivity coefficients due to their role in radical pool generation, their kinetic parameters serve as a validated foundation upon which the NH sub-mechanism is built.

Selection Criteria for Optimization

The top 20 reactions ranked by $S_i^{(\text{final})}$ were selected for subsequent uncertainty quantification and optimization. This threshold represents a compromise between:

- **Computational tractability**: The genetic algorithm search space scales exponentially with the number of reactions, making larger selections prohibitively expensive.
- **Chemical coverage**: The top 20 reactions encompass the primary NH_3 decomposition pathways (R24–R26), nitrogen chemistry intermediates (R40–R50, R70–R80), NO_x formation routes (R75–R77, R86–R90), and N_2O chemistry (R188–R190), providing comprehensive representation of the $NH_3 - H_2$ oxidation network.

- **Sensitivity threshold**: The vast majority of these reactions typically exhibits $S_i^{(\text{final})} > 0.10$, indicating non-negligible influence across at least one experimental observable type. Reactions below this threshold show minimal impact on model predictions within experimental uncertainty bounds.

This multi-dataset, multi-observable sensitivity analysis framework ensures that the selected reactions for optimization are those with demonstrated importance across the full range of combustion conditions relevant to $NH_3 - H_2$ co-firing applications, from low-temperature oxidation to high-temperature ignition and transient species evolution.

3.4 Uncertainty factor definition

Following the identification of sensitive reactions, uncertainty factors were assigned to define physically meaningful search bounds for the genetic algorithm optimization. The uncertainty factor f_r quantifies the permissible variation in the pre-exponential factor A of the Arrhenius expression for reaction r:

$$k(T) = AT^n \exp\left(-\frac{E_a}{RT}\right)$$

The determination of f_r followed a literature-based methodology that balances experimental validation with practical optimization constraints.

Literature-Based Approach

For each of the top 20 reactions identified through sensitivity analysis, experimental kinetic data were systematically retrieved from the NIST Chemical Kinetics Database [26]. The procedure involved three distinct stages:

Data Collection

A semi-automated Python script (`NIST_RC_collector_v1.7.py`) was developed to download all available experimental rate constant measurements $k(T)$ from NIST for each target reaction. The script parsed the NIST web interface and extracted:

- Experimental rate constant data as a function of temperature
- Temperature range of validity for each measurement
- Literature reference and experimental technique
- Reported uncertainties (when available)

The data were organized in reaction-specific subdirectories for subsequent analysis, with each subfolder containing CSV files for individual literature sources.

Comparative Visualization

For each reaction, the reference mechanism's rate expression (Stagni_2023 [20]) was plotted against the experimental literature data over the relevant temperature range (typically 500–2500 K, encompassing shock tube and flow reactor conditions). Uncertainty envelopes were constructed by applying multiplicative factors to the pre-exponential factor while maintaining the temperature exponent n and activation energy E_a unchanged:

$$k^{\pm}(T) = \left(A \cdot 10^{\pm f_r}\right) T^n \exp\left(-\frac{E_a}{RT}\right)$$

A dedicated plotting routine (`plotter.py`) generated Arrhenius plots displaying:

- The reference mechanism rate expression (solid line)
- Experimental data points from multiple literature sources (scatter points)
- Uncertainty envelopes for candidate values of f_r (shaded regions)

An exemple is here provided for reaction R26: $NH_3 + OH = NH_2 + H_2O$

Optimal Coverage Determination

The uncertainty factor f_r was selected as the minimum value for which the resulting $k^{\pm}(T)$ envelope encompassed the majority of experimental data points across the temperature range of interest. The selection criterion prioritized:

1. **Data Coverage**: The envelope should include at least 80% of available experimental measurements, accounting for outliers and potentially erroneous data points.
2. **Kinetic Plausibility**: The resulting rate constant bounds must remain physically realistic, avoiding regions that would violate collision theory limits or thermodynamic consistency.
3. **Conservative Constraints**: In cases where experimental data exhibited large scatter, preference was given to smaller f_r values to prevent over-parameterization and maintain model robustness.

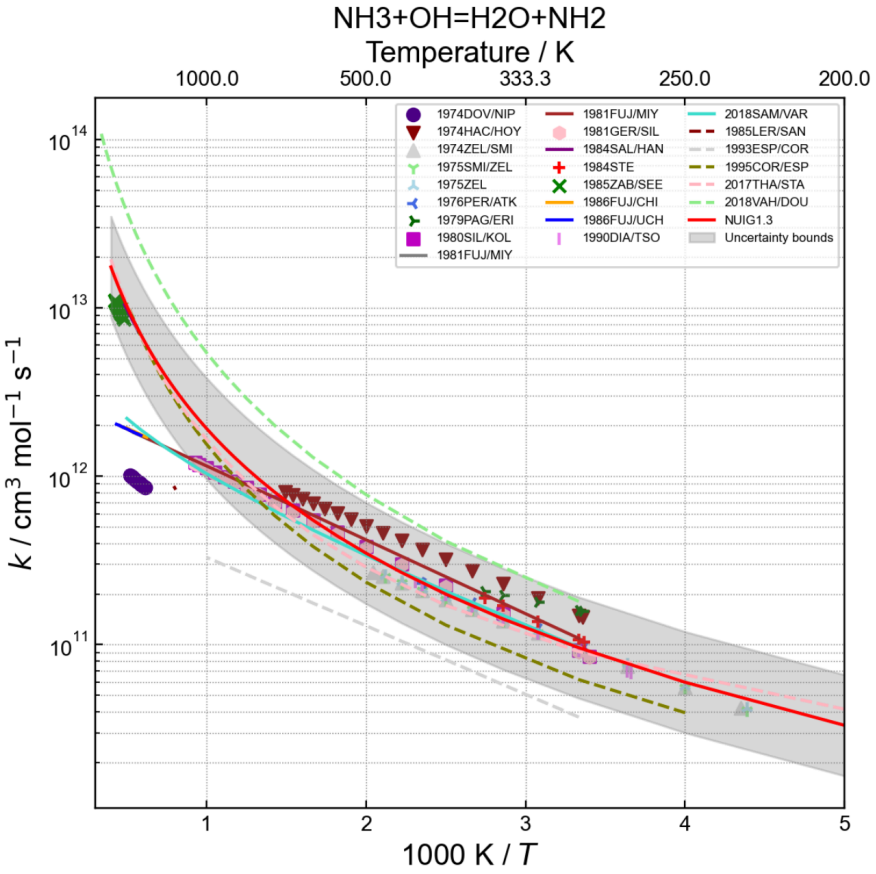

Figure 3.2. Uncertainty factor analysis for reaction #26 in Stagni2023 mechanism file: $NH_3 + OH = NH_2 + H_2O$, considedring an Uncertianty factor $f_r = 0.3$

Default Uncertainty Factor

For reactions lacking sufficient experimental kinetic data in the NIST database (fewer than three independent literature sources), a conservative uncertainty factor of $f_r = 0.5$ was assigned. This value represents approximately a factor of 3.16 variation in either direction ($10^{0.5} \approx 3.16$), which is consistent with typical uncertainties for elementary reactions in complex combustion mechanisms [21].

The choice of $f_r = 0.5$ as the default value was based on the following considerations:

- **Balance between flexibility and constraint**: The factor allows sufficient parameter space exploration to improve model predictions without permitting excessive deviations from the reference mechanism.
- **Theoretical kinetics considerations**: Uncertainties larger than one order of magnitude ($f_r > 1.0$) would be scientifically unjustifiable in the absence of experimental validation, as they could lead to kinetically implausible rate expressions that violate fundamental principles of transition state theory.

- **Optimization stability**: Excessively large uncertainty factors increase the risk of the genetic algorithm converging to local optima that improve fit to the target datasets but degrade predictive capability for untested conditions.

3.5 Genetic Algorithm and Optimisation

PyGAD Introduction

PyGAD (Python Genetic Algorithm [59]) is a Python library that implements genetic algorithms (GAs) for solving optimization problems. Genetic algorithms are population-based metaheuristic optimization techniques inspired by natural evolution, particularly well-suited for complex, high-dimensional, non-linear optimization problems where gradient-based methods may fail or become trapped in local optima. In this work, PyGAD optimizes kinetic parameters by maximizing the mean CM score.

$$\overline{\mathrm{CM}} = \frac{1}{N} \sum_{i=1}^{N} \mathrm{CM}_i \tag{3.8}$$

The algorithm maintains a *population* of candidate solutions, where each solution represents a complete set of kinetic parameters encoded as continuous floating-point genes, each defined within its gene space determined by the REPRICA methodology (Sec. 3.6).

The **fitness function** evaluates each candidate solution through a multi-stage process. First, parameter reconstruction is performed via inverse ICA transformation to obtain Arrhenius parameters (A, n, E_a) for each target reaction. A feasibility check verifies that the resulting rate constants remain within prior uncertainty bounds across the temperature range of 300–3000 K. For feasible solutions, a temporary Cantera mechanism file is created with the proposed parameters, and ignition delay times are simulated for all experimental conditions using a 0D reactor model. Solutions that violate feasibility constraints receive a large penalty (CM = 0.0), effectively excluding them from reproduction. Fitness is in fact coincident with the mean CM score, ensuring that best correspondence between experimental and simulation results receive higher fitness values (approaching 1), while poor or infeasible solutions receive fitness values approaching 0.

Tournament selection ($K = 5$) determines which individuals become parents for the next generation. In each tournament, K individuals are randomly sampled from the population, and the fittest among them is selected as a parent. This process repeats until the required number of parents (defined by `num_parents_mating`) is obtained. The

tournament size K controls selection pressure: smaller values maintain population diversity but slow convergence, while larger values accelerate convergence but risk premature stagnation. A moderate value of $K = 5$ provides balanced exploitation and exploration.

Genetic Operators

The evolution of the population is governed by three main genetic operators: *crossover*, *mutation*, and *elitism*. These operators determine how genetic information is propagated, modified, and preserved across generations, directly influencing the algorithm's ability to explore the search space and converge to optimal solutions.

Crossover combines genetic material from two parent solutions to create offspring, mimicking biological reproduction. This operator enables the algorithm to explore combinations of successful traits from different solutions. For this work, *uniform crossover* is employed, where each gene is independently inherited from either parent with equal probability (50%). This is particularly well-suited for the ICA-transformed parameter space, where independent components are statistically uncorrelated by construction, allowing gene-wise mixing without disrupting meaningful parameter correlations.

The **crossover probability** determines how frequently crossover occurs versus direct cloning of parents. When crossover is applied (e.g., with probability 0.85), offspring receive a mix of genes from both parents; when crossover is skipped (15% of cases), offspring are exact copies of one parent. Higher crossover probabilities (0.8–1.0) promote exploration by creating diverse offspring through recombination, while lower values (0.5–0.7) favor exploitation by preserving successful parent genotypes more frequently. This parameter directly affects the balance between maintaining proven solutions and exploring new parameter combinations.

Mutation introduces random variations in offspring genes, preventing premature convergence and enabling exploration of new regions in the parameter space. This operator is essential for maintaining genetic diversity and escaping local optima. For each offspring, a specified **mutation percentage** of genes is randomly selected and replaced with entirely new random values sampled from their allowed ranges.

The mutation percentage controls the exploration-exploitation trade-off at the individual solution level. Higher percentages (e.g., 50%–60%) provide stronger exploration by introducing more novelty, but risk disrupting good solutions found through crossover. Lower percentages (e.g., 15%–30%) enable finer refinement of existing solutions but may not provide sufficient exploration to escape local optima. For high-dimensional problems like kinetic mechanism optimization, finding the optimal mutation percentage is critical.

Only **random mutation** was explored in this work. Random mutation maintains constant exploration pressure throughout optimization regardless of population state, providing consistent perturbation strength in all generations. This is particularly effective for escaping local optima in rugged fitness landscapes. **Adaptive mutation**, which dynamically adjusts mutation intensity based on population diversity, was deliberately excluded. For high-dimensional problems with many genes, adaptive strategies tend to reduce mutation intensity too aggressively as the population converges, increasing the risk of premature convergence. Additionally, the ICA-constrained parameter space (Sec. 3.6) already provides adaptive behavior by pre-filtering the search region to physically feasible solutions, reducing the need for further adaptation during evolution.

Elitism preserves the best solutions from the current generation by directly copying them to the next generation without modification. The parameter `keep_elitism_percent` specifies what percentage of the population consists of elite individuals. For example, with 15% elitism in a population of 500, the top 75 individuals are guaranteed to survive unchanged, while the remaining 425 slots are filled with new offspring created through crossover and mutation.

This mechanism provides several benefits: it guarantees monotonic improvement in best fitness (the best solution can never worsen), prevents loss of high-quality solutions due to stochastic genetic operations, and accelerates convergence by maintaining "islands of quality" that guide population evolution. However, a trade-off exists: too few elite individuals (e.g., 5%) risk losing good solutions to random chance, while too many (e.g., $> 25\%$) reduce diversity by limiting the number of new offspring, potentially causing premature convergence. The elitism percentage therefore balances solution preservation against population renewal.

Hyperparameter Exploration

The performance of genetic algorithms is highly sensitive to hyperparameter configuration, and optimal settings are problem-dependent. There is no universal "best" configuration because performance depends critically on the specific characteristics of the fitness landscape. Chemical kinetic parameter optimization presents particular challenges: high dimensionality (20 reactions, 60 independent components), non-linear objective functions with complex relationships between Arrhenius parameters and ignition delay or speciation predictions, and expensive fitness evaluations (3–5 seconds per simulation requiring multiple reactor calculations).

Different hyperparameter combinations navigate the exploration-exploitation trade-off in

fundamentally different ways. High mutation combined with weak selection pressure promotes broad exploration but results in slower convergence, while low mutation combined with strong selection pressure enables fast convergence but carries significant risk of becoming trapped around less performing candidate solutions. With a limited computational budget (20–50 generations for preliminary exploration; 500–10,000 generations for final optimization), suboptimal hyperparameters can lead to premature convergence, wasted evaluations on infeasible solutions, or failure to adequately refine solutions.

Isolated Effect of Genetic Operators

Before conducting full factorial hyperparameter exploration, it is essential to understand how each genetic operator independently influences optimization performance. This *one-at-a-time* (OAT) analysis isolates the effect of each operator by varying it across three levels (low, medium, high) while holding the other two operators constant at baseline values. This approach reveals the individual contribution of each operator to convergence behaviour, helping to identify which parameters most critically affect performance and guiding the selection of ranges for subsequent full exploration.

Three independent scenarios isolate mutation, crossover, and elitism effects:

1. **Mutation Effect**: Varying mutation percentage (25%, 50%, 75%) while keeping crossover probability = 0.75 and elitism = 50% constant.
2. **Crossover Effect**: Varying crossover probability (0.25, 0.50, 0.75) while keeping mutation percentage = 50% and elitism = 50% constant.
3. **Elitism Effect**: Varying elitism percentage (25%, 50%, 75%) while keeping mutation percentage = 50% and crossover probability = 0.75 constant.

The baseline configuration (mutation = 50%, crossover = 0.75, elitism = 50%) represents moderate operator intensities, providing a neutral reference point for comparison. Each scenario was evaluated with 60 individuals over 100 generations, yielding 6,000 fitness evaluations per run, enough to observe full convergence behavior. The test configuration used 20 target reactions (60 independent components) with 1 experimental set per observable (STIDT, JSRSP, STSP), matching the optimization problem's complexity.

Performance Metrics. For each configuration, convergence curves were recorded across all 100 generations, tracking the CM score of the best solution at each generation. Key performance indicators extracted from convergence curves include:

- *Final CM (Generation 100)*: Best solution quality after full convergence.

- *Convergence rate*: Steepness of CM improvement in early generations (Gen 0–20), indicating exploration efficiency.
- *Plateau behavior*: CM stability in late generations (Gen 70–100), revealing whether the algorithm is still improving or has stagnated.
- *Smoothness of convergence*: Variability in generation-to-generation CM changes, reflecting balance between exploration (noisy progress) and exploitation (smooth refinement).

The one-at-a-time analysis provides critical insights that inform the design of the subsequent full hyperparameter grid search. By identifying which operators most strongly influence performance and revealing appropriate parameter ranges, this preliminary study ensures that computational resources in the full exploration are allocated efficiently to the most promising hyperparameter regions.

Fine Exploration Strategy

Given the results of the OAT analysis on the genetic operators, a systematic grid search was conducted to identify optimal configurations for the ICA-constrained NH_3-H_2 mechanism optimization problem. The hyperparameter grid tested 27 combinations ($3 \times 3 \times 3$) varying:

- **Mutation percentage**: 20%, 30%, 40%
- **Crossover probability**: 0.2, 0.3, 0.4
- **Elitism percentage**: 30%, 40%, 50%

Fixed parameters included: mutation type (random), crossover type (uniform), tournament size ($K = 5$), number of parents mating (5), population size (60 individuals), and generations (50, sufficient to observe convergence trends). The exploration was performed over 20 target reactions (60 independent components), with 1 experimental set per observable (STIDT [44], JSRSP [49], STSP [52]).

3.6 REPRICA: REaction Parameter Reconstruction via Independent Component Analysis

As demonstrated in Section 2.5, both the Bertolino [33] and Fürst [34] gene space definitions suffer from a fundamental limitation: their rectangular bounds in (α, n, ϵ) space include large fractions of infeasible parameter combinations.

As shown in Table 2.2, it was seen that feasibility rates for a single reaction range from 7% to 61%, with joint feasibility potentially dropping below 1% for multi-reaction optimization.

This infeasibility problem motivates a paradigm shift: rather than defining gene spaces through *analytical bounds* derived from temperature extremes, we propose to *learn the feasible parameter space* directly from data. The REPRICA methodology (REaction Parameter Reconstruction via Independent Component Analysis), developed by Dr. Murakami and Prof. Parente at the Université Libre de Bruxelles [60], achieves this through:

1. *Statistical sampling*: Generate a large ensemble ($N_{\text{samples}} = 10^4$) of candidate parameter sets within selected type of bounds and filter for feasibility, creating a representative dataset of the feasible region.

2. *Dimensionality reduction*: Apply Independent Component Analysis (ICA) to extract statistically independent directions in parameter space, capturing the intrinsic dimensionality of the feasible manifold.

3. *Transformed gene space*: Define genetic algorithm genes as coefficients in the ICA basis rather than raw Arrhenius parameters, ensuring that random combinations of genes naturally produce feasible parameters.

By construction, this approach achieves feasibility rates of 90-95% across all reactions, eliminating correlation-induced infeasibility, and reduces the effective dimensionality of the optimization problem.

Mathematical Formulation

For each target reaction r, we begin by sampling available analytical bounds (e.g., Fürst):

$$\alpha \sim \mathcal{U}(\alpha_{\min}, \alpha_{\max}), \quad n \sim \mathcal{U}(n_{\min}, n_{\max}), \quad \epsilon \sim \mathcal{U}(\epsilon_{\min}, \epsilon_{\max})$$

where the bounds are given by Equations 2.5, 2.8, and 2.9. Each sampled parameter set $\boldsymbol{\theta}_i = (\alpha_i, n_i, \epsilon_i)$ is then tested for feasibility using the envelope constraint (Eq. 2.10):

$$k_{\min}(T) \leq k(T; \boldsymbol{\theta}_i) \leq k_{\max}(T) \quad \forall T \in [T_{\min}, T_{\max}]$$

The feasibility check is performed on a discretized temperature grid with $N_T = 100$ points logarithmically spaced between $T_{\min} = 300$ K and $T_{\max} = 3000$ K. Parameter sets that satisfy the constraint at all grid points are retained, forming the training matrix:

$$\mathbf{X}_{r,\text{train}} = \begin{bmatrix} \alpha_{r,1} & n_{r,1} & \epsilon_{r,1} \\ \alpha_{r,2} & n_{r,2} & \epsilon_{r,2} \\ \vdots & \vdots & \vdots \\ \alpha_{r,N_{\text{feas}}^{(r)}} & n_{r,N_{\text{feas}}^{(r)}} & \epsilon_{r,N_{\text{feas}}^{(r)}} \end{bmatrix}$$

Starting from $N_{\text{samples}} = 10^4$ initial samples, typical feasibility rates (Section 2.5, Table 2.2) yield $N_{\text{feas}}^{(r)} \approx 750$ feasible parameter sets per reaction, providing sufficient statistical power for ICA decomposition.

Independent Component Analysis (ICA)

Independent Component Analysis is a statistical technique that decomposes a multivariate dataset into maximally independent components. Unlike Principal Component Analysis (PCA), which seeks orthogonal directions of maximum variance, ICA identifies directions along which the data has minimal mutual information, capturing non-Gaussian structure and higher-order statistical dependencies. In the context of chemical kinetic parameter optimization, ICA seeks a linear transformation that relates observed parameter variations to statistically independent components.

Data Preprocessing: Before applying ICA, the feasible parameter samples in $\mathbf{X}_{r,\text{train}}$ are preprocessed. In fact, each parameter is shifted so that the average of all feasible samples becomes zero.

$$\mathbf{x}_{r,\text{centered}} = \mathbf{x}_r - \boldsymbol{\mu}_r, \quad \text{where} \quad \boldsymbol{\mu}_r = \frac{1}{N_{\text{feas}}^{(r)}} \sum_{j=1}^{N_{\text{feas}}^{(r)}} \mathbf{x}_{r,j}$$

The mean vector $\boldsymbol{\mu}_r = (\bar{\alpha}_r, \bar{n}_r, \bar{\epsilon}_r)$ represents the geometric center (centroid) of the feasible parameter region for reaction r.

ICA Decomposition: The ICA algorithm (FastICA [61], implemented in `scikit-learn`) is applied separately to each reaction's training data.

The algorithm estimates the mixing matrix $\mathbf{W}_r$ that best separates the original parameter variations into statistically independent components by maximizing their non-Gaussianity, finding directions in parameter space where variation patterns are most distinct and uncorrelated. The complete matrix relationship for reaction r is:

$$\mathbf{X}_{r,\text{centr}} = \mathbf{S}_{r,\text{train}} \mathbf{W}_r^T$$

Parameter Reconstruction and Gene Space Definition

For any feasible parameter candidate of reaction r, the ICA model expresses it as a linear mixture of independent components:

$$\mathbf{x}_r = \mathbf{s}_r \mathbf{W}_r^T + \boldsymbol{\mu}_r \tag{3.9}$$

where:

- $\mathbf{x}_r \in \mathbb{R}^3$ is a feasible parameter vector $(\alpha_r, n_r, \epsilon_r)$ for reaction r;
- $\mathbf{s}_r \in \mathbb{R}^3$ is the vector of independent component coefficients $(c_{r,1}, c_{r,2}, c_{r,3})$;
- $\mathbf{W}_r \in \mathbb{R}^{3\times 3}$ is the mixing matrix for reaction r;
- $\boldsymbol{\mu}_r \in \mathbb{R}^3$ is the mean vector representing the centroid of the feasible parameter region.

In expanded form, the reconstruction equations are:

$$\begin{aligned}
\alpha_r &= \bar{\alpha}_r + w_{r,11}c_{r,1} + w_{r,21}c_{r,2} + w_{r,31}c_{r,3} \\
n_r &= \bar{n}_r + w_{r,12}c_{r,1} + w_{r,22}c_{r,2} + w_{r,32}c_{r,3} \\
\epsilon_r &= \bar{\epsilon}_r + w_{r,13}c_{r,1} + w_{r,23}c_{r,2} + w_{r,33}c_{r,3}
\end{aligned}$$

where $w_{r,jk}$ represents the (j,k)-th element of the transpose of the mixing matrix $\mathbf{W}_r^T$.

Gene Space Bounds: The key insight of REPRICA is to use the independent component coefficients $(c_{r,1}, c_{r,2}, c_{r,3})$ as genetic algorithm genes rather than the Arrhenius parameters themselves. The bounds for each independent component are derived directly from the training data $\mathbf{S}_{r,\text{train}}$:

$$c_{r,j} \in [c_{r,j,\min},\, c_{r,j,\max}] \quad \text{where} \quad \begin{aligned} c_{r,j,\min} &= \frac{100 - \text{cutoff}}{100} \min_{k=1,\ldots,N_{\text{feas}}^{(r)}} s_{r,k,j}, \\ c_{r,j,\max} &= \frac{100 - \text{cutoff}}{100} \max_{k=1,\ldots,N_{\text{feas}}^{(r)}} s_{r,k,j}. \end{aligned}$$

with $s_{r,k,j}$ denoting the (k,j)-th element of $\mathbf{S}_{r,\text{train}}$, and a typical cutoff value of 5% to ensure robustness by preventing sampling at extreme boundaries.

Optimization Process

This entire procedure is repeated independently for all 20 target reactions, producing 20 independent ICA models, each with its own mixing matrix $\mathbf{W}_r$, independent component

dataset $\mathbf{S}_{r,\text{train}}$, and IC bounds. For an optimization target of 20 reactions, the total gene space dimensionality is $3 \times 20 = 60$ independent coefficients.
During optimization:

1. The GA proposes a candidate vector of IC coefficients: $\mathbf{s}_r^{(\text{cand})} = (c_{r,1}, c_{r,2}, c_{r,3})$
2. The Arrhenius parameters are reconstructed using Equation 3.9
3. The rate constant is evaluated using the Arrhenius rate law
4. The mechanism is simulated and the objective function (CM score) is computed

Key Advantage: The mixing matrix $\mathbf{W}_r$ captures the correlation structure of the feasible parameter space, while $\mathbf{S}_{r,\text{train}}$ represents the same feasible data in a decorrelated coordinate system. This transformation enables efficient sampling during optimization: random combinations in the independent component space have a much higher probability ($> 95\%$) of yielding feasible parameters compared to direct sampling in $(\alpha_r, n_r, \epsilon_r)$ space ($\sim 7\%$).

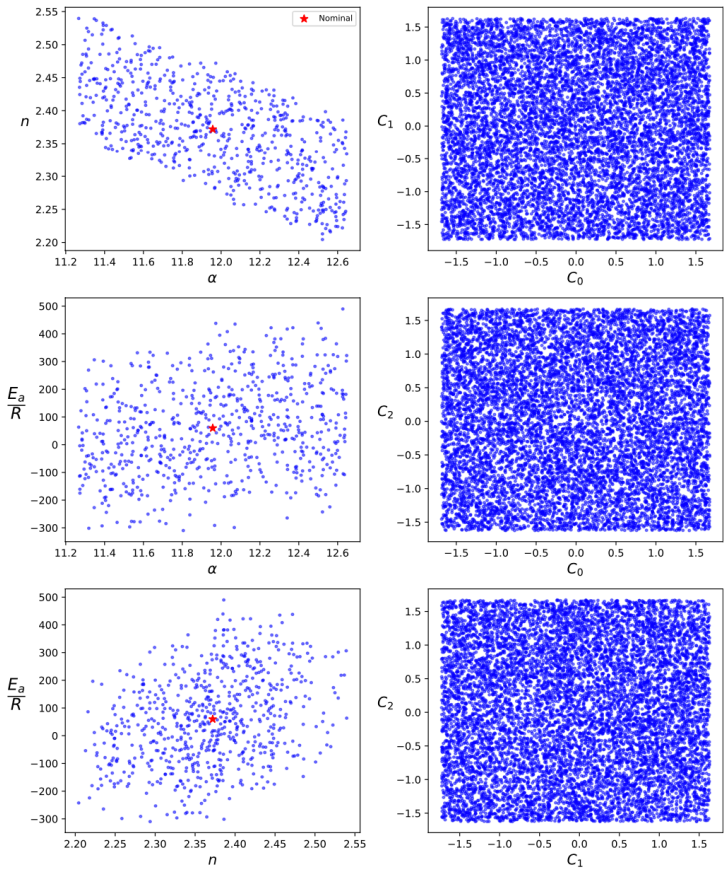

Figure 3.3. Visualization of the ICA transformation for reaction R26. **Left**: Feasible parameter samples in (α, n), (ϵ, α) and (ϵ, n): strong correlation. **Right**: Same samples transformed in independent components: minimal correlation.

Figure 3.3 illustrates the ICA transformation for a representative reaction (R26: NH_2 + NO $\leftrightarrow$ N_2 + H_2O). The feasible parameter samples form a curved, correlated manifold in $(\alpha_r, n_r, \epsilon_r)$ space, which becomes decorrelated and uniformly distributed in the transformed $(c_{r,1}, c_{r,2}, c_{r,3})$ space. This transformation enables more efficient sampling and optimization by working in a coordinate system that respects the natural correlations among kinetic parameters for each individual reaction.

Comparison with Bertolino and Fürst Approaches

The REPRICA approach trades upfront computational cost (ICA sampling and decomposition) for improved optimization efficiency. To quantitatively assess this trade-off and validate the methodology, a systematic comparison was conducted between the three gene space approaches, introduced in Sec. 2.5.

1. **Fürst bounds**: nonlinear, thermodynamically exact at $T_{\min}$ and $T_{\max}$. Extremely low feasibility rates (6.7–7.3%)
2. **Bertolino bounds**: linear approximation, more conservative than Fürst, shows intermediate feasibility (60.7–61.2%)
3. **REPRICA bounds**: data-driven ICA transformation), shows the highest feasibility rates by construction (90-95%)

The comparison was designed to stress-test the methods under optimization conditions. Six experimental datasets were selected spanning the three complementary observables. The choice of a large population size (3000 individuals) evaluated over 50 generations serves two critical purposes. First, it enables the Fürst bounds approach to remain viable despite its severe feasibility constraints: with joint feasibility rates as low as 0.0002% for multi-reaction systems (Section 2.2), such a large population ensures sufficient sampling to discover at least some feasible solutions per generation, thereby preventing complete optimizer stagnation. Second, it provides statistical robustness to the comparison: evaluating 3000 candidates per generation over 50 generations ensures that convergence trends and penalty rate statistics are well-resolved across all three methodologies.
The optimization was configured identically for all three approaches:

- **Target reactions**: Top 3 reactions selected via sensitivity analysis (Sec. 4.1)
- **Hyperparameters**: fixed values of genetic operators (mutation, crossover and elitism) were applied to every optimization run, to allow for the fairest comparison.
- **Experimental targets**: 2 STIDT datasets ([44],[40]), 2 JSRSP datasets ([49], [46]) and 2 STSP datasets ([54], [51])

Advantages

1. **High feasibility rate**: By learning the feasible region directly from envelope-constrained samples, REPRICA achieves more than 95% feasibility across all reactions, compared to 6–7% for Fürst bounds. This drastically reduces penalty evaluations during optimization, allowing computational resources to focus on CM score improvement.

2. **Correlation capture**: The ICA mixing matrix $\mathbf{W}$ automatically encodes nonlinear correlations that would be ignored by rectangular bounds. Genetic mutations in IC space $\mathbf{c} = (c_1, c_2, c_3)$ naturally follow physically plausible trajectories on the Arrhenius surface.

3. **Improved convergence**: Genetic algorithm populations remain in feasible regions more consistently, allowing more effective exploration and faster convergence to higher CM scores.

Limitations

1. **Upfront cost**: ICA sampling requires evaluating 10^4 candidate parameter sets per reaction to build a robust feasible sample. For each candidate, the envelope constraint (Equation 2.10) must be verified across 100 temperature points, in the temperature range of validity.

2. **Statistical dependence on sampling**: The ICA transformation depends on the random sample $\mathcal{D}_{\text{feas}}$. While 10^4 initial samples provide robust statistics, they still stem from Fürst analytical bounds, which do not represent the full available gene space. Therefore, different sampling approaches could also be explored.

3. **Interpretability trade-off**: Unlike Arrhenius parameters $(\ln A, n, E_a/R)$, independent components (c_1, c_2, c_3) lack direct physical meaning.

Despite these limitations, the REPRICA methodology represents a significant advancement in kinetic mechanism optimization, particularly for multi-reaction problems where joint feasibility becomes the dominant bottleneck. The gene space comparison (Chapter 4) provides empirical validation that the trade-offs are favorable.

Connection to Other Workflow Components

The REPRICA gene space integrates seamlessly with the broader optimization workflow:

- **Sensitivity analysis (Section 3.3)**: Target reactions for REPRICA are selected based on sensitivity coefficients, ensuring that computational effort (both ICA preprocessing and GA optimization) is focused on kinetically influential reactions. Reactions with negligible sensitivity are excluded, reducing problem dimensionality.
- **Uncertainty factors (Section 3.4)**: The uncertainty factor f_r defines the initial Fürst bounds used for ICA sampling. The envelope bounds $k_{\min}(T)$ and $k_{\max}(T)$ depend on both nominal parameters and f_r, linking literature uncertainty quantification to the REPRICA gene space construction.
- **Genetic algorithm (Section 3.5)**: The GA operates on IC coefficients $\mathbf{c} = (c_1, c_2, c_3)$ rather than Arrhenius parameters. Mutations and crossovers in IC space automatically respect parameter correlations encoded in each reaction matrix $\mathbf{W}$.
- **Feasibility check**: Even with REPRICA's high feasibility rate, an explicit rate constant check is still performed during fitness evaluation. The low rejection rate captures candidates near gene space boundaries that extrapolate slightly beyond the sampled feasible region due to discretization and ICA linearity limitations.

The *penalty rate* (fraction of candidates rejected) has a significant impact on genetic algorithm efficiency:

- **High penalty rate (Fürst, 90–95%)**: The GA wastes computational effort evaluating infeasible candidates that are immediately discarded. This slows convergence and biases the population toward feasible boundary regions, where parameter space density is highest. The effective population size is reduced by 5–10×, severely limiting exploration.
- **Moderate penalty rate (Bertolino, 40–45%)**: A significant fraction of evaluations are wasted, but the GA can still make progress. Convergence is slower than REPRICA, but not catastrophically so. The conservative linear bounds may exclude globally optimal solutions in high-curvature regions of the Arrhenius surface.
- **Low penalty rate (REPRICA, 1–5%)**: Most candidates are feasible, allowing the GA to focus exploration on CM score optimization. The occasional rejected candidate provides useful information about gene space boundaries without dominating the search. Convergence is both faster and reaches higher CM scores.

The gene space comparison (Chapter 4) quantitatively demonstrates these efficiency differences.

Chapter 4

Results and Discussion

4.1 Sensitivity and uncertainty analysis results

The sensitivity analysis framework described in Section 3.3 was applied to three distinct experimental observables: shock tube ignition delay times (STIDT), jet-stirred reactor species profiles (JSRSP), and spatially-resolved species profiles in premixed flames (SPST). Figure 4.1 presents the global normalized sensitivity coefficients for the top 20 most influential reactions across all three target observables.

Several key observations emerge from the comparative sensitivity analysis:

Shock Tube Ignition Delay Time (STIDT) For the high-temperature shock tube conditions (1000 – 2500 K, 1 – 30 atm), the ignition delay time exhibits dominant sensitivity to two reaction classes:

1. *NH_3 consumption pathways* (R26, R25, R24): The hydroxyl radical attack on ammonia (R26: $NH_3 + OH \leftrightarrow H_2O + NH_2$) shows the highest sensitivity coefficient, reflecting its role as the primary chain-initiating step under these conditions. The additional sensitivity results of ammonia decomposition (R24) and hydrogen abstraction (R25) indicate that both pathways contribute comparably to NH_2 radical formation, which governs the overall ignition delay.

2. *N_2H_2 intermediates* (R111, R113, R112): The elevated sensitivity of diazene chemistry reflects its role as a branch point in the nitrogen radical pool. The decomposition of N_2H_2 to NNH (R111, R112) competes with hydrogen abstraction (R113), with the product NNH radicals subsequently controlling the rate of molecular nitrogen formation versus further oxidation.

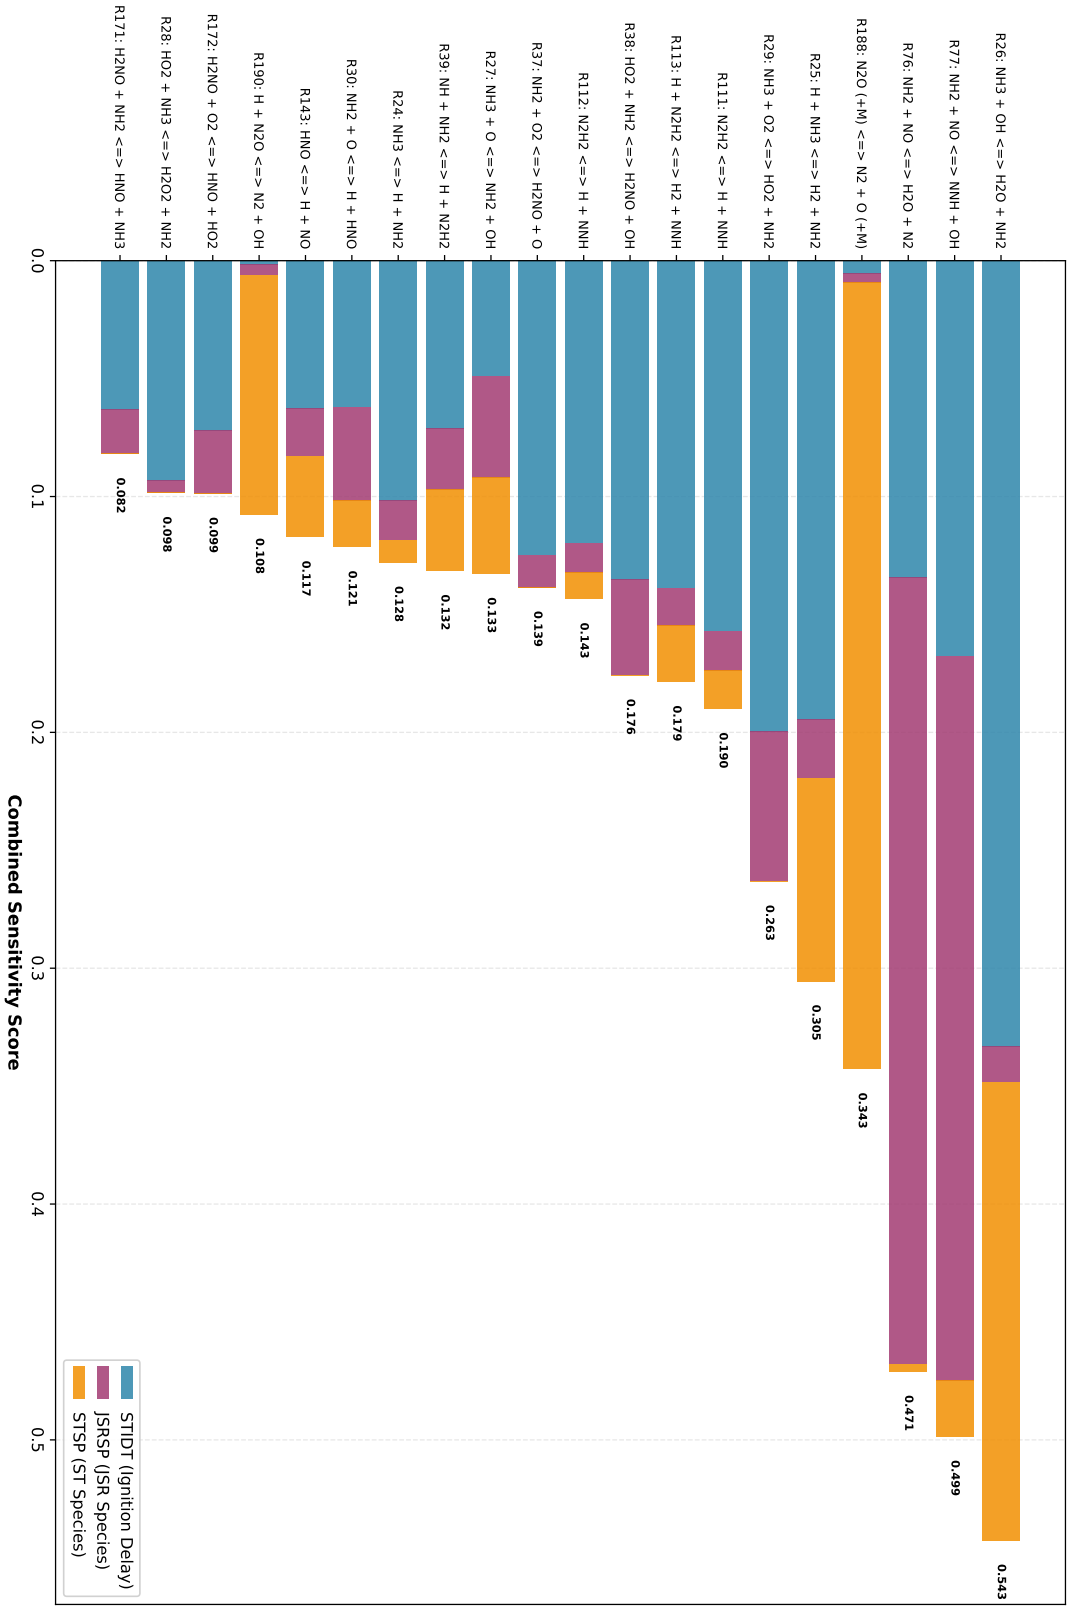

Figure 4.1. Global normalized sensitivity coefficients for the top 20 reactions influencing three key combustion observables: shock tube ignition delay time (STIDT - blue contribution), jet-stirred reactor species profiles (JSRSP - purple contribution), and spatially-resolved flame species profiles (SPST - yellow contribution).

3. *O_2 interaction* (R29, R37): R29 ($NH_3 + O_2 \leftrightarrow HO_2 + NH_2$) allows ammonia to react directly with oxygen molecules to produce NH_2 radicals, without requiring OH radicals. R37 ($NH_2 + O_2 \leftrightarrow H_2NO + O$) then consumes those NH_2 radicals by reacting with more oxygen, producing highly reactive O atoms

Jet-Stirred Reactor Species Profiles (JSRSP) The JSR sensitivity patterns differ substantially from the shock tube results, reflecting the influence of longer residence times and lower temperatures (500 – 1400 K):

1. *NH_2 + NO chemistry dominance* (R76, R77): The most striking feature of JSRSP sensitivity is the overwhelming importance of the NH_2/NO interaction pathways. Reactions R76 ($NH_2 + NO \leftrightarrow N_2 + H_2O$) and R77 ($NH_2 + NO \leftrightarrow NNH + OH$) exhibit sensitivities an order of magnitude higher than all other reactions. This reflects the accumulation of NO species in the quasi-steady reactor environment, where the branching ratio between R76 (forming stable $N_2 + H_2O$) and R77 (producing chain-branching NNH and OH radicals) critically determines both overall reactivity and nitrogen product distribution.

2. *NH_3 oxidation initiation* (R29, R27): Unlike shock tube ignition conditions, R29 ($NH_3 + O_2 \leftrightarrow HO_2 + NH_2$) emerges as the third most sensitive reaction in JSR environments. This pathway serves as the primary NH_3 consumption route at lower temperatures where H/OH radical concentrations are insufficient for rapid fuel decomposition. The elevated sensitivity of R27 ($NH_3 + O \leftrightarrow NH_2 + OH$) further emphasizes the importance of direct fuel attack mechanisms in maintaining steady-state radical populations under oxidizer-limited conditions.

3. *Reduced primary NH_3 decomposition sensitivity* (R26, R25): In contrast to shock tube ignition delay times where R26 ($NH_3 + OH \leftrightarrow H_2O + NH_2$) and R25 ($H + NH_3 \leftrightarrow H_2 + NH_2$) dominate the rankings, these reactions exhibit moderate sensitivity in JSR conditions. This reduction reflects the shift from radical-pool-limited ignition dynamics to oxidizer-limited steady-state chemistry, where rate-determining steps involve subsequent NH_2 radical reactions rather than initial NH_3 decomposition.

Shock Tube Species Profiles (STSP) The STSP sensitivity patterns reveal the importance of high-temperature transient chemistry (1500 - 3000 K) that is not captured by ignition delay measurements alone:

1. *N_2O chemistry dominance* (R188, R190): The most striking feature of STSP sensitivity is the overwhelming importance of N_2O decomposition and consumption

pathways. R188 (N_2O (+M) $\leftrightarrow$ N_2 + O (+M)) emerges as the most sensitive reaction, followed by R190 (H + N_2O $\leftrightarrow$ N_2 + OH) in fifth position. This reflects the role of N_2O in the fuel initial composition [54] as a key species especially for those pyrolisys experimental data for which N_2O serves as the only source of Oxygen in the process.

2. *NH_3 decomposition pathways* (R26, R25, R27): The primary fuel consumption reactions exhibit high sensitivity, consistent with shock tube ignition delay measurements. However, the relative ranking differs from STIDT results, with R26 (NH_3 + OH $\leftrightarrow$ H_2O + NH_2) showing the strongest influence, followed by R25 (H + NH_3 $\leftrightarrow$ H_2 + NH_2) and R27 (NH_3 + O $\leftrightarrow$ NH_2 + OH). This pattern reflects the importance of OH-radical-driven oxidation in the high-temperature, post-ignition environment where species profiles are measured.

The sensitivity analysis results presented here establish the foundation for the uncertainty factor assignment. As a matter of fact, following the semi-automated workflow described in Chapter 3, uncertainty factors were determined for each reaction by comparing the reference mechanism rate constants from `Stagni_2023` with experimental data retrieved from the NIST Chemical Kinetics Database. For each reaction, f_r was selected as the smallest boundary that adequately encompassed the available experimental rate constant data while maintaining physical plausibility. Reactions lacking sufficient experimental coverage were conservatively assigned a default value of 0.5.

The resulting uncertainty factors ensure that the GA optimization search space remains scientifically grounded and consistent with experimental knowledge, preventing exploration of unphysical parameter regions.

Ranked Sensitivity and Uncertainty Factors

Table 4.1 presents the final ranking of the 20 reactions in the `Stagni_2023` mechanism with the greatest influence on ignition delay time, as assessed using Mathieu's experimental shock-tube datasets. For each reaction, both the global normalized sensitivity coefficient and its corresponding uncertainty factor f_r are reported. These values form the quantitative basis for the subsequent kinetic parameter optimization.

Rank.	Reaction	Sensitivity	Uncertainty factor
1	R26: $NH_3 + OH \leftrightarrow H_2O + NH_2$	0.543	0.3
2	R77: $NH_2 + NO \leftrightarrow NNH + OH$	0.499	0.2
3	R76: $NH_2 + NO \leftrightarrow H_2O + N_2$	0.471	0.2
4	R188: N_2O (+M) $\leftrightarrow N_2 + O$ (+M)	0.343	0.5
5	R25: $H + NH_3 \leftrightarrow H_2 + NH_2$	0.305	0.5
6	R29: $0_2 + NH_3 \leftrightarrow HO_2 + NH_2$	0.263	0.5
7	R111: $N_2H_2 \leftrightarrow H + NNH$	0.190	0.5
8	R113: $H + N_2H_2 \leftrightarrow H_2 + NNH$	0.179	0.5
9	R38: $HO_2 + NH_2 \leftrightarrow H_2NO + OH$	0.176	0.5
10	R112: $N_2H_2 \leftrightarrow H + NNH$	0.143	0.5
11	R37: $NH_2 + O_2 \leftrightarrow H_2NO + O$	0.139	0.5
12	R27: $NH_3 + O \leftrightarrow NH_2 + OH$	0.133	0.5
13	R39: $NH + NH_2 \leftrightarrow N_2H_2 + H$	0.132	0.2
14	R24: $NH_3 \leftrightarrow NH_2 + H$	0.128	0.5
15	R30: $NH_2 + O \leftrightarrow H + HNO$	0.121	0.3
16	R143: $HNO \leftrightarrow H + NO$	0.117	0.5
17	R190: $H + N_2O \leftrightarrow N_2 + OH$	0.108	0.5
18	R172: $H_2NO + O_2 \leftrightarrow HNO + HO_2$	0.099	0.5
19	R28: $NH_3 + HO_2 \leftrightarrow H_2O_2 + NH_2$	0.098	0.5
20	R171: $H_2NO + NH_2 \leftrightarrow HNO + NH_3$	0.082	0.5

Table 4.1. Sensitivity analysis and uncertainty factor determination results.

4.2 Gene space comparison

Figure 4.2 shows the evolution of cumulative penalty accumulation and CM score convergence across generations for the three gene space approaches described in Sec. 3.6. The analysis focuses on a reduced optimization scenario with only 3 target reactions to demonstrate the fundamental differences in gene space efficiency.

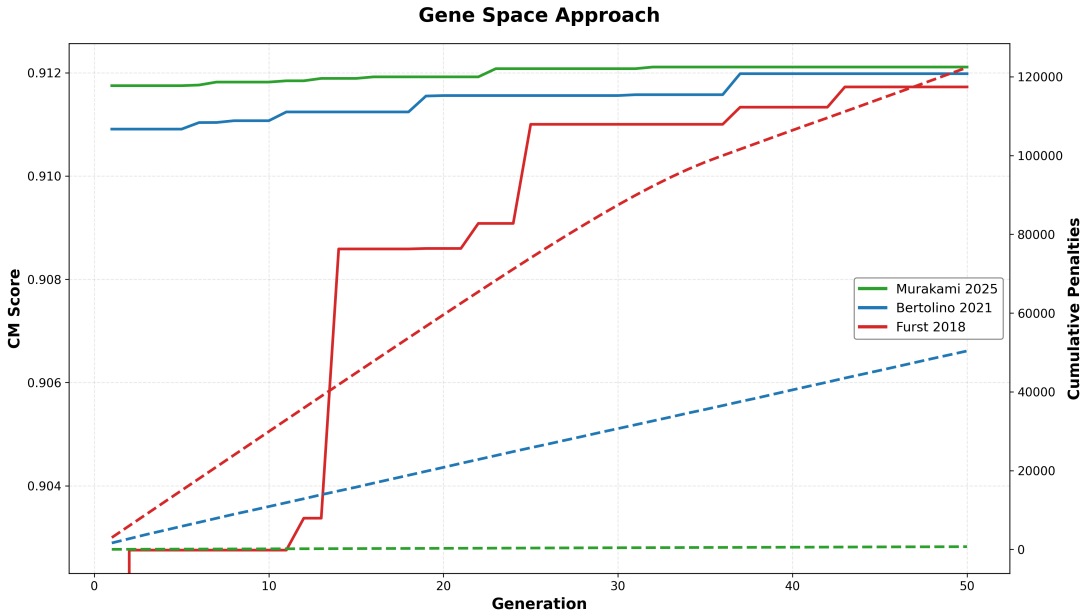

Figure 4.2. Evolution of cumulative penalties (dashed lines, right axis) and best CM score per generation (solid lines, left axis) during genetic algorithm optimization for three gene space approaches: REPRICA (green), Bertolino (blue) and Fürst (red). Population size: 3000 individuals, 6 experimental datasets, top 3 target reactions optimized over 50 generations.

Convergence and Solution Quality

REPRICA demonstrates faster convergence and superior solution quality compared to both competing approaches:

- **Best CM scores**: REPRICA achieves 0.9121 (generation 32), Bertolino 0.9120 (generation 37), Fürst 0.9117 (generation 43).
- **Convergence speed**: REPRICA converges 5 generations earlier than Bertolino (13.5% reduction) and 11 generations earlier than Fürst (25.6% reduction).
- **Initial performance**: From generation 1, REPRICA exhibits superior CM scores (0.9118 vs. 0.9109 for Bertolino; Fürst had zero feasible candidates), maintaining this advantage throughout optimization.

Importantly, **REPRICA is the slowest method per generation** because nearly all 3000 individuals undergo full Cantera simulations, whereas Fürst and Bertolino reject 82% and 34% of candidates before expensive evaluations. However, REPRICA's faster convergence to a higher-quality solution represents better overall resource utilization: reaching CM = 0.9121 in 32 generations versus Fürst requiring 43+ generations to reach only 0.9117.

Penalty Rates and GA Operation

The cumulative penalty curves (dashed lines) reveal fundamental differences in how the three gene spaces interact with genetic operators:

- **Fürst**: 122,325 total penalties over 50 generations (81.6% penalty rate), leaving on average ∼554 feasible individuals per generation for selection.
- **Bertolino**: 50,375 total penalties (33.6% penalty rate), with ∼1,993 feasible individuals per generation.
- **REPRICA**: 697 total penalties (0.47% penalty rate), maintaining ∼2,986 feasible individuals per generation.

For Fürst and Bertolino, the penalty rate decreases over generations. This is a direct consequence of elitism in the genetic algorithm: once feasible solutions with higher CM scores are discovered, they are preserved and propagated, gradually replacing infeasible individuals. This effect is visible as a "kink" or change in slope in the Fürst dashed line, indicating a transition from high initial penalty rates to lower, stabilized rates as the population becomes enriched in feasible solutions. The penalty curve for Fürst is therefore not strictly linear, but shows a marked reduction in slope after the first generations.

This difference in effective population size (2,986 vs. 554) enables REPRICA's genetic operators to function as designed: selection operates on genuine fitness differences rather than on feasibility discovery, and mutations explore parameter correlations encoded in the ICA-transformed gene space rather than randomly probing infeasible boundary regions.

The near-flat cumulative penalty curve for REPRICA (green dashed line) versus the initially steep but then reduced accumulation for Fürst (red dashed line) demonstrates that the ICA gene space maintains alignment with the feasible manifold throughout optimization, while analytical bounds initially generate many infeasible offspring but improve as feasible solutions are preserved.

Computational Efficiency

REPRICA's low penalty rate provides three advantages:

1. **Correct GA operation**: With 99.5% feasibility, genetic operators perform fitness-based evolution rather than feasibility-driven random search.
2. **Larger effective population**: The 2985 feasible individuals per generation maintain genetic diversity and prevent premature convergence.

3. **Superior optimization outcome**: REPRICA achieves the highest CM score, the ultimate optimization goal, demonstrating that effectiveness matters more than per-generation speed.

Early stopping criteria (e.g., terminating when CM improvement falls below 0.0001 for 5 consecutive generations) could reduce REPRICA's runtime by 36% (stopping at generation 32 instead of 50), while Fürst and Bertolino exhibit erratic late-stage convergence unsuitable for early termination.

Summary

The results reveal that gene space design directly controls optimization efficiency. REPRICA's data-driven construction achieves:

1. Faster convergence (32 vs. 37–43 generations)
2. Higher solution quality (CM = 0.9121 vs. 0.9120/0.9117)
3. Near-zero penalty waste (0.5% vs. 34%/82% infeasibility)

These advantages compound when scaling to the full 20-reaction mechanism, where convergence speed and solution quality become critical for computational tractability.

4.3 Hyperparameters exploration results

The one-at-a-time analysis revealed distinct performance patterns for each genetic operator, providing critical insights into their individual contributions to convergence behavior. Figure 4.3 presents the convergence curves across 100 generations for the three scenarios, showing how mutation, crossover, and elitism independently affect optimization performance when varied across low, medium, and high intensity levels.

Mutation Percentage Effect

Mutation percentage exhibited a clear inverse relationship with final solution quality. The low mutation configuration (25%) achieved the highest final CM score of 0.8525, representing a 6.6% improvement over the initial population. In contrast, the high mutation rate (75%) yielded the poorest performance with a final CM of only 0.8430, corresponding to a minimal 0.6% improvement. The medium mutation baseline (50%) produced intermediate performance at CM = 0.8498 (4.9% improvement).

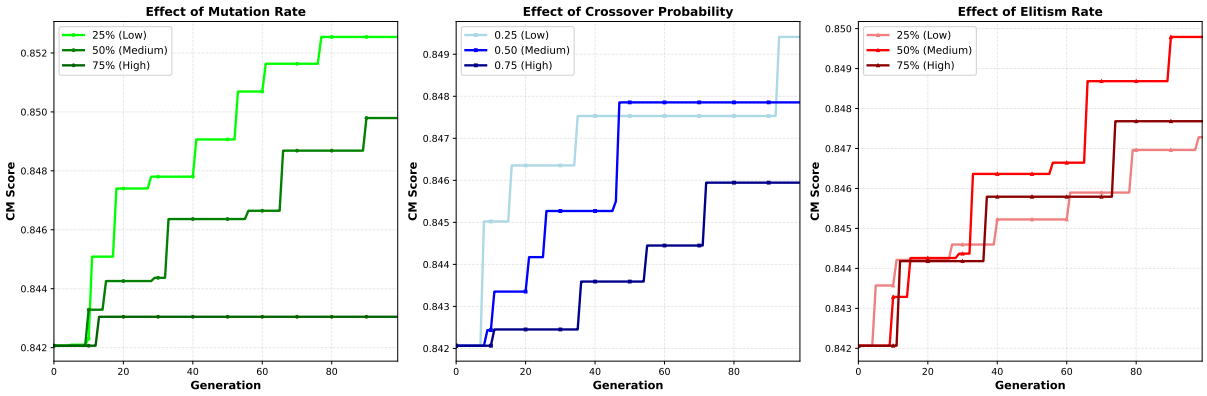

Figure 4.3. Convergence curves for one-at-a-time hyperparameter analysis. Each panel isolates the effect of a single genetic operator on optimization trends.

The convergence dynamics revealed important differences in exploration-exploitation balance: Low mutation demonstrated steady, monotonic improvement with significant advances occurring regularly, suggesting efficient refinement of high-quality solutions with minimal disruption. The high mutation configuration, conversely, showed very limited improvement after generation 13, with the CM score essentially plateauing for the remaining 87 generations. This behavior indicates that excessive mutation introduces too much randomness, preventing the algorithm from effectively exploiting promising regions of the parameter space.

Data suggest that lower mutation intensities (20–30%) better preserve the structured relationships among ICA components while still allowing sufficient exploration. For the subsequent fine exploration strategy, this finding motivated testing mutation percentages in the range of 20–40%, focusing on the lower end of the spectrum where performance was demonstrably superior.

Crossover Probability Effect

Crossover probability showed a similar inverse relationship with performance, though with more subtle differences compared to mutation.

The low crossover configuration (probability = 0.25) achieved the best final CM score of 0.8494 (4.7% improvement), while the high crossover setting (probability = 0.75, baseline) produced the lowest performance at CM = 0.8459 (2.5% improvement). The medium crossover (probability = 0.5) yielded intermediate results with CM = 0.8479 (3.7% improvement).

Convergence analysis revealed that all three crossover configurations exhibited relatively smooth, gradual improvement, with no dramatic differences in convergence rate or sta-

bility. However, the low crossover setting showed a notable improvement spike around generation 93, suggesting that reduced crossover frequencies allow more generations of independent mutation-driven exploration before recombining solutions, potentially discovering higher-quality regions that excessive crossover might average out. The inferior performance of high crossover probabilities likely stems from excessive disruption of well-adapted ICA component combinations. When crossover occurs too frequently (75% of offspring), beneficial candidates may be broken apart before they can be refined. Lower crossover rates (0.2–0.4) preserve these relationships while still enabling periodic exchange of genetic material between high-performing solutions.

This finding directly informed the fine exploration strategy's crossover range (0.2–0.4), deliberately avoiding the high crossover probabilities that proved suboptimal performance.

Elitism Percentage Effect

Elitism presented the most complex and counterintuitive results. The medium elitism configuration (50%) achieved the highest final CM score of 0.8498, while low elitism (25%) reached 0.8473 and high elitism (75%) achieved 0.8477. The differences in final performance were relatively minor, suggesting elitism has less impact on final solution quality than mutation or crossover, within the tested range.

However, elitism exhibited a striking computational efficiency effect: the high elitism configuration (75%) completed 100 generations in only 112 minutes, nearly 50% faster than the low elitism run (331 minutes) and substantially faster than the medium elitism baseline (225 minutes). This dramatic speedup occurs because high elitism preserves 45 of 60 individuals directly from the previous generation, requiring fitness evaluation for only 15 new offspring per generation, compared to 30 evaluations for the baseline configuration. Convergence dynamics revealed distinct behavioral patterns. Low elitism (25%) showed the most gradual, distributed improvement with small gains occurring across many generations, reflecting continuous introduction of new genetic material but risking loss of high-quality solutions. High elitism (75%) displayed more dramatic, discontinuous jumps separated by long plateaus, characteristic of strong exploitation with limited exploration. Medium elitism (50%) balanced these extremes with moderate improvement frequency and the best final performance.

The near-optimal performance of medium elitism (50%), combined with its reasonable computational cost, justified selecting the elitism range of 30–50% for the fine exploration strategy. This range maintains sufficient population diversity (testing 30%, 40%, 50%) while avoiding the slow convergence of low elitism and the potential premature convergence risks of very high elitism.

Synthesis and Implications for Fine Exploration

The OAT analysis yielded three critical insights that directly shaped the design of the subsequent full hyperparameter grid search:

1. **Lower mutation and crossover intensities outperform higher intensities** for this ICA-constrained optimization problem, contrary to typical GA recommendations for high-dimensional problems. This suggests the ICA transformation creates a structured parameter space where excessive stochastic operations are counterproductive.

2. **Mutation exhibits the strongest influence on final solution quality** (6.6% improvement range across tested values), followed by crossover (2.2% range) and elitism (minimal 0.3% range). This hierarchy indicates mutation control is the most critical hyperparameter for achieving optimal performance.

3. **Computational efficiency varies significantly with elitism** (3× difference between extremes), but final performance is relatively insensitive to elitism within the 25–75% range, suggesting this parameter can be tuned primarily for computational cost rather than solution quality.

Based on these findings, the fine exploration strategy focused computational resources on the lower ranges of mutation (20%, 30%, 40%) and crossover (0.2, 0.3, 0.4), where the OAT analysis demonstrated superior performance. The elitism range (30%, 40%, 50%) was centered on the baseline value that produced optimal results. This targeted approach, testing 27 combinations in the most promising regions identified by OAT analysis, maximizes the probability of discovering optimal configurations while avoiding wasteful exploration of parameter ranges already shown to yield inferior performance.

Hyperparameter	**Low**	**Medium**	**High**
Mutation percentage (%)	20	30	40
Crossover probability	0.2	0.3	0.4
Elitism percentage (%)	30	40	50
Total combinations: $3 \times 3 \times 3 = 27$			

Table 4.2. Hyperparameter grid for finer exploration: Ranges informed by the OAT-analysis

Fine Exploration Results: Performance Validation and Optimal Configuration

The focused grid search evaluated all 27 combinations across the identified optimal ranges, completing 1,350 generations of evolution (27 combinations × 50 generations each). Figure 4.4 presents the box-plot distribution of CM scores across the three hyperparameters.

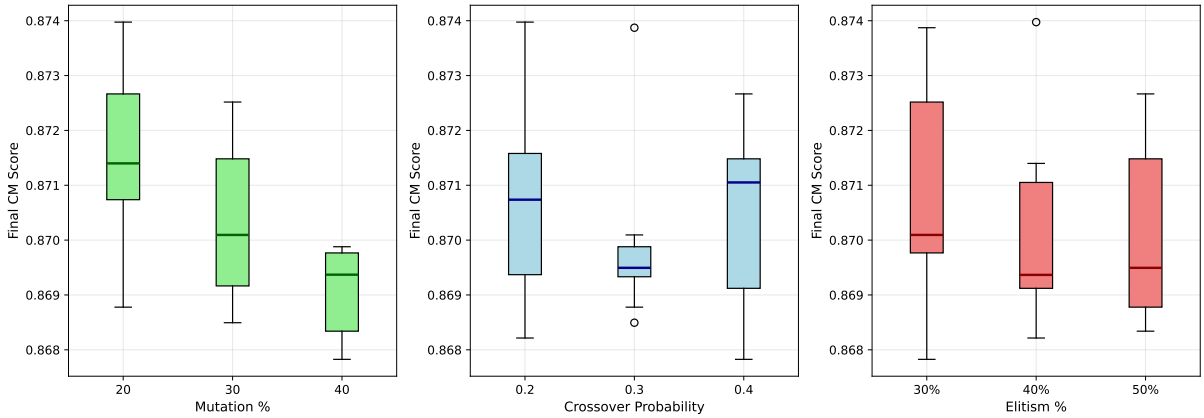

Figure 4.4. Distribution of final CM scores across hyperparameter values. Left: mutation percentage (green); Center: crossover probability (blue); Right: elitism percentage (red).

Box plots display the distribution of results for each hyperparameter value: the colored box encompasses the middle 50% of results, with the horizontal bold line inside marking the median. Vertical whiskers extend to the full range of typical results, while individual circles identify statistical outliers, combinations that performed notably different from their group. The mutation percentage plot reveals the strongest effect: 20% mutation achieves consistently higher CM scores (median ≈ 0.871) compared to 30% (median ≈ 0.870) and 40% (median ≈ 0.869), with tight clustering demonstrating reproducibility. The crossover probability plot shows a weaker but visible trend, with 0.2 slightly outperforming 0.4, consistent with OAT predictions; two outliers at 0.3 indicate specific parameter combinations achieved unexpectedly high performance. The elitism percentage plot proves minimal influence, with all three levels producing nearly overlapping distributions (medians within 0.0007 units), confirming the OAT-predicted insensitivity and validating the use of elitism primarily for computational tuning rather than performance optimization.

Rank	Mutation	Crossover prob.	Elitism	CM
1	20%	0.2	40%	0.8740
2	20%	0.3	30%	0.8739
3	20%	0.4	50%	0.8727
4	20%	0.4	30%	0.8726
5	30%	0.2	30%	0.8725

Table 4.3. Top 5 configurations ranked by CM score.

Table 4.3 compares the top 5 configurations ranked by CM score. The optimal configuration results as (20% mutation, 0.2 crossover, 40% elitism) and the predominant presence of 20% mutation in the top-5 CM performers confirms the critical importance of minimal mutation intensity for this ICA-transformed optimization problem.
The fine exploration strongly validated all three OAT predictions:

1. **Low mutation superiority confirmed**: The majority of top 5 configurations employ 20% mutation, achieving an average CM score of 0.8729 compared to 0.8693 for 40% mutation, demonstrating the superiority of minimal mutation intensity.

2. **Low crossover preference validated**: Configurations with 0.2 crossover probability (average: 0.8707) outperformed those with 0.4 (average: 0.8698), though differences were smaller than mutation effects, consistent with the parameter hierarchy identified in OAT analysis.

3. **Elitism insensitivity confirmed**: Performance variation across elitism levels (30%, 40%, 50%) was minimal (median range: 0.0007 CM units), while computational cost varied predictably (fastest: 50% elitism, slowest: 30% elitism).

The fine exploration's convergence on low mutation/crossover values reveals that ICA's dimensionality reduction creates a parameter landscape where structured search fundamentally outperforms stochastic exploration. Low genetic operator values prevent disruptive jumps in the compressed parameter space while simultaneously improving constraint satisfaction, reducing boundary violations from 163 penalties in the worst configuration to merely 79 in the optimal setting. For the final full-scale optimization, the most performing configuration was adopted as the definitive hyperparameter set, extended to 1000 generations with population size increased to 300 individuals to ensure thorough exploration of the 60-dimensional ICA-transformed space.

4.4 Optimization results

In this section, we present the results of the optimisation procedure, focusing on the CM scores as the primary metric of model performance.
The optimisation was conducted across 12 datasets, comprising 3 STIDT, 5 JSRSP, and 4 STSP cases (Table 3.3). The target reactions were selected as the top 20 identified by the sensitivity analysis, with their respective uncertainty factors (Table 4.1). The optimisation employed a population of 300 individuals per generation, using hyperparameter values determined from the prior hyperparameter exploration to ensure efficient and robust convergence (Table 4.3).

Table 4.4 compare the CM scores obtained from the baseline and optimised models, highlighting the impact of the optimisation strategy.

Observable Type	N cases	Baseline CM Score	Optimised CM Score
STIDT	21	0.926	0.930
JSRSP	85	0.892	0.921
STSP	54	0.878	0.902
TOTAL	160	**0.893**	**0.914**

Table 4.4. Comparison of average CM scores for baseline and optimised models.

The results demonstrate a clear improvement in the mean CM score following the optimisation procedure, with the overall average increasing from 0.893 to 0.914. This corresponds to an approximate relative improvement of 2.1% across all datasets, indicating that, on average, the optimisation strategy was effective in enhancing model performance for the diverse set of experimental cases considered.

A closer examination of the results by observable type reveals that the most substantial relative improvements were achieved for the JSRSP and STSP datasets, with increases of 3.3% and 2.7% in their average CM scores, respectively. In contrast, the STIDT datasets exhibited a more modest improvement of 0.4%. These findings suggest that the optimisation procedure was particularly beneficial for the JSRSP and STSP cases, while STIDT remains the most challenging observable type to optimise.

It is important to note that the reported improvements are based on average CM scores. As such, while the general trend is positive, some individual datasets may have experienced a reduction in performance following optimisation. This phenomenon is not uncommon in multi-objective optimisation scenarios, where trade-offs between different datasets or observables can occur.

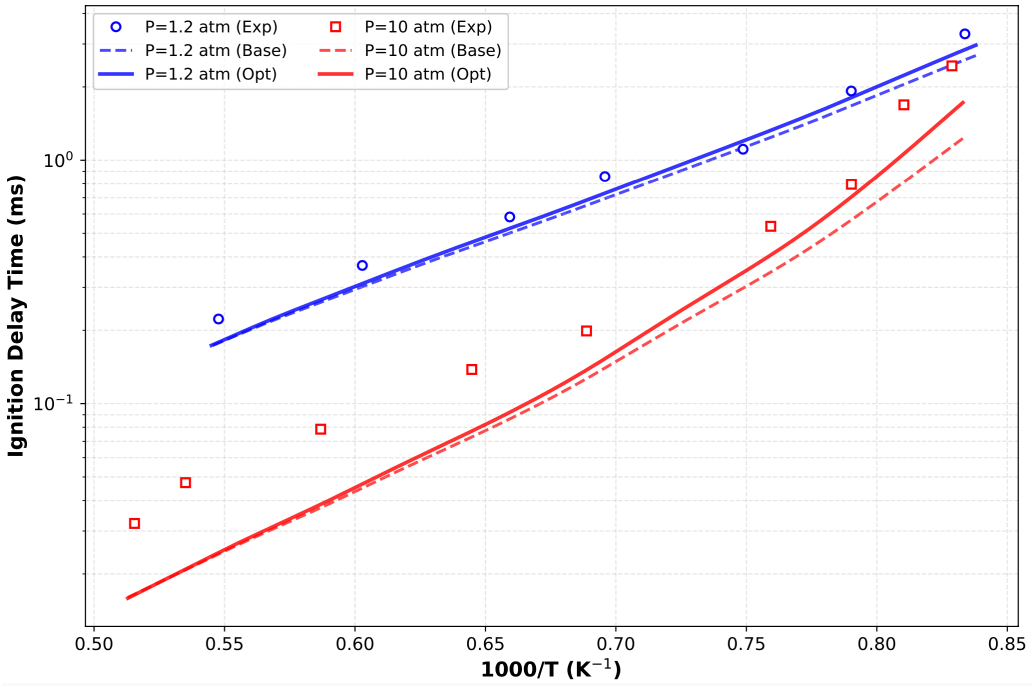

Figure 4.5. Comparison between the Stagni 2023 original mechanism (dashed line) and its optimised version (solid line) on Ignition Delay Time estimation in a Shock Tube. Fuel composition: 70% NH_3 (30% H_2). Results in blue refer to 1.2 atm, while results in red refer to 10 atm. Data from Chen et al.

Figure 4.5 presents a comparison between the Stagni 2023 original mechanism (dashed line) and its optimised version (solid line) on the ignition delay times estimation.

The dataset under discussion comes from Chen et al. [44], which investigates ignition delay times for a fuel mixture containing 70% NH_3 and 30% H_2 under varying pressures (1.2 atm in blue, 10 atm in red).

The results demonstrate that the optimisation yields a noticeable improvement in the agreement between simulation and experimental data at both pressures, with the effect being particularly significant at 10 atm (red curves). However, at higher temperatures (corresponding to the left side of the plot), the advantage of the optimisation diminishes, as the optimised mechanism's predictions converge towards those of the original Stagni 2023 mechanism. This suggests that while the optimisation enhances predictive accuracy overall, its impact is most pronounced at intermediate and lower temperatures, especially at elevated pressures.

Figure 4.6 presents a comparison between the Stagni 2023 original mechanism (dashed line) and its optimised version (solid line) for NO production in a Jet Stirred Reactor at 1 atm, using data from Zhang et al. [48] The mixture has an equivalence ratio of $\phi = 0.25$.

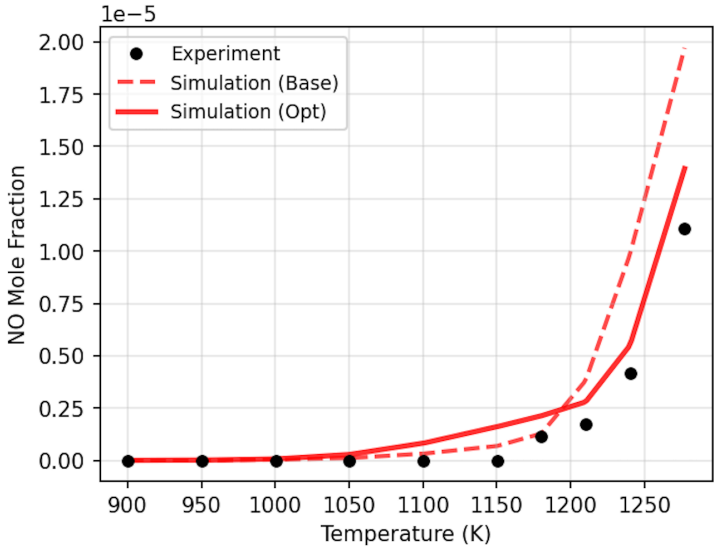

Figure 4.6. Comparison between the Stagni 2023 original mechanism (dashed line) and its optimised version (solid line) on NO production in a Jet Stirred Reactor at 1 atm (Data from Zhang et al.). Mixture equivalence ratio: $\phi = 0.25$.

Optimisation successfully reduced the overprediction of NO production, as evidenced by the improved agreement with experimental data. However, not all species profiles are in perfect alignment with the experiments. This discrepancy may arise from uncertainties on the experimental measurements as well as from limitations in the optimisation process. It is likely that some reactions not included in the optimisation, but which play a non-negligible role in the system's chemistry, should be considered in future optimisation efforts to further enhance model accuracy.

Figure 4.7 compares the performance of the Stagni 2023 original mechanism (dashed line) and its optimised version (solid line) in modelling NH_3 pyrolysis in a shock tube reactor at 2131 K and 1.22 atm, using experimental data from Alturaifi et al. [53]
The optimisation specifically targets the time evolution of NH_3, resulting in improved agreement between simulation and experiment. However, it is important to note that the available experimental data are subject to significant intrinsic uncertainties. In some cases, the experimental trends appear inconsistent or unreliable, and the presence of outliers cannot be excluded. These factors complicate the assessment of model performance and highlight the need for careful interpretation of both simulation and experimental results. Further experimental validation and uncertainty quantification would be beneficial to more robustly evaluate the efficacy of the optimisation.

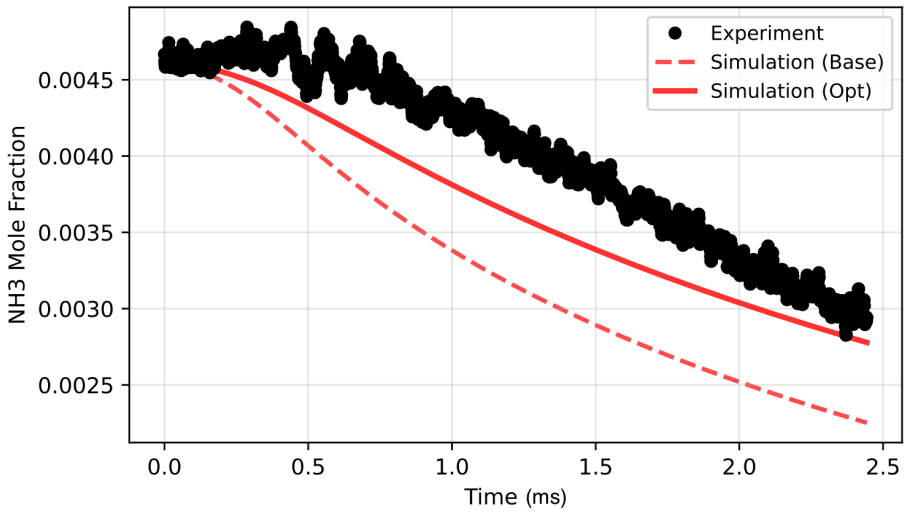

Figure 4.7. Comparison between the Stagni 2023 original mechanism (dashed line) and its optimised version (solid line) on NH_3 pyrolisis in a Shock Tube reactor at 2131 K and 1.22 atm

4.5 Validation results

Validation of the optimised mechanism is fundamental to assess the efficacy of the optimization procedure. The comparison between experimental measurements and simulated laminar burning velocities for ammonia-hydrogen mixtures as a function of the equivalence ratio (ϕ) is presented in Figure 4.8.

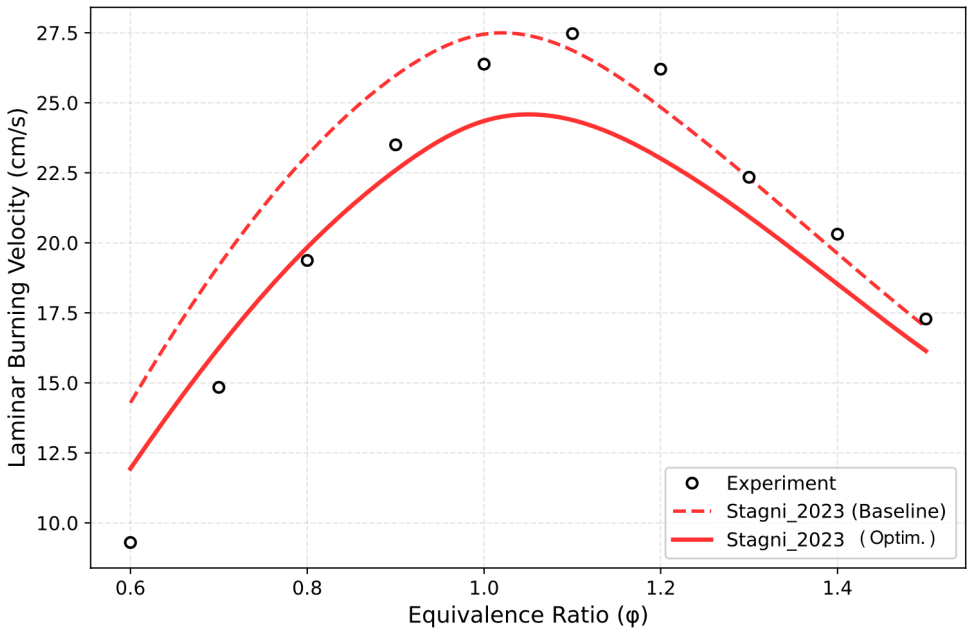

Figure 4.8. Comparison of simulated and experimental laminar burning velocities for ammonia-hydrogen mixtures as a function of equivalence ratio (ϕ).

The experimental data from Mei et al. [55] are depicted as black circles, while simulation results for the baseline Stagni 2023 mechanism are shown as a red dashed line, and those for the optimised mechanism as a solid red line. Both mechanisms capture the characteristic parabolic trend of burning velocity with equivalence ratio, reaching a maximum near stoichiometric conditions ($\phi \approx 1.1$), followed by a decline at fuel-rich mixtures. However, the optimised mechanism shows only marginal improvement over the baseline, particularly failing to capture the experimental peak accurately. This limited enhancement suggests that the reactions primarily governing laminar flame propagation may not have been adequately represented in the target reaction set of 20 species, indicating the need for a more comprehensive selection strategy, potentially through Principal Component Analysis (PCA) rather than relying solely on local sensitivity analysis.

In contrast, the validation results for ignition delay times demonstrate significantly more promising outcomes, as illustrated in Figure 4.9.

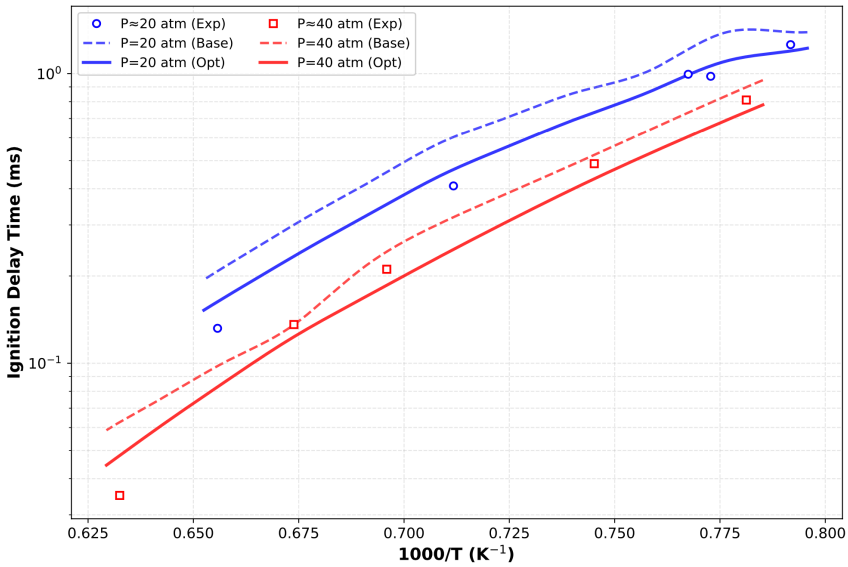

Figure 4.9. Ignition delay times for ammonia-hydrogen mixtures at 20 and 40 atm as a function of inverse temperature ($1000/T$).

This figure presents ignition delay times for ammonia-hydrogen mixtures at elevated pressures (20 and 40 atm) as a function of inverse temperature ($1000/T$), comparing experimental results from Shu et al. [57] with simulations from both the baseline and optimised Stagni 2023 mechanisms. Experimental data are shown as open symbols (circles for 20 atm, squares for 40 atm), while simulation results are represented by dashed (baseline) and solid (optimised) lines for each pressure condition. Both mechanisms reproduce the expected Arrhenius-type behaviour, with ignition delay decreasing exponentially as temperature increases.

Notably, the optimised mechanism yields substantially improved predictions that closely align with experimental values across the entire temperature range, particularly at higher pressures. This successful optimization of ignition delay times represents a significant achievement, as it validates the mechanism's performance in the high-pressure regime where pressure-dependent (PLOG) reactions play a crucial role. The marked improvement underscores the importance of accurately representing pressure-dependent kinetics in ammonia-hydrogen combustion systems and demonstrates the effectiveness of the optimization approach for this particular combustion phenomenon.

4.6 Limitations of the model

Target Reaction Selection and Sensitivity Analysis The present optimization approach relied on expert knowledge and literature-guided selection of target reactions for parameter adjustment. While this methodology successfully identified key rate-limiting steps in H_2/NH_3 combustion, it inherently suffers from potential bias and incomplete coverage of reaction sensitivities across the full experimental matrix. A more systematic approach would employ Principal Component Analysis (PCA) to decompose the sensitivity matrix and identify underlying reaction clusters that collectively govern observable responses. PCA-based sensitivity analysis becomes particularly valuable when analyzing large-scale mechanisms with numerous reactions across many operating conditions, where the resulting sensitivity matrix grows too large for manual interpretation. This statistical decomposition can reveal deeper structural relationships among reactions and identify non-obvious parameter correlations that influence multiple observables simultaneously. While PCA was not employed in the present study due to time constraints, it represents a powerful tool for future mechanism development efforts, particularly when extending the analysis to broader temperature and pressure ranges or additional fuel blends. The integration of PCA for structural relationships among reactions would provide a more comprehensive foundation for target reaction identification and optimization strategy formulation.

Experimental Target Diversity and Computational Constraints The current optimization framework was limited to four observable types (STIDT, JSRSP, STSP, with LBV reserved for validation). Expanding the experimental target spectrum to include Rapid Compression Machine (RCM) data, Plug Flow Reactor (PFR) measurements, and additional LBV datasets would significantly enhance mechanism robustness by constraining rate parameters across broader temperature, pressure, and residence time ranges.

RCM experiments probe the negative temperature coefficient (NTC) regime and intermediate-temperature chemistry, while PFR data capture high-temperature, short-residence-time kinetics relevant to gas turbine combustion. Incorporating these observables, alongside the LBV datasets currently used only for validation, would create a more comprehensive optimization target that spans the full spectrum of practical combustion conditions.

However, this expansion comes with substantial computational drawbacks. RCM simulations require complex heat transfer modeling and pressure-time history matching, while PFR calculations demand spatially-resolved species transport coupled to detailed chemistry. Each additional observable type increases the objective function evaluation time by an order of magnitude, transforming a 24-hour optimization run into a multi-week endeavor. Although high-performance computing clusters with multi-processing capabilities can parallelize individual simulations, the population-based nature of genetic algorithms still requires hundreds of generations to achieve convergence. Current cluster architectures impose strict job time limits (typically 24-168 hours), fragmenting long optimizations into multiple restart cycles and potentially disrupting convergence trajectories. Future optimization efforts must balance experimental target comprehensiveness against computational feasibility, potentially through surrogate modeling, reduced-order kinetic representations, or advanced parallelization strategies that can accommodate the extended wall times required for multi-observable mechanism refinement.

Chapter 5

Conclusions

This thesis presented a comprehensive optimization framework for H_2/NH_3 combustion mechanisms (Fig. 3.1), integrating sensitivity analysis, genetic algorithm optimization, and validation methodologies. The work successfully enhanced the predictive accuracy of the Stagni 2023 mechanism through data-driven parameter refinement while establishing robust methodological foundations for future kinetic mechanism development.

5.1 Summary of the results

Sensitivity Analysis and Target Reaction Identification

The global sensitivity analysis across three distinct combustion regimes, shock tube ignition delay times (STIDT), jet-stirred reactor species profiles (JSRSP), and shock tube species profiles (STSP) revealed fundamental differences in reaction importance across operating conditions (Fig. 4.1). NH_3 consumption pathways dominated STIDT sensitivity (R26, R25, R24), while NH_2/NO chemistry emerged as critical for JSRSP conditions (R76, R77), and N_2O decomposition proved essential for STSP predictions (R188, R190). This analysis successfully identified 20 target reactions spanning the complete combustion chemistry network, providing a scientifically grounded foundation for parameter optimization.

The systematic uncertainty factor assignment methodology, coupling sensitivity rankings with experimental rate constant databases, ensured that optimization remained within physically meaningful parameter bounds (Tab. 4.1). This approach successfully prevented exploration of unphysical parameter regions while maintaining sufficient freedom for meaningful mechanism refinement.

Optimization Methodology: REPRICA Gene Space Development

The development of the REPRICA gene space representation constituted a major methodological advancement, addressing fundamental limitations in existing optimization approaches. Through Independent Component Analysis (ICA) of feasible parameter combinations, REPRICA (Fig. 4.2) achieved near-zero penalty rates (0.47%) compared to traditional methods (34-82%), enabling genetic algorithms to perform true fitness-based evolution rather than feasibility-driven random search.

Comparative analysis demonstrated REPRICA's superior performance across multiple metrics: faster convergence (32 vs. 37-43 generations), higher solution quality (CM = 0.9121 vs. 0.9120/0.9117), and dramatically improved computational efficiency through maintaining 2,986/3000 feasible individuals per generation versus 554-1,993/3000 for competing approaches. These advantages compound significantly when scaling to full 20-reaction optimization problems, making previously intractable optimizations computationally feasible.

Hyperparameter Optimization and Algorithm Tuning

The systematic hyperparameter exploration revealed counterintuitive findings that challenged conventional genetic algorithm wisdom (Fig. 4.3). Lower mutation rates (20%) and crossover probabilities (0.2) consistently outperformed higher intensities, suggesting that ICA's dimensionality reduction creates structured parameter landscapes where excessive stochastic operations prove counterproductive.

The one-at-a-time analysis established clear performance hierarchies: mutation percentage exhibited the strongest influence on solution quality (6.6% improvement range), followed by crossover probability (2.2% range), while elitism primarily affected computational cost rather than final performance. The subsequent 27-combination grid search (Fig. 4.4) validated all OAT predictions and identified the optimal configuration (20% mutation, 0.2 crossover, 40% elitism) that maximized computational efficiency.

Mechanism Optimization Results

The full-scale optimization achieved significant improvements across all target observables (Tab. 4.4), with the overall curve matching score increasing from 0.893 to 0.914 (2.1% relative improvement). More importantly, the improvements were distributed effectively: JSRSP predictions improved by 3.3%, STSP by 2.7%, and STIDT by 0.4%, indicating successful multi-objective optimization without severe performance trade-offs.

Detailed case studies demonstrated the optimization's effectiveness across diverse conditions. Chen et al. ignition delay time predictions improved substantially (Fig. 4.5), particularly at elevated pressures (10 atm), while maintaining accuracy at atmospheric conditions. Zhang et al. JSR species profile predictions showed reduced overprediction of NO formation (Fig. 4.6), and Alturaifi et al. shock tube pyrolysis simulations achieved better NH_3 consumption kinetics despite experimental uncertainty challenges (Fig. 4.7).

Independent Validation and Mechanism Robustness

The validation phase provided critical insights into optimization effectiveness and mechanism transferability. Laminar burning velocity predictions from Mei et al. (Fig. 4.8) experiments showed only marginal improvement, suggesting that flame propagation kinetics require broader reaction coverage than the 20 target reactions identified through ignition-based sensitivity analysis. This finding highlights the importance of observable-specific sensitivity analysis and motivates future PCA-based target selection strategies.

In contrast, high-pressure ignition delay time validation using Shu et al. data (Fig. 4.9) demonstrated exceptional success. The optimized mechanism achieved substantially improved predictions across the full temperature range at both 20 and 40 atm, validating the pressure-dependent reaction optimizations and demonstrating mechanism robustness under extreme conditions not included in the optimization targets.

5.2 Perspectives and future works

This work established several methodological advances with implications extending beyond H_2/NH_3 combustion:

REPRICA Gene Space Architecture: The ICA-based parameter transformation provides a generalizable framework for chemical kinetics optimization. The approach's ability to maintain feasibility while enabling efficient genetic algorithm operation makes it applicable to diverse mechanism optimization challenges across *different fuel systems*.

Integrated Sensitivity-Uncertainty Framework: The systematic coupling of global sensitivity analysis with experimental uncertainty quantification provides a robust methodology for target reaction identification and optimization bound determination.

Multi-Observable Optimization Strategy: The successful integration of three distinct combustion observables (STIDT, JSRSP, STSP) demonstrates the feasibility of comprehensive mechanism optimization across diverse operating regimes.

The improvement distribution validates the multi-objective approach while highlighting the challenges of simultaneous optimization across different combustion phenomena.

However, several limitations emerged that warrant future investigation:

Target Reaction Selection: While sensitivity analysis successfully identified key reactions for ignition-based observables, it proved insufficient for flame propagation kinetics. Future work should employ Principal Component Analysis to decompose sensitivity matrices and identify reaction clusters governing multiple observables simultaneously. This statistical approach would reveal deeper structural relationships and eliminate selection bias.

Experimental Target Diversity: The current framework's limitation to three observable types (STIDT, JSRSP, STSP) restricts mechanism validation scope. Incorporating Rapid Compression Machine data, Plug Flow Reactor measurements, and expanded LBV datasets would enhance robustness across broader temperature, pressure, and residence time ranges. However, the computational burden associated with these observables' simulation, as the order-of-magnitude evaluation time increases, require innovative solutions such as surrogate modeling or advanced parallelization strategies.

Final Remarks

This thesis successfully demonstrated that systematic, data-driven optimization can significantly enhance combustion mechanism predictive accuracy while maintaining scientific rigor and computational feasibility. The developed methodologies provide robust foundations for future mechanism development efforts, with particular relevance to the growing importance of alternative fuels in sustainable energy applications.

The H_2/NH_3 system's complexity, spanning multiple temperature regimes, pressure ranges, and combustion phenomena, makes it an ideal testbed for optimization methodology development. The successful improvements achieved across diverse observables, coupled with the methodological advances in gene space design and hyperparameter optimization, establish a comprehensive framework readily adaptable to other fuel systems.

As the energy sector increasingly emphasizes carbon-neutral fuels, the ability to rapidly and accurately optimize complex kinetic mechanisms becomes crucial for technology development and deployment. This work provides both immediate improvements to H_2/NH_3 combustion modeling and the methodological tools necessary for addressing future optimization challenges in sustainable combustion systems.

Acknowledgements

I would first like to thank my supervisor, Prof. Ernesto Salzano, for giving me the opportunity to carry out my thesis project abroad. This experience has been incredibly valuable both academically and personally, and I am very grateful for the trust and support he has shown me throughout my overall academic journey.

I would also like to thank my co-supervisor, Prof. Alessandro Parente, for his guidance, insightful suggestions, and feedbacks. A special thanks goes to Dr. Yuki Murakami for the time, patience, and help provided during the research process. His advice and availability made a real difference in completing this project.

I am also grateful to my colleagues and friends who shared this final 2 years with me. Their support, encouragement, and the many moments spent together made these challenging years much more enjoyable.

Finally, I would like to thank my family for their constant encouragement and understanding. These years have not always been easy, but they have been full of achievements, growth, and many meaningful encounters with wonderful people. Without you it would not have been the same, and for all of this, I am deeply thankful.

Appendix A

Supplementary Material

The optimized Stagni 2023 kinetic mechanism is available in the GitHub repository linked below.
To ensure reproducibility of the results presented in this thesis, all scripts and source code used for data processing, model development, and evaluation have also been made available.

The repository includes the full implementation of the methods described in this thesis, along with instructions for running the experiments and reproducing the reported results.

Repository: https://github.com/petermalpe00-svg/ATM_Optimization.git

Bibliography

[1] G. R. Group, "Simulation study on combustion performance of ammonia–hydrogen fuel engines," *Energies*, vol. 17, no. 10, p. 2337, 2024.

[2] P. Glarborg, J. A. Miller, B. Ruscic, and S. J. Klippenstein, "Modeling nitrogen chemistry in combustion," *Progress in Energy and Combustion Science*, vol. 67, pp. 31–68, 2018.

[3] C. K. Westbrook and F. L. Dryer, "Chemical kinetic modeling of hydrocarbon combustion," *Progress in Energy and Combustion Science*, vol. 10, pp. 1–57, 1984.

[4] S. Eckart, E. Salzano, A. Richter, M. Alnajideen, A. Valera-Medina, K. P. Shrestha, A. Yasiry, J. Wang, F. Bauer, C. Yu, H. Krause, and G. Pio, "Exploring the potential of ammonia as a fuel: Advances in combustion understanding and large-scale furnace applications," *Fuel*, vol. 403, p. 136746, 2026.

[5] J. J. Hernández, R. Ballesteros, and J. Sanz-Argent, "Reduction of kinetic mechanisms for fuel oxidation through genetic algorithms," *Mathematical and Computer Modelling*, vol. 52, no. 7–8, pp. 1185–1193, 2010.

[6] N. Sikalo, O. Hasemann, C. Schulz, A. Kempf, and I. Wlokas, "A genetic algorithm-based method for the automatic reduction of reaction mechanisms," *International Journal of Chemical Kinetics*, vol. 46, no. 1, pp. 41–59, 2014.

[7] B. Yan, Z. Wu, S. Zhou, J. Lv, X. Liu, W. Wu, and G. Chen, "A critical review of nh_3–h_2 combustion mechanisms," *Renewable and Sustainable Energy Reviews*, vol. 196, p. 114363, 2024.

[8] M. O'Conaire, H. J. Curran, J. M. Simmie, C. K. Westbrook, and W. J. Pitz, "A comprehensive modeling study of hydrogen oxidation," *International Journal of Chemical Kinetics*, vol. 36, no. 11, pp. 603–622, 2004.

[9] S. Zitouni, P. Brequigny, C. Mounaïm-Rousselle *et al.*, "Combustion performances of premixed ammonia/hydrogen/air laminar and swirling flames for a wide range of

equivalence ratios," *International Journal of Hydrogen Energy*, vol. 47, pp. 41 170–41 182, 2022.

[10] A. Stagni, C. Cavallotti, S. Arunthanayothin, Y. Song, O. Herbinet, F. Battin-Leclerc, and T. Faravelli, "An experimental, theoretical and kinetic-modeling study of the gas-phase oxidation of ammonia," *Reaction Chemistry & Engineering*, vol. 5, pp. 696–711, 2020.

[11] A. Alnasif, J. Jójka, S. Mashruk, T. Nagy, and A. Valera-Medina, "Analysis of the performance of kinetic reaction mechanisms in estimating n_2o mole fractions in 70/30 vol% nh_3/h_2 premixed flames," *Fuel*, 2024.

[12] K. Shrestha, C. Lhuillier, A. A. Barbosa, P. Brequigny, F. Contino, C. Mounaïm-Rousselle, L. Seidel, and F. Mauss, "An experimental and modeling study of ammonia with enriched oxygen content and ammonia/hydrogen laminar flame speed at elevated pressure and temperature," *Proceedings of the Combustion Institute*, 2021.

[13] A. Alnasif, S. Mashruk, M. Hayashi, J. Jójka, H. Shi, A. Hayakawa, and A. Valera-Medina, "Performance investigation of currently available reaction mechanisms in the estimation of no measurements: A comparative study," *Energies*, 2023.

[14] S. Chen, Z. Yu, Y. Wang, Y. Ai, and C. Liu, "Experimental investigation of the flame propagation speeds of ammonia/air and ammonia/hydrogen/air mixtures at elevated temperatures," *Journal of Thermal Science*, vol. 34, no. 4, pp. 1497–1511, 2025.

[15] G. J. Gotama, A. Hayakawa, E. Okafor, R. Kanoshima, M. Hayashi, T. Kudo, and H. Kobayashi, "Measurement of the laminar burning velocity and kinetics study of the importance of the hydrogen recovery mechanism of ammonia/hydrogen/air premixed flames," *Combustion and Flame*, 2022.

[16] J. A. Miller and C. T. Bowman, "Mechanism and modeling of nitrogen chemistry in combustion," *Progress in Energy and Combustion Science*, vol. 15, pp. 287–338, 1989.

[17] J. A. Miller and P. Glarborg, "Modeling the thermal de-nox process: Closing in on a final solution," *International Journal of Chemical Kinetics*, vol. 31, pp. 757–765, 1999.

[18] A. A. Konnov, "Implementation of the ncn pathway of prompt-no formation in the detailed reaction mechanism," *Combustion and Flame*, vol. 156, pp. 2093–2105, 2009.

[19] P. Dagaut, P. Glarborg, and M. Alzueta, "The oxidation of hydrogen cyanide and related chemistry," *Progress in Energy and Combustion Science*, vol. 34, pp. 1–46, 2008.

[20] A. Stagni, G. Savarese, S. Arunthanayothin, Y. Song, O. Herbinet, F. Battin-Leclerc, T. Faravelli, and A. Frassoldati, "Low- and intermediate-temperature ammonia/hydrogen oxidation in a flow reactor: experiments and a wide-range kinetic modeling," *Chemical Engineering Journal*, vol. 471, p. 144577, 2023.

[21] A. Tomlin and T. Turányi, "Investigation of rate coefficient uncertainty in chemistry models using sensitivity analysis," *Annual Reports Section C (Physical Chemistry)*, vol. 109, pp. 50–91, 2013.

[22] I. M. Sobol, "Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates," *Mathematics and computers in simulation*, vol. 55, no. 1-3, pp. 271–280, 2001.

[23] M. D. Morris, "Factorial sampling plans for preliminary computational experiments," *Technometrics*, vol. 33, no. 2, pp. 161–174, 1991.

[24] D. L. Baulch, C. T. Bowman, C. J. Cobos, R. A. Cox *et al.*, "Evaluated kinetic data for combustion modelling," *Journal of Physical and Chemical Reference Data*, vol. 34, no. 3, pp. 757–1397, 2005.

[25] S. J. Klippenstein, L. B. Harding, J. Zheng *et al.*, "Uncertainty-driven theoretical kinetics for combustion modeling," *Progress in Energy and Combustion Science*, vol. 55, pp. 1–57, 2017.

[26] J. A. Manion, R. E. Huie, R. D. Levin *et al.*, "Nist chemical kinetics database," National Institute of Standards and Technology, Tech. Rep., 2015.

[27] W. Tsang and R. F. Hampson, "Chemical kinetic data base for combustion chemistry," in *Journal of Physical and Chemical Reference Data*, 1986, vol. 15, pp. 1087–1279.

[28] A. Tomlin, T. Turanyi, and M. Pilling, "The role of sensitivity analysis in combustion modeling," *Combustion and flame*, vol. 115, pp. 436–461, 1998.

[29] T. Turányi, "Uncertainty analysis," *HDMR modeling*, vol. 101, pp. 80–96, 2009.

[30] D. A. Sheen and H. Wang, "The method of uncertainty quantification and its application to combustion chemistry," *Combustion and Flame*, vol. 162, no. 3, pp. 835–851, 2015.

[31] R. Zhai, J. Xu, Z. Luo *et al.*, "Combustion-optimized kinetic model development using uncertainty quantification and genetic algorithms," *Combustion and Flame*, vol. 226, pp. 83–96, 2021.

[32] C. Tülpan, T. Varga, and T. Turányi, "Uncertainty evaluation in chemical kinetic mechanisms," *International Journal of Chemical Kinetics*, vol. 51, pp. 637–665, 2019.

[33] A. Bertolino, M. Fürst, A. Stagni, C. Cavallotti, A. Parente, and T. Faravelli, "An evolutionary, data-driven approach for mechanism optimization: theory and application to ammonia combustion," *Combustion and Flame*, vol. 229, p. 111366, 2021.

[34] M. Fürst, P. Sabia, M. Lubrano Lavadera, G. Aversano, M. de Joannon, A. Frassoldati, and A. Parente, "Optimization of chemical kinetics for methane and biomass pyrolysis products in moderate or intense low-oxygen dilution combustion," *Energy & Fuels*, vol. 32, pp. 10 194–10 201, 2018.

[35] B. M. Adams, M. S. Ebeida, M. S. Eldred *et al.*, "Dakota: A multilevel parallel object-oriented framework for design optimisation, parameter estimation, uncertainty quantification, and sensitivity analysis," 2014, version 6.5 User's Manual.

[36] L. Elliott, D. B. Ingham, A. G. Kyne, N. S. Mera, and M. Pourkashanian, "Genetic algorithms for optimisation of chemical kinetics reaction mechanisms," *Progress in Energy and Combustion Science*, vol. 30, no. 3, pp. 297–328, 2004.

[37] D. A. DelVescovo, J. Li, D. A. Splitter, F. Dal Forno Chuahy, and P. Zhao, "Genetic algorithm optimization of a chemical kinetic mechanism for propane at engine relevant conditions," *Fuel*, vol. 338, p. 127371, 2023.

[38] M. S. Bernardi, M. Pelucchi, A. Stagni, L. M. Sangalli, A. Cuoci, A. Frassoldati, P. Secchi, and T. Faravelli, "Curve matching, a generalized framework for models/experiments comparison: An application to n-heptane combustion kinetic mechanisms," *Combustion and Flame*, vol. 168, pp. 186–203, 2016.

[39] M. E. O'Connor *et al.*, "Rate coefficients for oh + no (+n2) in the fall-off regime and the impact of water vapor," *Journal of Physical Chemistry A*, vol. 127, 2023.

[40] O. Mathieu and E. L. Petersen, "Experimental and modeling study on the high-temperature oxidation of ammonia and related NO_x chemistry," *Combustion and Flame*, vol. 162, no. 3, pp. 554–570, 2015.

[41] X. Han, Z. Wang, M. Costa, Z. Sun, Y. He, and K. Cen, "Experimental and kinetic modeling study of laminar burning velocities of NH_3/air, NH_3/H_2/air, NH_3/CO/air

and NH_3/CH_4/air premixed flames," *Combustion and Flame*, vol. 206, pp. 214–226, 2019.

[42] Y. Song, H. Hashemi, J. M. Christensen, C. Zou, P. Marshall, and P. Glarborg, "Ammonia oxidation at high pressure and intermediate temperatures," *Fuel*, vol. 181, pp. 358–365, 2016.

[43] A. Lucassen, P. Oßwald, U. Struckmeier, K. Kohse-Höinghaus, T. Kasper, N. Hansen, T. A. Cool, and P. R. Westmoreland, "Species identification in a laminar premixed low-pressure flame of morpholine as a model substance for oxygenated nitrogen-containing fuels," *Zeitschrift für Physikalische Chemie*, vol. 226, no. 11–12, pp. 1487–1513, 2012.

[44] J. Chen, X. Jiang, X. Qin, and Z. Huang, "Effect of hydrogen blending on the high temperature auto-ignition of ammonia at elevated pressure," *Fuel*, vol. 287, p. 119563, 2021.

[45] J. Baker, R. K. Rahman, M. Pierro, J. Higgs, J. Urso, C. Kinney, and S. Vasu, "Experimental ignition delay time measurements and chemical kinetics modeling of hydrogen/ammonia/natural gas fuels," *Journal of Engineering for Gas Turbines and Power*, vol. 145, no. 4, p. 041002, 2023.

[46] R. Rota, D. Antos, E. Zanoelo, and S. Carra, "Experimental study and kinetic modelling of nitric oxide reduction with ammonia," *Combustion Science and Technology*, vol. 163, no. 1, pp. 25–47, 2001.

[47] A. Stagni, C. Cavallotti, S. Arunthanayothin, Y. Song, O. Herbinet, F. Battin-Leclerc, and T. Faravelli, "An experimental, theoretical and kinetic-modeling study of the gas-phase oxidation of ammonia," *Reaction Chemistry & Engineering*, vol. 5, no. 4, pp. 696–711, 2020.

[48] X. Zhang, S. P. Moosakutty, R. P. Rajan, M. Younes, and S. M. Sarathy, "Combustion chemistry of ammonia/hydrogen mixtures: Jet-stirred reactor measurements and comprehensive kinetic modeling." *Combustion and Flame*, vol. 234, p. 111653, 2021, jSR experiments reported in supplementary material.

[49] K. N. Osipova, X. Zhang, S. M. Sarathy, O. P. Korobeinichev, and A. G. Shmakov, "Ammonia and ammonia/hydrogen blends oxidation in a jet-stirred reactor: experimental and numerical study," *Fuel*, vol. 310, p. 122202, 2022.

[50] R. Tang, Q. Xu, J. Pan, J. Gao, Z. Wang, H. Wei, and G. Shu, "An experimental and modeling study of ammonia oxidation in a jet stirred reactor," *Combustion and Flame*, vol. 240, p. 112007, 2022.

[51] K. Kohse-Hoinghaus, D. F. Davidson, A. Y. Chang, and R. K. Hanson, "A pyrolysis mechanism for ammonia," *International Journal of Chemical Kinetics*, vol. 22, no. 5, pp. 513–535, 1990, time-resolved NH and NH_2 measurements via laser absorption spectroscopy behind reflected shocks (2200–2800 K).

[52] C. Halat-Augier, P. Lecouvreur, C. Paillard, G. Dupré, F. Heilmann, and L. Lemoullec, "Shock tube study of NH_3 oxidation," in *Colloque International de Propulsion des Véhicules Spatiaux*, Arcachon, France, 1994, proceedings paper reporting shock-tube measurements of hydrazine (N_2H_4) thermal decomposition under inert conditions; no DOI available.

[53] S. A. Alturaifi, O. Mathieu, and E. L. Petersen, "An experimental and modeling study of ammonia pyrolysis," *Combustion and Flame*, vol. 235, p. 111694, 2022, includes NH_3 and H_2/NH_3 time-resolved species measurements.

[54] O. Mathieu, C. M. Gregoire, and E. L. Petersen, "Shock-tube study of the oxidation of ammonia by n_2o," *Proceedings of the Combustion Institute*, vol. 40, p. 105250, 2024, time-resolved NH_3 concentration measurements via laser absorption.

[55] B. Mei, X. Zhang, S. Ma, M. Cui, H. Guo, Z. Cao, and Y. Li, "Experimental and kinetic modeling investigation on the laminar flame propagation of ammonia under oxygen enrichment and elevated pressure conditions," *Combustion and Flame*, vol. 210, pp. 236–246, 2019.

[56] B. Mei, S. Ma, X. Zhang, X. Zhang, W. Li, and Y. Li, "Characterizing ammonia and nitric oxide interaction with outwardly propagating spherical flame method," *Proceedings of the Combustion Institute*, vol. 38, pp. 2477–2485, 2021.

[57] B. Shu, S. K. Vallabhuni, X. He, G. Issayev, K. Moshammer, A. Farooq, and R. X. Fernandes, "A shock tube and modeling study on the autoignition properties of ammonia at intermediate temperatures," *Proceedings of the Combustion Institute*, vol. 37, no. 1, pp. 205–211, 2019, published online 2018.

[58] M. P. Burke, M. Chaos, Y. Ju, F. L. Dryer, and S. J. Klippenstein, "Comprehensive H_2/O_2 kinetic model for high-pressure combustion," *International Journal of Chemical Kinetics*, vol. 44, no. 7, pp. 444–475, 2012.

[59] A. F. Gad, "PyGAD: An intuitive genetic algorithm python library," https://github.com/ahmedfgad/GeneticAlgorithmPython, 2021, version 3.x.

[60] Y. Murakami, A. Bertolino, and A. Parente, "Enhancing chemical kinetic mechanism optimization using independent component analysis," *Fuel*, vol. 418, 2026.

[61] A. Hyvärinen and E. Oja, "Independent component analysis: Algorithms and applications," *Neural Networks*, vol. 13, no. 4-5, pp. 411–430, 2000.