Control Engineering Practice 175 (2026) 107098

Contents lists available at ScienceDirect

Control Engineering Practice

journal homepage: www.elsevier.com/locate/conengprac

Explicit learning stochastic model predictive control with general ambiguity sets in homogeneous charge compression ignition applications

Xu Chen [a,*], Julian Bedei [b], Jakob Andert [b], Lorenz Dörschel [a], Heike Vallery [a]

[a] Institute of Automatic Control, RWTH Aachen University, Campus-Boulevard 30, 52074, Aachen, Germany
[b] Chair of Mechatronics in Mobile Propulsion, RWTH Aachen University, Forckenbeckstrasse 4, 52074, Aachen, Germany

ARTICLE INFO

Keywords:
Model predictive control
Stochastic MPC
Machine learning
Low-temperature combustion
Homogeneous charge compression ignition

ABSTRACT

In pursuit of more efficient and cleaner combustion, the field of Homogeneous Charge Compression Ignition is under active investigation. This combustion approach offers lower nitrogen oxide emissions and higher efficiency than conventional gasoline engines. However, two major challenges hinder its implementation: the stochastic and auto-regressive nature of the process, which complicates accurate prediction and control; and the difficulty of deploying controllers in real time, due to model complexity and stringent timing constraints. Existing MPC-based approaches primarily focus on control-oriented modeling to capture more intricate system dynamics. Meanwhile, many optimization-based controllers rely on outsourcing computations to meet real-time constraints. To overcome the first challenge, we introduce Nonlinear Moment-Based Distributionally Robust Stochastic MPC (NoM-DiR-SMPC), a novel stochastic MPC variant based on a general ambiguity set, and reformulate the MPC problem into a tractable optimization problem. To tackle the second challenge, we propose a learning-based control scheme inspired by the MPC structure, which efficiently approximates the control strategy using a parameterized model. In our test bench study, we achieved real-time control using a local control unit and observed improved combustion stability. In the simulation study, incorporating process uncertainties helped reduce combustion irregularities, specifically, constraint violation by 33 % and the CA50 variation by 16.3 % compared to the classical MPC strategy.

1. Introduction

Internal combustion engines continue to drive innovation aimed at improving energy conversion efficiency and reducing high emissions. In this context, Homogeneous Charge Compression Ignition (HCCI) remains a highly relevant platform for advanced combustion control research (Kulzer et al., 2006). HCCI offers the potential for improved combustion efficiency while reducing peak combustion temperatures and thereby suppressing the formation of nitrogen oxides (Saxena & Bedoya, 2013; Zhao, 2007). In this work, however, HCCI is not chosen because it represents the dominant production pathway, but because it provides a stringent benchmark for control design under strong stochastic variability, cycle-to-cycle coupling, and real-time constraints.

Unlike traditional spark- and compression-ignition engines, HCCI lacks a direct ignition mechanism. The air, residual gases, and fuel in the HCCI process are well-mixed, allowing the assumption of a homogenized mixture. When this mixture reaches the appropriate boundary conditions for auto-ignition, combustion occurs almost simultaneously throughout the chamber, thus improving the fuel conversion rate (Oakley et al., 2001).

This dependence on auto-ignition poses substantial control challenges to maintain stable combustion in HCCI engines. Auto-ignition depends strongly on the thermodynamic state, rendering the process sensitive to state variations, which could arise from external conditions and the changing properties of the residual gas (Saxena & Bedoya, 2013). Compression alone may not suffice to achieve the necessary conditions for auto-ignition. Therefore, additional thermal energy is required, such as the intake air heating and external or internal exhaust gas recirculation (EGR) (Wei et al., 2012).

One effective approach to increase thermal energy is Negative Valve Overlap (NVO), which involves temporarily closing both intake and exhaust valves around the Top Dead Center (TDC) to trap hot exhaust gases from the previous cycle (Koopmans et al., 2003). This strategy enables a TDC-symmetrical, throttle-free operation that reduces pumping losses and improves thermal efficiency, thereby promoting stable auto-ignition, particularly under low load.

* Corresponding author.
E-mail address: x.chen@irt.rwth-aachen.de (X. Chen).

https://doi.org/10.1016/j.conengprac.2026.107098
Received 15 September 2025; Received in revised form 16 April 2026; Accepted 1 June 2026
Available online 10 June 2026
0967-0661/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

The introduction of NVO also brings control challenges. Firstly, the gas composition and thermodynamic states across successive engine cycles are strongly coupled, which increases system complexity under state and input constraints and makes the dynamics highly sensitive to state uncertainty and external disturbances, leading to stochastic cycle-to-cycle variability. Secondly, since NVO serves as the primary control input, control decisions must be made before the desired exhaust valve closes. This timing constraint requires controllers with low computational complexity to ensure real-time feasibility.

As an effective method to handle state and input constraints, Model Predictive Control (MPC) has been used widely for HCCI (Bengtsson et al., 2006a,b; Ebrahimi & Koch, 2018), with commendable performance outcomes. However, due to the high complexity of the process and the stringent real-time computation constraints, guaranteeing combustion stability and extending the operating range remained major challenges. In an attempt to expand the operating range, several studies in HCCI explored MPC strategies augmented with spark plug assistance based on linear models (Albin et al., 2013, 2012; Ritter et al., 2018). Although these approaches improved controllability in certain regimes, experimental investigations still reported irregular combustion behavior, indicating that simplified linear control formulations may be insufficient. Neural Network-based Nonlinear MPC approaches (Gordon et al., 2024; Janakiraman et al., 2015) alleviate this shortcoming, showing high consistency between predicted and measured responses. However, these methods often require computational outsourcing to manage the complexity of NMPC optimization problems. In relation to this work, Nuss et al. (2019) provides a comprehensive simulation analysis of NMPC, while De Schutter et al. (2022) examines the feasibility of deploying NMPC on local controllers, taking into account computational capability and time constraints. These studies highlight that, although nonlinear predictive control is promising for HCCI combustion, achieving both sufficient control fidelity and real-time implementability remains difficult.

This difficulty becomes even more pronounced in the presence of combustion uncertainty. HCCI engines are characterized by strong cycle-to-cycle variations, which may lead to occasional but large deviations in combustion phasing and pressure-related variables. Under such conditions, nominal MPC may yield control actions that are too aggressive with respect to probabilistic constraint satisfaction. Robust MPC (RMPC) mitigates this issue by accounting for the worst-case effect of uncertainty over the prediction horizon (Bemporad & Morari, 2007). Such a strategy is practical when the uncertainty set is small or mainly parametric. However, for highly stochastic combustion systems, rare but large deviations may occur only infrequently, and a purely worst-case treatment can lead to overly conservative behavior or even feasibility problems.

To balance robustness and control performance, Stochastic MPC (SMPC) relaxes hard constraint satisfaction by permitting a small and quantifiable probability of violation (Mesbah et al., 2019). However, a central challenge in SMPC is the treatment of chance constraints. Existing methods often rely either on sampling-based approximations of the underlying probability distribution (Batina, 2004; Blackmore et al., 2010) or on restrictive parametric assumptions such as Gaussian disturbances (Oldewurtel et al., 2008; van Hessem & Bosgra, 2003). Sampling-based approaches may require a large number of scenarios to certify small violation probabilities, which substantially increases computational complexity. Parametric approaches, in contrast, may become inaccurate when the disturbance distribution is non-Gaussian or exhibits complex tail behavior, thereby resulting in either excessive conservatism or underestimated risk.

An effective strategy to handle stochastic uncertainties consists in so-called distributionally robust constraints, which entail constructing an ambiguity set over distributions with specified properties and evaluating the worst-case impact of these distributions on the constraints (Rahimian & Mehrotra, 2019). The two most commonly used types of ambiguity sets are distance-based sets (Gao & Kleywegt, 2023) and moment-based sets (Delage & Ye, 2010; Scarf et al., 1957; Wang et al., 2016). In the context of SMPC, moment-based ambiguity sets are particularly attractive because they of10 admit tractable reformulations (Popescu, 2007; Zymler et al., 2013). However, most existing formulations are constructed only from the first and second moments of the uncertainty. While this leads to computationally tractable optimization problems, it may also yield overly conservative decisions because mean and covariance alone cannot adequately capture more complex distributional features. This limitation is especially relevant for HCCI combustion, where cycle-to-cycle variations may exhibit non-Gaussian characteristics and pronounced dispersion (Chen et al., 2023). Recent studies (Lu et al., 2020; Shang & You, 2018) have shown that nonlinear functionals can be used to incorporate richer statistical information into ambiguity-set construction, thereby enabling a less restrictive characterization of uncertainty. However, the integration of such nonlinear moment information into SMPC for chance-constraint handling has not yet been sufficiently investigated.

In addition to uncertainty representation, real-time implementation remains a second major obstacle. Combustion processes evolve on a fast timescale, leaving only a limited time window for solving the MPC problem online. The trade-off between model fidelity, computational time, and hardware resources hampers the adoption of MPC. Explicit MPC addresses this difficulty through offline policy construction. For linear systems, explicit MPC partitions the state space into critical regions and stores the corresponding optimal control laws in lookup tables (Bemporad et al., 2002). To extend this idea to more complex nonlinear settings, policy approximation methods have been developed (Johansen, 2004). Compared with region-based tabulation, parameterized function approximators such as neural networks offer better scalability and can represent more complex control mappings with lower storage requirements. A common strategy is imitation learning, in which the MPC policy is treated as an expert and a neural network is trained to reproduce the control action (Chen et al., 2021; Reske et al., 2021). However, conventional approaches like behavioral cloning (Torabi et al., 2018) do not incorporate any structural knowledge of MPC, which may impair generalization and safety in out-of-sample regions. In this respect, numerous studies have focused on integrating MPC with learning frameworks (Ghezzi et al., 2023; Gros & Zanon, 2019, 2021), aiming to leverage the strengths of both MPC and learning-based methods.

In summary, although predictive control has shown strong potential for HCCI combustion control, most existing HCCI-oriented MPC methods are still developed in a nominal setting and do not explicitly account for the impact of cycle-to-cycle variability on constraint satisfaction. SMPC provides a natural extension, but its application to HCCI remains challenging because reducing conservatism must be balanced against tractability. Existing tractable approaches typically rely on restrictive distributional assumptions or ambiguity sets based only on first- and second-order moments, which may be overly conservative for highly stochastic HCCI combustion. In addition, the fast combustion timescale places strict limits on online computation. Therefore, there is still a need for an HCCI-oriented control framework that can handle uncertainty less conservatively while remaining tractable and suitable for real-time implementation.

To address the above limitations, this paper proposes a nonlinear moment-based distributionally robust SMPC approach, termed NoM-DiR-SMPC, for HCCI combustion control. Unlike conventional second-moment-based SMPC formulations, the proposed method employs a nonlinear ambiguity set to capture richer statistical information of the underlying uncertainty. It aims to enable a less conservative yet still tractable treatment of chance constraints under HCCI combustion uncertainty. The effectiveness of the proposed approach is validated in both simulation and engine experiments in comparison with a baseline nominal MPC controller.

The contributions of this work are summarized as follows. First, the theoretical contribution lies in the reformulation of chance constraints under general nonlinear statistics into a tractable optimization problem. Recent studies (Lu et al., 2020; Shang & You, 2018) have employed

nonlinear moment-based ambiguity sets using piecewise linear functions to characterize uncertainty in the context of distributionally robust optimization, with the main focus placed on robust optimization itself and the parameterization of the resulting solutions. In contrast, the present work adopts the same uncertainty-characterization principle and extends it to the treatment of chance constraints in stochastic predictive control, which are generally intractable. Specifically, the original chance constraints are approximated via Conditional Value-at-Risk (CVaR) and subsequently converted into tractable LP-based constraint-tightening conditions. This allows the nonlinear moment-based ambiguity-set description to be embedded into a practically solvable SMPC framework for HCCI control.

Second, to enable real-time implementation, we combine the proposed NoM-DiR-SMPC formulation with a learning-based policy approximation strategy following Ghezzi et al. (2023). In this way, the online computational burden associated with stochastic optimization can be largely reduced, making the proposed method more suitable for high-speed engine control.

Third, we deploy the resulting learned SMPC controller on a single-cylinder engine test bench and compare its performance against a learned nominal MPC controller that does not explicitly account for uncertainty. Through this implementation and validation, the present work links richer uncertainty modeling, tractable stochastic optimization, and practical real-time control within a unified framework.

The overall framework of the proposed method is illustrated in Fig. 1. It consists of four main components: uncertainty characterization, tractable chance-constraint reformulation, stochastic MPC design, and learning-based policy approximation for real-time implementation.

The remainder of this work is organized as follows: Section 2 details the methodological framework, beginning with an introduction to the SMPC strategy and the formulation of the equivalent optimization problem. It then proceeds to describe the imitation learning framework that takes into account the MPC structure, namely the objective and constraints. Section 3 introduces the simulative and experimental setup and the evaluation protocol. Section 4 presents the performance results of the proposed SMPC controller and the benchmark MPC controllers in both the simulation environment and the experimental test bench, accompanied by an analysis and discussion of the findings in Section 5. Section 6 concludes the paper by summarizing the potential of the proposed controller, discussing its advantages and limitations, and outlining future improvements.

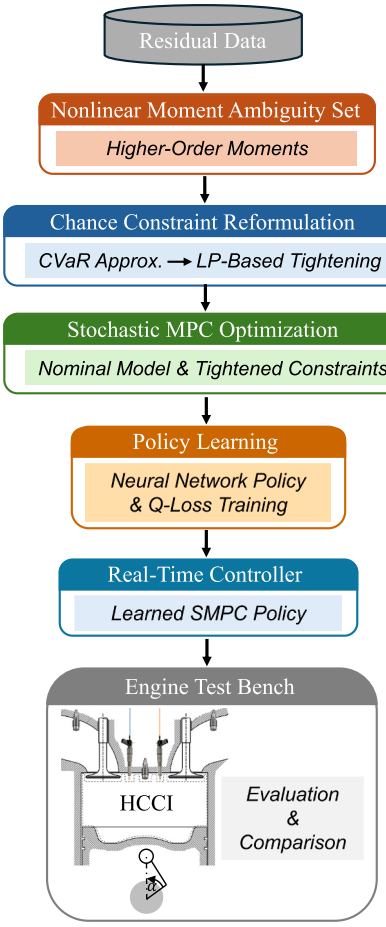


Fig. 1. Overview of the proposed NoM-DiR-SMPC. The approach consists of nonlinear moment-based uncertainty set construction, tractable reformulation of chance constraints via CVaR, stochastic MPC optimization, and learning-based policy approximation for real-time implementation, followed by experimental validation on an HCCI engine test bench.

2. Control strategy

This section consists of two parts. The first outlines the formulation of NoM-DiR-SMPC for handling system uncertainties in HCCI combustion, while the second presents a learning-based method for approximating the control policy to satisfy real-time requirements.

2.1. NoM-DiR-SMPC

We consider a discrete-time, time-invariant nonlinear system subject to additive uncertainties given by

$$X_{k+1} = f(X_k, U_k) + w_k, \tag{1}$$

where $X \in \mathbb{R}^{n_x}$ is the system state, $U \in \mathbb{R}^{n_u}$ is the control input, and $w \in \mathbb{R}^{n_w}$ is the disturbance. Moreover, we assume that state feedback is available for the system in (1). To address the nonlinearities in the MPC optimization objective, we linearize the system model at each time instant k and obtain the linearized model

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \tag{2}$$

where x and u represent the offsets from the linearization point at each time instant. The matrices A_k and B_k denote the Jacobians with respect to the current measured state and the last control input, respectively. Moreover, we assume that the disturbance w_k as a random variable has a compact support described by a convex polytopic set of the form

$$\mathcal{W} := \{w \in \mathbb{R}^{n_w} : Mw \leq m\}, \tag{3}$$

with $M \in \mathbb{R}^{n_m \times n_w}$ and $m \in \mathbb{R}^{n_m}$, and both the system and input are constrained to compact polytopic sets containing the origin given by

$$\mathcal{X} := \{x \in \mathbb{R}^{n_x} : Dx \leq d\}, \tag{4}$$

with $D \in \mathbb{R}^{n_d \times n_x}$ and $d \in \mathbb{R}^{n_d}$, and

$$\mathcal{U} := \{u \in \mathbb{R}^{n_u} : Gu \leq g\}, \tag{5}$$

with $G \in \mathbb{R}^{n_g \times n_u}$ and $g \in \mathbb{R}^{n_g}$. We adjust the feasible state set and input set accordingly with respect to the linearization point.

2.1.1. Problem formulation and policy parameterization

In contrast to standard MPC, SMPC explicitly accounts for the propagation of uncertainties in open-loop predictions, where these uncertainties arise from model mismatch, inherent stochasticity of the system, and noise. We consider the SMPC problem formulation

$$\min_{\substack{u_{0|k},\dots,\\ u_{N-1|k},x_{0|k},\\ \dots,x_{N|k}}} \mathbb{E}\left\{\sum_{i=0}^{N-1}\left(\left\|x_{i|k}\right\|_Q^2+\left\|u_{i|k}\right\|_R^2\right)+\left\|x_{N|k}\right\|_{Q_T}^2\right\}$$
$$\text{s.t. } x_{0|k}=x_k,$$
$$x_{i+1|k}=A_k x_{i|k}+B_k u_{i|k}+w_{i|k}, \quad \forall\, i\in\mathbb{Z}_{[0,N-1]}, \tag{6}$$
$$\mathbb{P}\left(x_{i|k}\in\mathcal{X}\right)\geq 1-\epsilon,\ \forall\, i\in\mathbb{Z}_{[1,N]},$$
$$u_{i|k}\in\mathcal{U}, \quad \forall\, i\in\mathbb{Z}_{[0,N-1]}.$$

In this formulation, the mean objective function covers a prediction horizon N, where Q, R, and Q_T are weighting matrices corresponding to the state, control input, and terminal state, respectively. We assume that the disturbances $w_{i|k}$ are independent and identically distributed. The system matrix propagates the disturbance at each time step to capture the coupling across the prediction horizon. Although its exact realization is unknown during prediction, we can estimate its statistical distribution from historical measurement data.

In real applications, actuators typically have strict output limits, so we treat the control input constraints as hard constraints. In contrast, we formulate state constraints as chance constraints, with $\epsilon\in[0,1)$ denoting the permissible probability of constraint violation.

First, we consider the propagation of uncertainty in the prediction. By exploiting a linear mapping in (2), we gain notable advantages, as the evolution of the nominal state is decoupled from the evolution of the uncertainty. If we parameterize the control input as

$$u_{i|k}=K_k\left(x_{i|k}-z_{i|k}\right)+c_{i|k}, \tag{7}$$

where z represents the nominal state with $z_{0|k}=x_k$, K_k is the feedback gain, and $c_{i|k}$ is the feedforward term. In this decomposition, $x_{i|k}$ represents the stochastic predicted state affected by uncertainty, whereas $z_{i|k}$ corresponds to the disturbance-free system evolution. Then, the nominal state dynamics are

$$z_{i+1|k}=A_k z_{i|k}+B_k c_{i|k}, \tag{8}$$

and the error dynamics are

$$e_{i+1|k}=\left(A_k+B_kK_k\right)e_{i|k}+w_{i|k}. \tag{9}$$

Although the parameterization in (7) enhances computational efficiency, it also leads to a suboptimal solution of the problem in (6). Here, we obtain the feedback gain K_k in (7) by solving an infinite-horizon algebraic Riccati equation, and then the weighting Q_T is the solution to the Lyapunov equation

$$\left(A_k+B_kK_k\right)^\top Q_T\left(A_k+B_kK_k\right)-Q_T=-Q-K_K^TRK_k. \tag{10}$$

Similarly, one could employ H-infinity or game theory approaches as in Lall and Glover (1994), Magni et al. (2001) to derive the corresponding feedback gain.

At each control instant, we obtain a distinct linear system representation, and therefore, the corresponding feedback gain is also time-varying.

2.1.2. Definition and reformulation of the ambiguity set

A key challenge in SMPC lies in addressing the chance constraints in (6). Given historical disturbance data, one can estimate their distribution by evaluating a multivariate convolution integral numerically to see how errors violate the constraints, thus informing the margin by which the original constraints must be tightened (Kouvaritakis et al., 2010). Consequently, one can transform the soft constraints into hard constraints on the nominal state. Although this approach can yield minimally conservative constraints, it becomes impractical when the feedback gain K_k varies over time due to repeated linearization, since each step would require another numerical estimation.

A commonly used alternative approach considers a set of distributions, referred to as an ambiguity set, where all distributions in the set satisfy specific properties, and evaluates the worst case within the set. For instance, one can form a set of distributions with the same first- and second-order moments (Calafiore & Ghaoui, 2006). This strategy, frequently employed in stochastic MPC, is often excessively conservative (Kouvaritakis et al., 2010; Mesbah et al., 2014). However, incorporating additional information into the set, thereby narrowing the range of possible distributions, can help reduce this conservativeness. The following part explicitly addresses techniques for achieving such a reduction.

We consider the space $\mathbb{R}^n$, and $\mathcal{P}(\mathbb{R}^n)$ denotes the set of Borel probability measures. We further introduce a family of general ambiguity sets in Wiesemann et al. (2014), formulated as

$$\mathcal{M}:=\left\{\mathbb{P}_w\in\mathcal{P}(\mathbb{R}^{n_w}):\ \begin{array}{l}\mathbb{P}_w\{w\in\mathcal{W}\}=1,\\ \mathbb{E}_{\mathbb{P}}[\,g_i(w)\,]\ \leq\ \gamma_i,\ i\in\mathbb{Z}_{[1,I]}\end{array}\right\}. \tag{11}$$

The first constraint in (11) ensures that all valid distributions are supported on the given compact set $\mathcal{W}$. In the second constraint, we employ different functions g_i with a total of I to capture more general moment information about the distributions, with each moment bounded by the value γ_i.

We use equivalent sets with an auxiliary variable $\phi\in\mathbb{R}^I$ to transform the chance constraints into tractable form (Shang & You, 2018):

$$\tilde{\mathcal{M}}:=\left\{\mathbb{P}_{w,\phi}\in\mathcal{P}(\mathbb{R}^{n_w+I}):\ \begin{array}{l}\mathbb{P}_{w,\phi}\{(w,\phi)\in\tilde{\mathcal{W}}\}=1,\\ \mathbb{E}_{\mathbb{P}}[\,\phi\,]\ \leq\ \boldsymbol{\gamma}\end{array}\right\}, \tag{12}$$

where we define the lifted support set as

$$\tilde{\mathcal{W}}:=\left\{(w,\phi):\ \begin{array}{l}w\in\mathcal{W},\\ g_i(w)\leq\phi_i,\ \forall i\in\mathbb{Z}_{[1:I]}\end{array}\right\}. \tag{13}$$

Having established the general description of the ambiguity set for disturbances in (12), we now revisit the chance constraint by formulating it in terms of the ambiguity set. Our goal is to recast the original chance constraints in (6) into a tractable convex formulation.

In general, only few cases of chance constraints admit tractable convex representations (Zymler et al., 2013). To tackle this, we introduce a CVaR constraint as in Rockafellar et al. (2000), which provides a convex conservative approximation for the chance constraints. By definition, given a measurable function $L:\mathbb{R}^{n_w}\to\mathbb{R}$, a probability measure $\mathbb{P}\in\mathcal{P}(\mathbb{R}^{n_w})$, and an allowable violation probability ϵ, one can express the CVaR as

$$\mathrm{CVaR}_\epsilon\left(L(w)\right)=\inf_{\beta\in\mathbb{R}}\left\{\beta\ +\ \frac{1}{\epsilon}\mathbb{E}_{\mathbb{P}}\left[\left(L(w)-\beta\right)_+\right]\right\}, \tag{14}$$

where the function $(\cdot)_+:=\max(\cdot,0)$.

For the individual chance constraints in (6), we consider the worst case over the given ambiguity set in (12), written as

$$\inf_{\mathbb{P}\in\tilde{\mathcal{M}}^i}\ \mathbb{P}\left(Dz_{i|k}+De_{i|k}(w_{0:i-1|k})-d\leq 0\right)\geq 1-\epsilon \tag{15}$$

at each prediction step i, where $\tilde{\mathcal{M}}^i$ denotes the i-Cartesian product of $\tilde{\mathcal{M}}$. In the present work, the independent and identically distributed (i.i.d.) assumption is adopted for the residual disturbance term over the prediction horizon as a tractable approximation for stochastic MPC reformulation.

From a physical standpoint, CVaR quantifies the expected value of losses beyond the $(1-\epsilon)$-quantile of the distribution, and this expectation bound sufficiently exceeds the feasibility bound imposed by the chance constraints. Hence, we can always establish the following conclusion:

$$\begin{aligned}&\sup_{\mathbb{P}\in\tilde{\mathcal{M}}^i}\mathrm{CVaR}_\epsilon\left(Dz_{i|k}+De_{i|k}(w_{0:i-1|k})-d\right)\leq 0\Longrightarrow\\ &\inf_{\mathbb{P}\in\tilde{\mathcal{M}}^i}\ \mathbb{P}\left(Dz_{i|k}+De_{i|k}(w_{0:i-1|k})-d\leq 0\right)\geq 1-\epsilon.\end{aligned} \tag{16}$$

CVaR furnishes a conservative estimate for chance constraints. Under certain assumptions regarding both the ambiguity set and the loss function, the converse implication also holds (Zymler et al., 2013).

Next, we employ CVaR as an approximation for the chance constraints and derive its tractable representation under the given ambiguity set in (12). We begin by examining the following Lemma.

Lemma 1. *(Shang & You, 2018) Consider the ambiguity set described in (12) and a piecewise linear function l. The Lagrange dual problem of the maximization problem $\sup_{\mathbb{P}\in\tilde{\mathcal{M}}} \mathbb{E}_{\mathbb{P}}[l(w)]$ is given by*

$$\begin{aligned} \min_{\eta,\mu} \quad & \eta + \gamma^\top \mu \\ s.t. \quad & \mu \geq 0, \\ & \eta + \phi^\top \mu \geq l(w), \ \forall (w,\phi) \in \tilde{\mathcal{W}}, \end{aligned} \tag{17}$$

where η and μ are associating Lagrange multipliers, and the strong duality holds. A rigorous justification of the strong duality is provided in Appendix A.

Based on Lemma 1, we transform the problem posed with infinite-dimensional optimization variables into one involving infinitely many constraints. We arrive at the following Proposition.

Proposition 1. *Consider the CVaR condition in (16). When we express the ambiguity set as in (12), the condition is equivalent to*

$$\begin{aligned} & \beta\epsilon + \eta + \gamma^\top \mu \leq 0, \\ & -\mu \leq 0, \\ & -\eta - \phi^\top \mu \leq 0, \ \forall (w,\phi) \in \tilde{\mathcal{W}}^i, \\ & Dz_{i|k} + De_{i|k}(w_{0:i-1|k}) - d - \beta - \eta - \phi^\top \mu \leq 0, \\ & \forall (w,\phi) \in \tilde{\mathcal{W}}^i. \end{aligned} \tag{18}$$

Proof. We reformulate the CVaR condition in (16) as

$$\begin{aligned} & \sup_{\mathbb{P}\in\mathcal{M}^i} \inf_{\beta\in\mathbb{R}} \left\{ \beta + \frac{1}{\epsilon} \mathbb{E}_{\mathbb{P}} \left[\left(Dx_{i|k} - d - \beta \right)_+ \right] \right\} \leq 0 \\ \overset{(a)}{\iff} & \\ & \inf_{\beta\in\mathbb{R}} \sup_{\mathbb{P}\in\mathcal{M}^i} \left\{ \beta + \frac{1}{\epsilon} \mathbb{E}_{\mathbb{P}} \left[\left(Dx_{i|k} - d - \beta \right)_+ \right] \right\} \leq 0 \\ \iff & \\ & \inf_{\beta\in\mathbb{R}} \left\{ \beta\epsilon + \sup_{\mathbb{P}\in\mathcal{M}^{i+1}} \mathbb{E}_{\mathbb{P}} \left[\left(Dx_{i|k} - d - \beta \right)_+ \right] \right\} \leq 0 \\ \overset{(b)}{\iff} & \begin{cases} \inf_{\beta,\eta,\mu} \ \beta\epsilon + \eta + \gamma^\top \mu \leq 0, \\ \mu \geq 0, \\ \eta + \phi^\top \mu \geq \left(Dx_{i|k} - d - \beta \right)_+, \ \forall (w,\phi) \in \tilde{\mathcal{W}}^i, \end{cases} \\ \overset{(c)}{\iff} & \begin{cases} \beta\epsilon + \eta + \gamma^\top \mu \leq 0, \\ \mu \geq 0, \\ \eta + \phi^\top \mu \geq 0, \ \forall (w,\phi) \in \tilde{\mathcal{W}}^i, \\ \eta + \phi^\top \mu \geq Dx_{i|k} - d - \beta, \ \forall (w,\phi) \in \tilde{\mathcal{W}}^i. \end{cases} \end{aligned} \tag{19}$$

In equivalence (a), the interchange of the max and min operators guaranteed by strong duality follows from the reference Shapiro and Kleywegt (2002), and we obtain equivalence (b) directly by applying Lemma 1. A rigorous justification of the strong duality is provided in Appendix A. Furthermore, in equivalence (c), the min-problem in the inequality is derived from checking the existence of a feasible solution. □

Remark 1. The CVaR-based reformulation adopted in this work serves as a tractable sufficient condition for the original chance constraint, and therefore introduces a degree of conservatism. In general, the gap between the CVaR constraint and the corresponding chance constraint depends on the underlying uncertainty distribution, particularly on the behavior of its tail.

For distributions with heavy tails, the CVaR constraint may lead to a noticeable level of conservatism, since it controls not only the violation probability but also the magnitude of violations in the tail region. In contrast, for bounded or light-tailed distributions, where extreme violations are limited and the tail probability mass is relatively small, the gap introduced by the CVaR approximation is typically moderate.

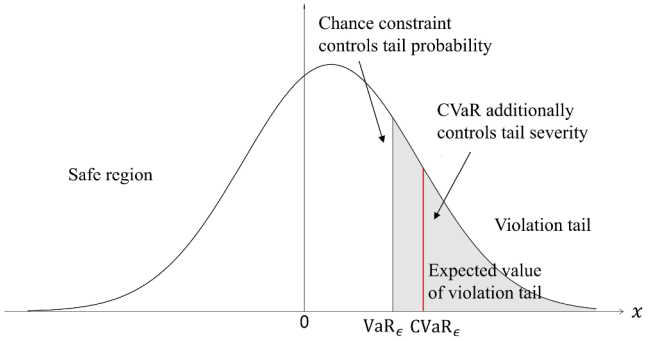


Fig. 2. Illustration of the relationship between the chance constraint and the CVaR approximation. The chance constraint limits the probability of the violation tail, while the CVaR constraint additionally limits the average severity of tail violations.

In the considered HCCI application, the uncertainty is characterized by residual data with bounded support and no evident heavy-tail behavior in the available dataset. Therefore, the CVaR-based reformulation provides a reasonable approximation of the original chance constraint while maintaining computational tractability.

From an interpretation perspective, the CVaR constraint can be viewed as imposing a bound on the expected severity of violations beyond a given probability level, as illustrated in Fig. 2. This means that, in addition to limiting how often constraint violations occur, the formulation also restricts how severe those violations can be, which explains the conservative nature of the approximation.

2.1.3. Chance constraint handling

Before we continue to explore the explicit formulation of the ambiguity set, we first need to specify the support set $\mathcal{W}$. For simplicity, we impose constraints along the principal directions u_r identified from the data via Principal Component Analysis, capturing the most significant modes of variability in the uncertainty. Using these directions, we derive the outer approximation of the support set

$$\mathcal{W} = \left\{ w : w_{\min,r} \leq u_r^\top w \leq w_{\max,r}, \ \forall r \in \mathbb{Z}_{[1,n_w]} \right\}, \tag{20}$$

where the thresholds $w_{\min,r}$ and $w_{\max,r}$ are the minimal and maximal projected values along the direction u_r, respectively.

To characterize the moments for (13), we adopt the same choice of the function g proposed in Shang and You (2018), namely a piecewise linear function of the form

$$g_{rl}(w) = \max(p_r^\top w - q_{rl}, 0), \ r \in \mathbb{Z}_{[1,n_w]}, \ l \in \mathbb{Z}_{[1,I']}. \tag{21}$$

In this interpretation, we view this form as the first-order deviation of disturbance in the direction p_r with the threshold q_{rl} and specify the number I' later. By employing multiple piecewise linear functions in (21), we can capture higher-order moments of the distribution, thereby incorporating richer information about the distribution and also enabling a more tractable reformulation of the chance constraints.

According to Shang and You (2018), the reasonable choice of direction p and threshold value q is informed by historical disturbance data. Specifically, assume we have N_w realizations of w drawn from measurement data, from which we can estimate its mean vector $\hat{\mu}$ and covariance matrix $\hat{\Sigma}$. A PCA decomposition on $\hat{\Sigma}$ yields $\hat{\Sigma} = U\Lambda U^\top$, where U is an orthogonal matrix, and Λ is a diagonal matrix. We select p_r to be the rth principal direction u_r, with u_r denoting the rth column of U. In total, there are n_w directions. For each direction p_r, we choose $I' := 2K+1$ threshold values, equally spaced by the distance δ_r. Here, we define δ_r as the rth element on the diagonal of Λ. Hence, the $2K+1$ threshold

values along the direction p_r take the form

$$p_r^\top \hat{\mu}, p_r^\top \hat{\mu} + \delta_r, p_r^\top \hat{\mu} - \delta_r, \dots, p_r^\top \hat{\mu} + K\delta_r, p_r^\top \hat{\mu} - K\delta_r. \tag{22}$$

Up to this point, we can express the lifted support set of the disturbance w as

$$\tilde{\mathcal{W}} = \left\{ (w,\phi) : \begin{array}{c} U^\top w - w_{\max} \le 0, \\ -U^\top w + w_{\min} \le 0, \\ -\phi_{rl} \le 0, \\ p_r^\top w - q_{rl} - \phi_{rl} \le 0, \\ \forall r \in \mathbb{Z}_{[1,n_w]}, \forall l \in \mathbb{Z}_{[1,2K+1]} \end{array} \right\}, \tag{23}$$

which we can rewrite in a compact matrix form

$$\tilde{\mathcal{W}} = \{(w,\varphi) : H_1 w + H_2 \varphi - h \le 0\}, \tag{24}$$

with

$$H_1 = \begin{bmatrix} U^\top \\ -U^\top \\ \mathbf{0}_{(2K+1)n_w \times n_w} \\ \mathbf{p} \end{bmatrix}, H_2 = \begin{bmatrix} \mathbf{0}_{n_w \times (2K+1)n_w} \\ \mathbf{0}_{n_w \times (2K+1)n_w} \\ -\mathbf{I}_{(2K+1)n_w} \\ -\mathbf{I}_{(2K+1)n_w} \end{bmatrix}, \quad h = \begin{bmatrix} w_{\max} \\ -w_{\min} \\ \mathbf{0}_{(2K+1)n_w \times 1} \\ \mathbf{q} \end{bmatrix}, \tag{25}$$

where $\mathbf{p} \in \mathbb{R}^{(2K+1)n_w \times n_w}$, and $\mathbf{q} \in \mathbb{R}^{(2K+1)n_w \times 1}$.

Given the decoupling between the nominal system and the error dynamics, we determine the minimal tightening parameter α_{ij} for each step i and each constraint j. This ensures that, if the nominal trajectory satisfies the tightened constraint, the original system satisfies the associated chance constraint. Consequently, we write the tightened constraint on the nominal system as

$$D_j z_{i|k} - d_j + \alpha_{ij} \le 0, \ \forall i \in \mathbb{Z}_{[1,N]}, j \in \mathbb{Z}_{[1,n_d]}. \tag{26}$$

When substituting (26) into (18) and seeking the smallest possible α_{ij} while ensuring that the constraints in (18) are satisfied so as to guarantee the chance constraints, we obtain the following linear program:

$$\begin{aligned} \min_{\eta,\mu,\beta,\alpha_{ij}} \quad & \alpha_{ij} \\ \text{s.t.} \quad & \beta\epsilon + \eta + \gamma^\top \mu \le 0, \\ & -\mu \le 0, \\ & -\eta - \phi^\top \mu \le 0, \ \forall (w,\phi) \in \tilde{\mathcal{W}}^i, \\ & D_j e_{i|t} - \alpha_{ij} - \beta - \eta - \phi^\top \mu \le 0, \ \forall (w,\phi) \in \tilde{\mathcal{W}}^i. \end{aligned} \tag{27}$$

At this stage, one problem remains unresolved, namely that the optimization problem in (27) involves infinitely many constraints. Therefore, we further transform it into a tractable form stated in the following Proposition.

Proposition 2. *Considering the lifted support set in (24), we reformulate the optimization problem in (27) as*

$$\begin{aligned} \min_{\eta,\mu,\beta,\alpha_{ij},\xi_1,\xi_2} \quad & \alpha_{ij} \\ s.t. \quad & \beta\epsilon + \eta + \gamma^\top \mu \le 0, \\ & -\eta + \mathbf{h}_i^\top \xi_1 \le 0, \\ & -\alpha_{ij} - \beta - \eta + \mathbf{h}_i^\top \xi_2 \le 0, \\ & \mathbf{H}_{1i}^\top \xi_1 = 0, \ \mathbf{H}_{2i}^\top \xi_1 + \mu = 0, \\ & \mathbf{H}_{1i}^\top \xi_2 - \mathbf{T}_i^\top D_j^\top = 0, \ \mathbf{H}_{2i}^\top \xi_2 + \mu = 0, \\ & \mu \ge 0, \ \xi_1 \ge 0, \ \xi_2 \ge 0, \alpha_{ij} \ge 0, \end{aligned} \tag{28}$$

where $\mathbf{T}_i := \left[A_K^{i-1}, \dots, A_K, I\right]$, $\mathbf{H}_{1i} := \mathbf{I}_i \otimes H_1$, $\mathbf{H}_{2i} := \mathbf{I}_i \otimes H_2$, *and* $\mathbf{h}_i := \mathbf{1}_{i\times 1} \otimes h$.

Proof. First, let us consider the error dynamics in (9). At the ith prediction step, we express the error as

$$e_{i|k} = \underbrace{\left[A_K^{i-1}, \dots, A_K, I\right]}_{\mathbf{T}_i} \underbrace{\left[w_{0|k}, \dots, w_{i-2|k}, w_{i-1|k}\right]^\top}_{\mathbf{w}_{ik}}. \tag{29}$$

Furthermore, we also define $\Phi_i = \left[\phi_0, \dots, \phi_{i-2}, \phi_{i-1}\right]^\top$ and write the lifted set $\tilde{\mathcal{W}}^i$ as

$$\tilde{\mathcal{W}}^i = \left\{(\mathbf{w}_{ik}, \Phi_i) : \mathbf{H}_{1i}\mathbf{w}_{ik} + \mathbf{H}_{2i}\Phi_i - \mathbf{h}_i \le 0\right\}, \tag{30}$$

with $\mathbf{H}_{1i} := \mathbf{I}_i \otimes H_1$, $\mathbf{H}_{2i} := \mathbf{I}_i \otimes H_2$, and $\mathbf{h}_i := \mathbf{1}_{i\times 1} \otimes h$. In the problem in (27), we reformulate the final two constraints as

$$\begin{aligned} &\max_{(\mathbf{w}_i,\Phi_i)\in\tilde{\mathcal{W}}^i} -\eta - \Phi_i^\top \mu \le 0, \\ &\max_{(\mathbf{w}_i,\Phi_i)\in\tilde{\mathcal{W}}^i} D_j \mathbf{T}_i \mathbf{w}_i - \alpha_{ij} - \beta - \eta - \Phi_i^\top \mu \le 0. \end{aligned} \tag{31}$$

It is worth noting that we deal with two robust constraints with identical $\mathbf{w}_i$ and Φ_i. As for the worst-case for each constraint, we must address them separately. For the first constraint, its dual problem is

$$\begin{aligned} &-\eta + \mathbf{h}_i^\top \xi_1 \le 0, \quad \xi_1 \ge 0, \\ &\mathbf{H}_{1i}^\top \xi_1 = 0, \quad \mathbf{H}_{2i}^\top \xi_1 + \mu = 0, \end{aligned} \tag{32}$$

where ξ_1 is the dual variable. Then, we can reformulate the second constraint as

$$\begin{aligned} &-\alpha_{ij} - \beta - \eta + \mathbf{h}_i^\top \xi_2 \le 0, \quad \xi_2 \ge 0, \\ &\mathbf{H}_{1i}^\top \xi_2 - \mathbf{T}_i^\top D_j^\top = 0, \quad \mathbf{H}_{2i}^\top \xi_2 + \mu = 0, \end{aligned} \tag{33}$$

with the dual variable ξ_2. □

Similarly, for the hard constraints on the control input, we consider the robust case, i.e., we solve the optimization problem

$$\begin{aligned} \min_{\kappa_{ij}} \quad & \kappa_{ij} \\ \text{s.t.} \quad & \max_{\mathbf{w}_{ik}\in\mathcal{W}^i} G_j K_k \mathbf{T}_i \mathbf{w}_{ik} - \kappa_{ij} \le 0, \\ & \kappa_{ij} \ge 0, \end{aligned} \tag{34}$$

with the tightened margin κ, which is also equivalent to

$$\begin{aligned} \min_{\kappa_{ij},\xi_3} \quad & \kappa_{ij} \\ \text{s.t.} \quad & -\kappa_{ij} + \hat{\mathbf{h}}^\top \xi_3 \le 0, \\ & \xi_3^\top \hat{\mathbf{H}} = G_j K_k \mathbf{T}_i, \\ & \kappa_{ij} \ge 0, \ \xi_3 \ge 0, \end{aligned} \tag{35}$$

where ξ_3 is the dual variable. Here, we define $\mathcal{W}^i = \{w : \hat{H}w - \hat{h} \le 0\}$ as the compact linear form of the formulation in (20), and $\hat{\mathbf{H}} := \mathbf{I}_i \otimes \hat{H}$ and $\hat{\mathbf{h}} := \mathbf{1}_{i\times 1} \otimes \hat{h}$. Notably, for the optimization in (34), if the disturbance w has a small dimension and the prediction horizon is also short, one can directly assess the impact of the vertices of the uncertainty set on the constraints, which could be more efficient.

2.1.4. Implementation of NoM-DiR-SMPC

Once we have determined the tightened margins α_{ij} and κ_{ij}, we can rewrite the original optimization problem in (6) in terms of the nominal system as

$$\begin{aligned} \min_{\substack{c_{0|k},\dots,\\ c_{N-1|k},z_{0|k},\\ \dots,z_{N|k}}} \quad & \sum_{i=0}^{N-1} \left(\left\|z_{i|k}\right\|_Q^2 + \left\|c_{i|k}\right\|_R^2 \right) + \left\|z_{N|k}\right\|_{Q_T}^2 \\ \text{s.t.} \quad & z_{0|k} = x_k, \\ & z_{i+1|k} = A_k z_{i|k} + B_k c_{i|k}, \ \forall i \in \mathbb{Z}_{[0,N-1]}, \\ & D z_{i|k} - d + \alpha_i \le 0, \ \forall i \in \mathbb{Z}_{[1,N]}, \\ & G c_{i|k} - g + \kappa_i \le 0, \quad \forall i \in \mathbb{Z}_{[0,N-1]}, \end{aligned} \tag{36}$$

with $\kappa_0 = 0$. In this case, the resulting formulation reduces to a quadratic program, greatly simplifying the computational process.

Finally, one additional point merits clarification, where when using a fixed feedback gain in (7), the original expected objective function in (6) approximately reduces to the nominal objective function in (36). The reason for this becomes apparent when we only account for the second-order moment of the distribution. Specifically, at each step, we can express the error covariance with respect to the dynamics in (9) as

$$\Sigma_e^{i+1} = (A_k + B_k K_k)\Sigma_e^i (A_k + B_k K_k)^\top + \Sigma_w, \tag{37}$$

where $\Sigma_e^0 = 0_{n_x}$ at the time instant $i = 0$, and Σ_w is the covariance of the disturbance. Additionally, the state covariances Σ_x are equal to the error covariances, and when the feedback gain K_k is determined, all of them are constant. Then, we can adopt the expected objective function as

$$\begin{aligned}
&\mathbb{E}\left\{ \sum_{i=0}^{N-1} \left(\left\|x_{i|k}\right\|_Q^2 + \left\|u_{i|k}\right\|_R^2 \right) + \left\|x_{N|k}\right\|_{Q_T}^2 \right\} \\
&= \sum_{i=0}^{N-1} \left(\left\|z_{i|k}\right\|_Q^2 + \mathrm{tr}\left(\Sigma_e^i Q\right) + \left\|c_{i|k}\right\|_R^2 + \mathrm{tr}\left(K_k \Sigma_e^i K_k^\top R\right) \right) \\
&\quad + \left\|z_{N|k}\right\|_{Q_T}^2 + \mathrm{tr}\left(\Sigma_e^N Q_T\right).
\end{aligned} \tag{38}$$

By retaining only the terms that depend on the optimization variables, the original objective function degenerates to the nominal objective.

Finally, we arrive at the following SMPC algorithm:

Algorithm 1 Stochastic model predictive control algorithm.

Require: Initial state x_0, system dynamics in (1), state constraints $\mathcal{X}$, input constraint $\mathcal{U}$, lifted support set $\tilde{\mathcal{W}}$, objective function, and prediction horizon N.

1: **for** time step $k = 0, 1, 2 \ldots$ **do**
2: Measure the current state x_k
3: Linearize the nonlinear system to get the linear system in (2)
4: Determine the state and input constraints in terms of the linearized system
5: Determine the feedback gain K_k in (7) and terminal cost weight Q_T in (6)
6: Obtain the state and input safety margin α_{ij} and κ_{ij} by solving the linear programs in (28) and (35)
7: Solve the finite-horizon optimal control problem in (36)
8: Apply the first optimal control input $c_{0|k}^*$ to the plant
9: **end for**

Remark 2. In Step 6, the linear program solved at each time instant exhibits an identical structure, with variations occurring only in specific parameters such as the matrices T_i and D_j. Given that the constraint structure and dimensionality remain fixed, we can employ a warm-start strategy to initialize the solver with the previous solution. This approach reduces the number of iterations required for convergence and enhances overall computational efficiency.

Remark 3. In the proposed formulation, recursive feasibility is ensured in practice through the use of slack variables shown in Section 3.4, which soften the state constraints and prevent infeasibility caused by model mismatch or aggressive constraint tightening.

Regarding stability, the proposed controller is based on a locally linearized prediction model and implemented in a receding-horizon fashion. Therefore, stability can only be expected in a local and practical sense. A corresponding theoretical analysis is provided in Appendix B, where it is shown that, under the linearized system and standard terminal cost assumptions, the closed-loop system is practically stable in the sense of bounded long-term average quadratic cost.

Remark 4. The proposed tightening should be interpreted as a local receding-horizon approximation rather than a global certificate for the original nonlinear system. In the HCCI application, except for transient operating phases, the closed-loop system evolves predominantly in a neighborhood of a local operating point. Therefore, the locally linearized model provides a meaningful approximation over the corresponding prediction horizon, which is consistent with the standard operating regime of linearization-based predictive control.

At each sampling instant, the nonlinear dynamics are re-linearized around the current operating point, and the corresponding feedback gain and tightening margins are recomputed for the current prediction horizon. In this way, the tightening is always constructed from the most recent local model, rather than being propagated indefinitely from a previous linearization.

It should also be noted that the present formulation does not explicitly model the error introduced by repeated local linearization. Accordingly, the proposed tightening is not interpreted as a strict robustness certificate against all linearization mismatch. Instead, its validity relies on the practical operating condition that the closed-loop trajectory remains in a local neighborhood of the current operating point over one prediction horizon, so that the linearized model remains sufficiently representative for constraint tightening and prediction.

2.2. Imitation learning

To enable real-time deployment of the NoM-DiR-SMPC algorithm, we replace online optimization, which defines an implicit mapping from the state x_k to the control input $u_0^\star$, with an explicit policy representation. At each time instant k, instead of solving the optimization problem in (36), we approximate the mapping $x_k \to u_0^\star(x_k)$ using a parameterized model $\pi_\theta : \mathbb{R}^{n_x} \to \mathbb{R}^{n_u}$. This model directly generates the optimal control input $u_0^\star$ as a function of the current state. We choose a neural network to represent π_θ due to its universal approximation capabilities and high flexibility. The parameter vector θ is trainable, and its dimensionality depends on the network configuration. This parameterization reduces the online computational burden and makes the controller efficient and simple to deploy in high-frequency HCCI control tasks.

To train the policy model, we adopt an imitation learning framework. Its core is a loss function that quantifies how well the policy mimics the optimal MPC behavior. In general, we can formulate the objective as

$$L(\theta) := \mathbb{E}_{(x,y)\sim D}\left[\ell\left(y, \pi_\theta(x)\right)\right], \tag{39}$$

where x and y denote the input-output samples from an expert, and D denotes the sample distribution. Our goal is to find the optimal parameter vector θ to minimize the objective in (39).

A common strategy is behavioral cloning, which minimizes the squared error between expert and predicted actions

$$l(\theta) := \left\| u_i^{\text{expert}} - \pi_\theta(x_i) \right\|^2, \tag{40}$$

where u_i^{expert} represents the expert control input corresponding to the state x_i.

Typical imitation learning methods prioritize data efficiency, as they distill control policies directly from experimental datasets (Hussein et al., 2017). However, when applied to MPC, these methods often neglect its intrinsic objective and constraint structure, which raises concerns regarding safety and out-of-sample performance. To overcome these limitations, we adopt a learning framework specifically tailored for MPC introduced in Ghezzi et al. (2023), which integrates the original MPC formulation into the training process via an optimization-based loss termed as Q-loss. The specific form of this loss ℓ_{MPC} is

$$
\begin{aligned}
&\ell_{\mathrm{MPC}}\big(x_k, \pi_\theta(x_k)\big) := \\
&\min_{\substack{c_{0|k},\dots,\\ c_{N-1|k}, z_{0|k},\\ \dots, z_{N|k}}} \sum_{i=0}^{N-1} \left(\|z_{i|k}\|_Q^2 + \|c_{i|k}\|_R^2 \right) + \|z_{N|k}\|_{Q_T}^2 \\
&\quad \text{s.t. } z_{0|k} = x_k, \\
&\qquad c_{0|k} = \pi_\theta(x_k), \\
&\qquad z_{i+1|k} = A_k z_{i|k} + B_k c_{i|k},\ \forall\, i \in \mathbb{Z}_{[0,N-1]}, \\
&\qquad D z_{i|k} - d + \alpha_i \le 0,\ \forall\, i \in \mathbb{Z}_{[1,N]}, \\
&\qquad G c_{i|k} - g + \kappa_i \le 0, \quad \forall\, i \in \mathbb{Z}_{[0,N-1]}.
\end{aligned} \tag{41}
$$

In other words, to compute this loss, we need to solve the discrete optimal control problem in (41) for each sample scenario, where the policy model determines the first control input. In practice, we add slack variables to the constraints during training to ensure problem feasibility, especially during the initial stage of training. For this loss function in (41), we define it as an implicit function of the state x_k and the policy $\pi_\theta(\cdot)$ characterized by the optimization process. If the policy function closely approximates the optimal control input of the original problem in (36), the loss function value approaches its global minimum. Thus, our training goal is to parametrize the θ so as to minimize the objective function

$$
L_{\mathrm{MPC}}(\theta) := \mathbb{E}_{x\sim D}\big[\ell_{\mathrm{MPC}}\big(x, \pi_\theta(x)\big)\big]. \tag{42}
$$

During the learning process, we update the parameters θ iteratively and express the gradient of the objective function with respect to θ as

$$
\nabla_\theta L_{\mathrm{MPC}}(\theta) = \mathbb{E}_{x\sim D}\big[\nabla_\theta \pi_\theta(x)\ \nabla_\pi \ell_{\mathrm{MPC}}\big(x, \pi_\theta(x)\big)\big], \tag{43}
$$

where the first gradient term $\nabla_\theta \pi_\theta(x)$ comes from the backpropagation of the parameterized model π_θ, and the second term $\nabla_\pi \ell_{\mathrm{MPC}}\big(x, \pi_\theta(x)\big)$ is equal to the Lagrangian multiplier λ_π with respect to the initial input constraint $c_{0|k} - \pi_\theta(x_k) = 0$ in (41). We can derive it from the sensitivity analysis of the perturbed problem in (41) or by checking the first-order necessary condition of optimality (Bertsekas, 2012). In practice, approximation errors are inevitable.

For linear MPC problems, the optimal control law is a piecewise affine function of the current state x_k. In contrast, for the linearized nonlinear MPC problem, the control law generally lacks a closed-form structure and exhibits higher complexity. The effectiveness of the learned policy in approximating such control laws depends critically on the expressive capacity of the chosen neural network architecture.

The result of this training is an explicit policy π_θ, which requires only a forward pass through the neural network to compute the control action from the current state in HCCI applications. This approach completely eliminates the need for solving optimization problems online and enables fast and reliable deployment on embedded hardware.

3. Setup and protocol

To evaluate the performance and real-time feasibility of the proposed NoM-DiR-SMPC control framework, we conduct both simulation and experimental studies. This section describes the system modeling, experimental setup, and evaluation protocol used in both simulation and real-world validation to ensure a comprehensive performance assessment.

3.1. System modeling and control objectives

In this work, we adopt the neural network model framework for the MPC prediction model described in Chen et al. (2023), and we employ the zero-order model from Nuss et al. (2019) as our simulation model. The prediction model is formulated in a cycle-to-cycle manner, mapping the current combustion state and control input to the combustion state of the subsequent engine cycle. For online optimization, this nonlinear predictor is locally linearized around the current operating point, yielding the model used in the MPC formulation.

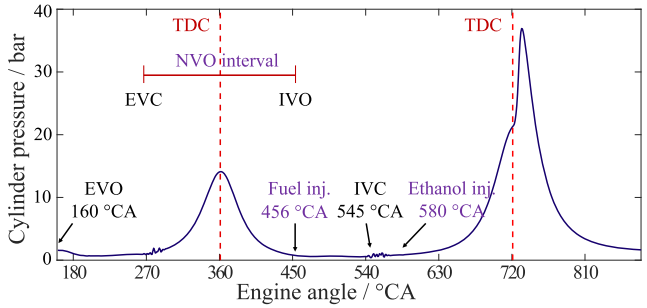


Fig. 3. Representative in-cylinder pressure trace of the HCCI operation in crank-angle coordinates at CA50 = 6 °CA aTDC and IMEP = 2.9 bar. The valve events (EVO, EVC, IVO, and IVC), the start timings of gasoline and ethanol injection, and the NVO interval are indicated.

Regarding the prediction model, we define the system state as $x = \big[x_{\mathrm{CA50}}, x_{\mathrm{imep}}, x_{\mathrm{DPmax}}\big]^\top$, where CA50 represents the crank angle at which 50 % of the cumulative heat release is achieved during combustion, serving as an indicator of combustion efficiency, and its standard deviation serves as an indicator of combustion stability. Indicated mean effective pressure (IMEP) is correlated with the output torque, and the driver input determines its actual reference. Additionally, DPmax, which represents the maximum in-cylinder pressure gradient, is closely related to mechanical stress and noise of the engine, and we need to ensure it remains below a specific upper bound. These physical quantities can be computed online from the high-resolution in-cylinder pressure trace with high accuracy and exhibit strong cycle-to-cycle correlations. Therefore, for the considered experimental setup, cycle-resolved estimates of the control-oriented states are available to the controller. Based on these considerations, they are selected as the states of the prediction model.

The HCCI operation considered in this study relies on a dual-fuel strategy together with NVO. The system input is $u = \big[u_{\mathrm{nvo}}, u_{\mathrm{doi,f}}, u_{\mathrm{doi,e}}\big]^\top$, where u_{nvo} denotes the negative valve overlap, $u_{\mathrm{doi,f}}$ specifies the duration of gasoline injection, and $u_{\mathrm{doi,e}}$ specifies the duration of ethanol injection. The starts of fuel and ethanol injection are fixed in this study with 456 °CA and 580 °CA. These actuators were selected as control inputs because they directly affect the thermodynamic state of the in-cylinder charge and, consequently, the combustion phasing.

NVO traps hot residual gases from the previous cycle and increases the thermal energy of the in-cylinder charge, which is necessary for auto-ignition in throttle-free HCCI operation. Under this operating concept, gasoline mainly determines the released chemical energy, while ethanol provides an additional degree of freedom for combustion phasing control. Since ethanol is injected after intake valve closing, its evaporation reduces the temperature of the compressed unburned mixture before ignition. This effect is beneficial in the considered operating regime because it suppresses excessively advanced combustion, limits rapid heat release, and enhances cycle-to-cycle stability. Accordingly, the ethanol injection duration is introduced as an additional control input to regulate CA50 and broaden the stable operating range.

To facilitate the understanding of the employed HCCI operating concept, Fig. 3 shows a representative in-cylinder pressure trace together with the valve events, the start timings of gasoline and ethanol injection, and the NVO interval. In this study, the NVO interval is defined as the crank-angle interval between EVC and IVO and is arranged symmetrically around TDC.

The state and input constraints are:

$$
x_{\min} = \begin{bmatrix} 2\,°\mathrm{CA} \\ 1.5\,\mathrm{bar} \\ 0\,\mathrm{bar}/°\mathrm{CA} \end{bmatrix}, \quad x_{\max} = \begin{bmatrix} 13\,°\mathrm{CA} \\ 4.5\,\mathrm{bar} \\ 5\,\mathrm{bar}/°\mathrm{CA} \end{bmatrix}, \tag{44}
$$

Table 1
Single-cylinder research engine parameters (Bedei et al., 2025).

	Parameter	Value
Geometry	Displacement	499 cm^3
	Compression Ratio	12:1
	Bore	84 mm
	Stroke	90mm
Operation	Intake Pressure	1.013bar
	Exhaust Pressure	1.013bar
	Engine Speed	1500 min^{-1}
	Intake Temperature	50°
	Oil Temperature	105 °
	Coolant Temperature	90°
	Exhaust Valve Open	160 ° CA
	Intake Valve Close	545 ° CA
Gasoline	Research Octane Num.	96
	Motor Octane Num.	85
	Lower Calorific Value	44.3 MJ kg^{-1}
	Ethanol Mass Fraction	10.4 %
Ethanol	Research Octane Num.	106
	Motor Octane Num.	89
	Lower Calorific Value	26.7MJ kg^{-1}
Injection	Fuel Rail Pressure	100 bar
	Ethanol Rail Pressure	60 bar

$$u_{\min} = \begin{bmatrix} 167\,°\mathrm{CA} \\ 0.3\,\mathrm{ms} \\ 0\,\mathrm{ms} \end{bmatrix}, \quad u_{\max} = \begin{bmatrix} 219\,°\mathrm{CA} \\ 0.91\,\mathrm{ms} \\ 0.6\,\mathrm{ms} \end{bmatrix}. \tag{45}$$

All physical quantities are normalized in the control framework. Hence, all subsequent formulations and problem statements pertain to the normalized variables.

The control objective of the system is threefold. First, we need to maintain the CA50 value, which serves as an indicator of combustion phasing, within a specified range to avoid combustion that is either too early (potentially leading to knocking) or too late (resulting in combustion delay and reduced efficiency). Second, IMEP, which is directly related to engine output torque, should accurately track a time-varying reference reflecting the driver's demand. Third, we must keep DPmax, a critical indicator of engine mechanical stress, below a predefined safety threshold to ensure reliable and smooth operation.

3.2. Test bench configuration and engine specifications

We experimentally evaluate our algorithms on a single-cylinder research engine (Fig. 4) equipped with a fully variable electro-hydraulic valve train, in which two electromagnetic coils alternately control valve opening and closing.

The RCP control unit used is the dSPACE Microautobox III with Multi I/O Board DS1552B1, equipped with an ARM Cortex-A15 quad-core processor running at 1.4 GHz and a Kintex 7 field-programmable gate array from Xilinx with a computation rate of 12.5 ns. The processor handles high-level software control and data processing, whereas the FPGA is responsible for low-latency, highly parallel real-time signal processing tasks.

Table 1 lists the detailed parameters of the single-cylinder engine shown in Fig. 4.

3.3. Controller configuration

In this subsection, we introduce the formulations related to uncertainty, the controller design, and the setup of the MPC optimization problem.

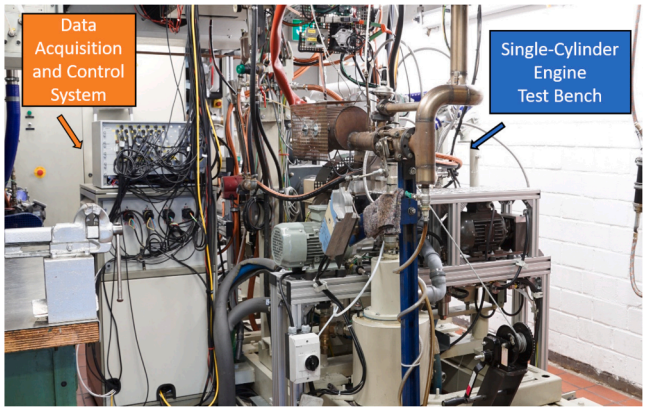


Fig. 4. Single-cylinder research engine at the chair of thermodynamics of mobile energy conversion systems, RWTH Aachen University.

3.3.1. Residual dataset

In the HCCI system, our primary concern is the impact of uncertainty on CA50 and IMEP. The direct influence of the disturbance term on DPmax is neglected. Instead, DPmax is handled primarily through the nominal prediction model together with the imposed safety constraint. This treatment is justified because, in HCCI combustion, DPmax is strongly correlated with combustion phasing and the overall heat-release process. Since the dominant uncertainty is characterized through CA50 and IMEP, a large portion of the induced variability in DPmax is already captured implicitly by the model.

The disturbance term w should be interpreted as a lumped residual uncertainty identified from the mismatch between model predictions and experimental measurements. It mainly reflects unmodeled cycle-to-cycle combustion variability, model mismatch, unmeasured variations in thermodynamic conditions such as residual gas effects, as well as measurement and estimation uncertainty. Accordingly, this does not imply that disturbance effects are isolated across cycles in the HCCI combustion process. Instead, the dominant inter-cycle coupling is already captured by the cyclic state transition model. Moreover, the i.i.d. assumption is only imposed on the residual term w over the short prediction horizon used in this work, which provides a practical approximation that keeps the distributionally robust chance-constrained formulation tractable.

3.3.2. Construction of the ambiguity set

The ambiguity set is constructed and implemented in four steps.

Step 1: Construction of the support set $\mathcal{W}$ of the uncertainty. From the residual dataset described above, the disturbance sample set is formed as $\mathcal{W}_{\text{data}} = \{w^{(s)}\}_{s=1}^{N_w}$, where each sample is given by $w^{(s)} = \begin{bmatrix} e^{(s)}_{\text{CA50}} & e^{(s)}_{\text{IMEP}} \end{bmatrix}^\top$. Based on these samples, the empirical mean and covariance are estimated as $\hat{\mu} = \frac{1}{N_w}\sum_{s=1}^{N_w} w^{(s)}$, and $\hat{\Sigma} = \frac{1}{N_w}\sum_{s=1}^{N_w}\left(w^{(s)} - \hat{\mu}\right)\left(w^{(s)} - \hat{\mu}\right)^\top$. PCA is then applied to $\hat{\Sigma}$ to identify the dominant directions of variability, denoted by $\{p_r\}$. The support set is constructed along these directions by projecting the residual samples onto each principal direction and determining the corresponding lower and upper bounds,

$$\underline{w}_r = \min_{s=1,\ldots,N_w} p_r^\top w^{(s)}, \qquad \overline{w}_r = \max_{s=1,\ldots,N_w} p_r^\top w^{(s)}. \tag{46}$$

This yields a data-driven support region that reflects the observed uncertainty in the experimental dataset.

Step 2: Construction of the lifted support set $\tilde{\mathcal{W}}$. To further enrich the distributional information, nonlinear moment constraints are introduced along each principal direction. Specifically, the piecewise-linear threshold functions $g_{rl}(w)$ are constructed according to (21), with the thresholds q_{rl} defined in (22). In this work, the number of thresholds is chosen as $K = 3$, corresponding to seven symmetric points along each

principal direction p_r. This results in multiple moment functions per direction, which capture distributional features beyond first- and second-order statistics. Based on these functions, the lifted support set is formulated as in (24).

Step 3: Construction of the ambiguity set $\tilde{\mathcal{M}}$. The corresponding moment bounds in (12) are computed as empirical expectations over the residual dataset,

$$\gamma_{rl} = \frac{1}{N_w}\sum_{s=1}^{N_w} g_{rl}\big(w^{(s)}\big), \qquad \forall l \in \mathbb{Z}_{[1,2K+1]}. \tag{47}$$

The ambiguity set in (12) is then defined as the family of probability distributions supported on the constructed support set and satisfying the resulting nonlinear moment constraints. In other words, it contains all distributions consistent with the support information and the empirical moment bounds extracted from the residual data. In the subsequent development, this ambiguity set is directly utilized to reformulate the chance constraints in (16) into the equivalent constraints in (18), thereby enforcing worst-case probabilistic guarantees over all admissible distributions. The violation probability is chosen as $\epsilon = 0.1$ to balance robustness and performance.

Step 4: Computation of the tightened margins α and κ. Based on the ambiguity set constructed above, the distributionally robust chance constraints are further converted into a deterministic tightened form. To separate the effects of the nominal dynamics and the stochastic error dynamics, the predicted state is decomposed as $x_k = z_k + e_k$, where z_k denotes the nominal state and e_k the disturbance error. The effect of the stochastic error is then incorporated through the tightening margins α_k and κ_k, such that the original chance constraints are replaced by tightened deterministic constraints on the nominal trajectory. The margins α_k and κ_k are obtained by solving the auxiliary optimization problems in (28) and (35). They are chosen as the smallest values that guarantee satisfaction of the original constraints for all admissible distributions in the ambiguity set. As a result, the original SMPC problem is transformed into a deterministic MPC problem with tightened constraints.

Remark 5. By construction, the ambiguity set $\tilde{\mathcal{M}}$ is guaranteed to be non-empty. Specifically, let $\hat{P}$ denote the empirical distribution induced by the residual samples $\{w^{(s)}\}_{s=1}^{N_w}$. Since the support set $\mathcal{W}$ is constructed directly from these samples, it follows that all realizations satisfy $w^{(s)} \in \mathcal{W}$. Moreover, the nonlinear moment bounds γ_{rl} are defined as empirical expectations under $\hat{P}$, i.e., $\gamma_{rl} = \mathbb{E}_{\hat{P}}[g_{rl}(w)]$. The use of empirical moment bounds provides a data-driven characterization of uncertainty and ensures consistency with the observed residual data. Therefore, $\hat{P}$ itself satisfies all constraints defining $\tilde{\mathcal{M}}$, which ensures that the ambiguity set is non-empty and consistent with the observed data. The resulting ambiguity set contains all probability distributions supported on $\mathcal{W}$ that are consistent with the empirical nonlinear moment information extracted from the residual data.

Furthermore, since the support set $\mathcal{W}$ is compact and each function $g_{rl}(w)$ is piecewise linear, all nonlinear moments are finite, which guarantees that the ambiguity set is well-posed. The richness of the ambiguity set is primarily controlled by the number of thresholds K, which determines the expressiveness of the nonlinear moment functions and thus the level of conservatism. A larger K leads to a more detailed characterization of the uncertainty distribution and tends to reduce conservatism in practice.

Remark 6. The parameter K determines the number of threshold points introduced along each principal direction in the ambiguity-set construction in (22), and thus directly affects the expressiveness of the uncertainty representation of the LP-based tightening step.

In this work, $K = 3$ is selected as a practical compromise. Under the PCA-based scaling, the diagonal elements of Λ reflect the uncertainty level along each principal direction, and the resulting threshold points already capture the dominant spread of the projected disturbance samples. Increasing K further mainly introduces additional threshold points in regions where relevant samples are already very sparse or correspond to outliers, which typically provides limited additional information while increasing the computational burden.

In general, the choice of K should be adapted to the distribution of the uncertainty. For disturbances with more complex or heavy-tailed distributions, a denser placement of threshold points or a larger K may be beneficial to improve the representation accuracy.

Remark 7. The violation probability ϵ determines the level of conservativeness in the chance-constraint approximation, and thus directly affects the trade-off between robustness and control performance.

A smaller ϵ leads to stricter constraint tightening, which improves robustness against uncertainty and reduces the risk of constraint violations, but may also result in more conservative control actions and degraded tracking performance. In contrast, a larger ϵ allows more aggressive control behavior and can improve nominal performance, but at the expense of reduced robustness. This effect becomes more critical when the uncertainty distribution contains large outliers or heavy tails, since an overly small ϵ may cause the controller to account for extreme disturbance realizations, leading to substantial performance degradation.

In this work, $\epsilon = 0.1$ is selected as a practical compromise between robustness and performance for the considered HCCI control problem.

3.4. MPC configuration

Since the HCCI system exhibits rapid dynamic response, we select a prediction horizon $N = 4$ for the optimal control problem in (36), and the chosen discrete-time dynamical model in (1) shows the cyclic state transition.

The objective function adjusted according to the control targets is

$$\sum_{i=0}^{N-1}\Big(\big\|z_{i|k} - z_{\text{ref}}\big\|_Q^2 + \big\|\Delta c_{i|k}\big\|_R^2\Big) + \big\|z_{N|k} - z_{\text{ref}}\big\|_{Q_T}^2 + \rho_1 s^2 + \rho_2 s. \tag{48}$$

Here, we penalize the discrepancies between the system states and their corresponding references, as well as the variation in the control inputs. Specifically, the real reference values for CA50 and DPmax are set to 7 °CA and 0 bar/°CA, while the real reference for IMEP varies with time. In order to ensure the feasibility of the overall optimization problem, we also introduce a global slack variable s to quantify the maximum constraint violation margin, namely the infinity norm of the violations. The weighting matrices are defined as follows: the state weight matrix is given by $Q = \text{diag}(1, 5, 0.1)$, where we assign a higher weight to tracking performance of IMEP, the input weight matrix is $R = \text{diag}(0.5, 0.5, 0.5)$, reflecting an equal penalization of variations in control inputs, the terminal state weight matrix comes from (10), and the slack variable s is penalized with a quadratic weight ρ_1 of 2 and a linear weight ρ_2 of 4. We select the penalty ρ_2 to satisfy $\rho_2 > \mu^*$ in most cases, where μ^* is the optimal dual multiplier, thereby satisfying the condition for exact penalty. Still, we cannot choose ρ_2 excessively large to prevent numerical issues during optimization.

3.5. Learning configuration

On the experimental test bench, we collected data from 1000 continuous cycles across 65 scenarios (65000 cycles in total) using the measurement algorithm in Bedei et al. (2023) that yields a highly diverse dataset over the considered load range. We use 70 % of the collected data for training, 20 % for validation, and the remaining 10 % for testing.

During the imitation learning phase, we randomly sample 10000 cycles from the training set for policy learning. This choice is motivated by the computational burden of the adopted Q-loss-based framework, since evaluating the loss in (41) requires solving an optimal control problem for each sample.

Each data point used for imitation learning is constructed from one engine cycle. Specifically, each data point used for policy learning consists of the current measured state $x_k = \left[x_{\mathrm{CA50},k}, x_{\mathrm{imep},k}, x_{\mathrm{DPmax},k}\right]^\top$, the previous control input $u_{k-1} = \left[u_{\mathrm{nvo},k-1}, u_{\mathrm{doi,f},k-1}, u_{\mathrm{doi,e},k-1}\right]^\top$, and the current reference $r_k = x^{\mathrm{ref}}_{\mathrm{imep},k}$. In this work, the reference input to the policy corresponds only to the IMEP target, which is sampled from a uniform distribution within the range of 2 to 4bar. The CA50 reference is fixed at 7 °CA and is not included as an input to the neural network, but is taken into account in the optimization problem (41). For each sampled data point, the loss is defined through the optimal control problem in (41). By solving this problem, we obtain the gradient of the loss with respect to the network output from the corresponding Lagrange multipliers, which is then used to update the network parameters through backpropagation.

To improve training efficiency and learning performance, the policy model is trained to predict a control increment rather than the absolute control input. Specifically, the network outputs a residual control action Δu_k, and the final control input is computed as $u_k = u_{k-1} + \Delta u_k$. In addition, all variables are normalized consistently with the control framework before being used for training.

In practical implementation, we can precompute the tightened margin α and κ, and the linearized system with feedback gain before training. We limited the number of iterations used to compute ℓ_{MPC} to balance training time and convergence performance. Specifically, the OCP solver employs an interior-point method with a maximum of 6 iterations. Although the generated gradients of the objective with respect to the policy model output are approximate and contain random noise due to the limited iterations, the parameter update is faster and requires less computation.

The training process updates the neural network parameters using the Adam optimizer. The policy model π_θ uses a fully connected architecture with three hidden layers, each containing 40 neurons. The activation functions for the first two layers are leaky-ReLU, and the final layer uses hyperbolic tangent to ensure that the output satisfies the imposed control input constraints. In this work, the adopted network architecture is selected as a practical compromise between approximation capability and computational complexity, with the aim of maintaining sufficient expressive power while preserving computational efficiency for real-time implementation.

3.6. Evaluation protocol

In this subsection, we first evaluate the control performance of different MPC methods on the HCCI engine system under random additive uncertainty in a simulation environment. Then, we assess the closed-loop performance of the learned controller, trained using the imitation learning framework in (41), on the experimental test bench.

In the HCCI system controller design, we focus on two objectives. First, we aim to keep the CA50 value within a reasonable range and to minimize the variance of CA50. Thus, we seek to evaluate whether the SMPC, with the incorporation of uncertainty propagation, adopts appropriate actions to proactively avoid potential constraint violations. Second, since the noise of IMEP is smaller compared to CA50, we primarily examine its mean performance to assess the reference tracking capability of the different controllers.

To ensure a meaningful and problem-driven evaluation, the benchmark controllers are not selected as a collection of competing methods, but rather to systematically investigate the role of uncertainty information in predictive combustion control.

First, the nominal MPC is used as the primary baseline. It represents a widely used deterministic predictive control formulation in HCCI applications, where model-based optimization is performed without explicit stochastic uncertainty propagation. To the best of our knowledge, most existing real-time HCCI control implementations rely on such deterministic MPC formulations.

Second, a second-order moment-based SMPC (SM-SMPC) is included to represent the most commonly adopted approach for incorporating uncertainty in stochastic distributionally robust MPC. In this formulation, the ambiguity set is constructed using only the first- and second-order statistics (mean and covariance), which are widely used due to their tractability and minimal data requirements.

Third, the proposed NoM-DiR-SMPC extends the framework based on second-order moment information by further enriching the ambiguity set. It incorporates additional structured moment information to better capture the distributional characteristics of the disturbances, while preserving computational tractability through a reformulation that leads to LP-based constraint tightening and a nominal QP problem.

With this benchmark design, the objective of the comparison is not only to evaluate performance differences, but more fundamentally to answer the following questions: (i) whether incorporating uncertainty information improves HCCI control performance compared to deterministic MPC, and (ii) whether higher-order uncertainty representations can reduce the conservatism associated with conventional second-moment-based approaches.

In the SM-SMPC case, where the ambiguity set is defined as a set of all distributions with the same first and second-order moments, the violation margin is

$$\alpha_{ij} = \upsilon\sqrt{D_j \Sigma_e^i D_j^\top}, \quad \upsilon^2 = \frac{\epsilon}{1-\epsilon} \tag{49}$$

where Σ_e^i is the covariance matrix of error at prediction instant i (Calafiore & Ghaoui, 2006).

In the simulation study, we selected 60 simulation runs, with all disturbances generated randomly. Each run consists of 120 consecutive cycles, and the IMEP reference is time-varying.

In order to test different MPC controllers on the test bench, we need to ensure real-time feasibility. To achieve this, we utilize the imitation learning framework presented in (41) to train two policy models that respectively learn the MPC and NoM-DiR-SMPC strategies. To quantify performance, we assess how closely the learned models reproduce the control actions of the original MPC policy with the use of 1000 test samples.

We experimentally compare MPC and NoM-DiR-SMPC in Algorithm 1 on our single-cylinder research engine. The trained policy models were integrated into the Rapid Control Prototyping hardware of the test bench. In the experimental tests, we primarily evaluate the controllers based on their ability to track the IMEP reference, the stability of CA50, and the effective operating range of the engine. For the test scenarios, we employ time-varying IMEP reference profiles, including sinusoidal, step, and ramp profiles, to assess the controllers' adaptability and performance under various operating conditions. The reference of CA50 is 7 °CA to ensure efficient combustion, and we need to hold the DPmax within a specified bound of 5 bar/°CA to guarantee mechanical integrity.

4. Results

This section shows the results of the NoM-DiR-SMPC for the HCCI system in both simulation and experimental studies. We compare it against a nominal MPC baseline and a second-order moment-based SMPC variant, focusing on their performance across different operating conditions.

4.1. Simulation results

Fig. 5 presents the closed-loop performance of CA50 with the introduced control methods. The box plot in Fig. 5 illustrates the variation in CA50, including the interquartile ranges, whiskers, and outliers.

As illustrated in the box plot of Fig. 5, when employing both SMPCs, the interquartile ranges (25 % to 75 %) and whiskers of CA50 variations are smaller compared to those observed with the classical MPC. Additionally, both SMPCs reduce the number of outliers relative to MPC.

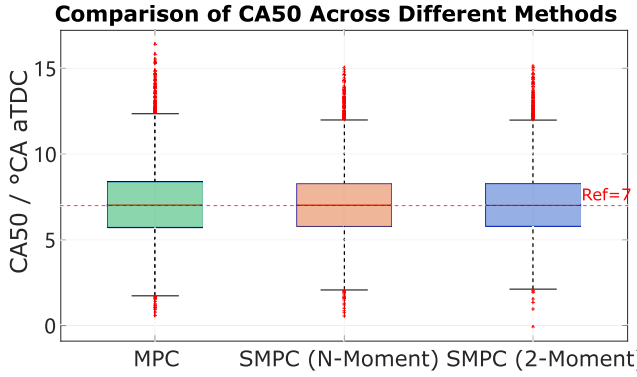


Fig. 5. Comparison of CA50 across different control methods. The box plot illustrates the distribution of CA50 (°CA) for three approaches: MPC, NoM-DiR-SMPC (Nonlinear-Moment), and SM-SMPC (2-Moment) under 7200 sample data. The horizontal red dashed line represents the reference value of CA50 at 7 °CA. The plot indicates extreme cases as red dots above and below the whiskers. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 2
Comparison of CA50 (° CA) variation under different controllers.

	Mean	Variance	≥ 13	≤ 2
MPC	7.13	4.55	1.01%	0.47%
NoM-DiR-SMPC	7.12	3.81	0.73%	0.25%
SM-SMPC	7.12	3.83	0.69%	0.08%

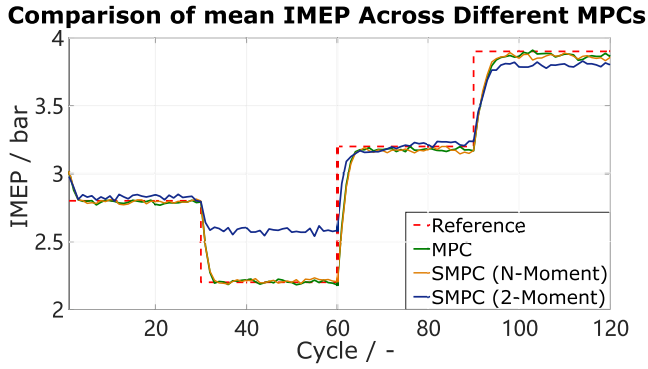


Fig. 6. Comparison of mean IMEP among various control methods, including MPC, NoM-DiR-SMPC (N-Moment), and SM-SMPC (2-Moment). The red dashed line represents the reference IMEP trajectory. The results show the response of each control strategy to different operating conditions. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The analysis presented in Table 2 shows no quantifiable difference in the mean CA50 values among the various controllers. However, regarding the variance of CA50, MPC results in a larger variance compared to SMPCs. Furthermore, in terms of constraint violations, the frequency of violations under SMPCs is lower than that under MPC. Therefore, in terms of CA50 related metrics, both SMPC approaches provide better stability and fewer constraint violations than MPC. Specifically, the CA50 variance is reduced from 4.55 under MPC to 3.81 under NoM-DiR-SMPC, corresponding to a reduction of approximately 16.3 %. In addition, the total CA50 constraint violation rate, defined as the sum of upper-bound and lower-bound violations, decreases from 1.48 % to 0.98 %, corresponding to a reduction of approximately 33 %.

In Fig. 6, we show the mean performance of IMEP under different controllers, where the reference is set to a step-change profile. This illustrates the tracking capability of each controller.

Table 3
Comparison of IMEP (bar) deviations from the reference under different controllers.

	MPC	NoM-DiR-SMPC (N)	SM-SMPC
RMSE	0.1560	0.1563	0.2288

Table 4
Comparison of learned model and real controller outputs.

	MPC R^2	SMPC R^2	MPC NRMSE	SMPC NRMSE
Fuel	0.9878	0.9810	0.0108	0.0134
NVO	0.9755	0.9546	0.0027	0.0037
Ethanol	0.9421	0.9320	0.0416	0.0447

Table 5
worst-case deviation of the learned policy from the MPC/SMPC solutions.

	Fuel (ms)	NVO (° CA)	Ethanol (ms)
SMPC	0.0226	1.8317	0.0395
MPC	0.0265	1.3632	0.0324

At low IMEP values, the SM-SMPC using second-order moment information exhibits a notable decline in tracking capability. This SMPC is fails to follow the reference at low IMEP levels. Similarly, at high IMEP references, SM-SMPC also shows relatively poor tracking performance compared to the other two controllers.

When evaluating the IMEP reference tracking performance in Table 3, the difference between the standard MPC and NoM-DiR-SMPC is not particularly pronounced. NoM-DiR-SMPC exhibits slightly larger offsets under both low and high IMEP operating conditions. However, these offsets remain within the small range.

In the simulation study, all three controllers are compared to examine their behavior under different ways of handling uncertainty. The results show that, although SM-SMPC improves CA50 robustness, it is considerably more conservative and leads to degraded IMEP tracking performance, especially near the operating limits.

In contrast, the experimental study is intended to evaluate real-time feasibility and practical closed-loop performance on the engine test bench. For this reason, only the nominal MPC and the proposed NoM-DiR-SMPC are carried forward to the experimental validation. These two controllers are more representative of practically deployable solutions in the considered HCCI operating regime, while SM-SMPC is mainly retained as a simulation benchmark to illustrate the limitation of using only second-order moment information, which in this case results in excessive conservatism and unsatisfactory tracking performance.

We employ imitation learning methods in (41) to learn the strategies of both the standard MPC and the proposed NoM-DiR-SMPC separately and evaluate the learned models on the given test samples to compare their control inputs with those of the true MPC solutions. Fig. 7 shows a comparison between the control actions generated by the learned models and those computed by solving the optimization problem online. The comparison focuses on three key control inputs, namely DOI fuel, NVO, and DOI ethanol, of explicit (learning) and implicit (MPC) methods to evaluate the fidelity of the learned policies.

Fig. 7 illustrates the differences between the strategies generated by the parameterized models and those obtained by solving the MPC directly. From the comparisons of the DOI fuel and NVO, it is evident that the learned models closely match the true strategies, nearly overlapping with the red reference lines. In the case of DOI ethanol, the generated values are slightly lower than the true values at high DOI levels, yet still remain within a reasonable range.

Table 4 presents quantitative statistics using R^2 and normalized RMSE to assess the effectiveness of the trained parameterized models in fitting the true MPC control inputs. Both the learned SMPC and MPC

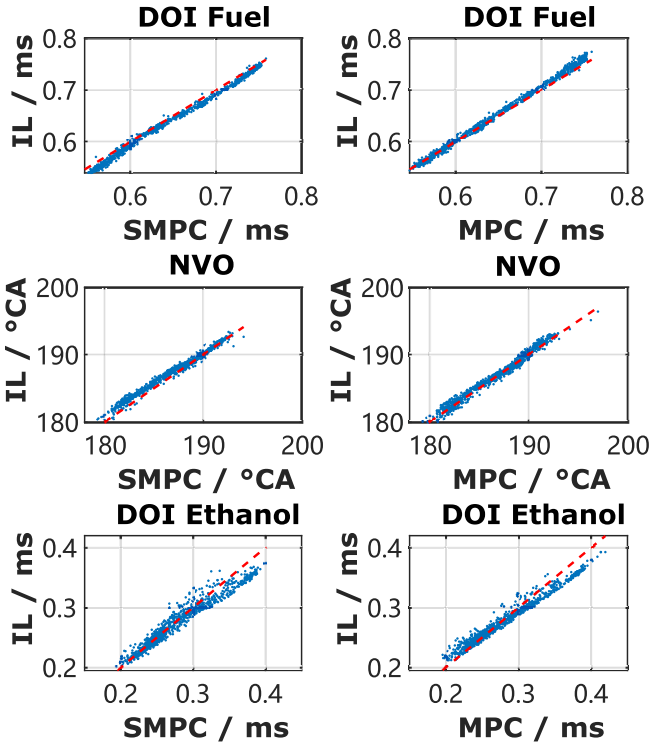


Fig. 7. Comparison of control inputs generated by imitation learning and MPC methods. The control inputs via MPC and NoM-DiR-SMPC are set to the ground truth. The red dashed line represents the perfect agreement. The blue dots show the data points indicating the learning accuracy. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

models achieve R^2 values exceeding 0.95 for DOI fuel and NVO, and above 0.90 for DOI ethanol, indicating a strong fit that effectively captures the variability in the data. Additionally, all normalized RMSE values remain below 5 %, demonstrating minimal prediction errors and superior model fitting performance, particularly for inputs DOI fuel and NVO. In addition, Table 5 summarizes the maximum absolute errors between the learned policies and the corresponding MPC solutions, highlighting the worst-case deviations. These results provide additional insight into the worst-case approximation performance and show that the deviations remain within an acceptable range for the considered control inputs.

4.2. Experimental results

In our test bench experiments, we primarily compared the performance of the learned MPC and NoM-DiR-SMPC controllers. Fig. 8 illustrates the tracking performance of MPC and SMPC under time-varying IMEP, with the black dashed line representing the reference. The upper and lower IMEP values that can be stably tracked determine the effective operating range of the engine under the given control.

As depicted in the IMEP trajectories of Fig. 8, the system operating under the MPC strategy experienced multiple incomplete late combustion cycles, evidenced by drops in IMEP values. This indicates that the combustion stability under MPC is inferior to that achieved using NoM-DiR-SMPC.

Regarding reference tracking capability, both MPC methods produce IMEP outputs that closely align with the reference, demonstrating satisfactory performance. To determine the effective operating ranges of the two MPC methods, we observe that for MPC, the operating range is between 3 to 4 bar, whereas for SMPC, the effective range extends from

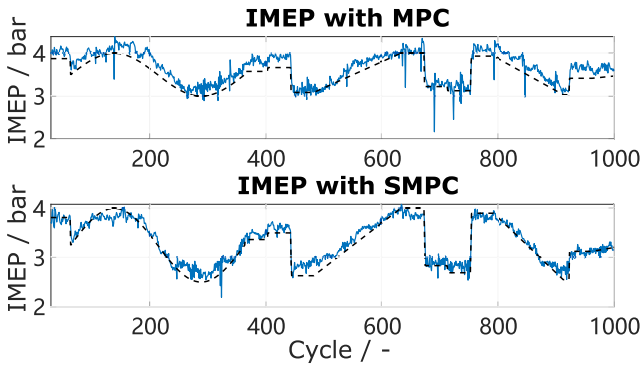


Fig. 8. Comparison of IMEP trajectories over 1000 engine cycles under MPC and NoM-DiR-SMPC methods. The solid blue lines represent the actual IMEP responses, while the dashed black lines indicate the reference trajectories. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 6
Comparison of tracking error and fuel-related actuation.

	IMEP RMSE (bar)	average Fuel (ms)	average Ethanol (ms)
SMPC	0.1351	0.4997	0.3462
MPC	0.2066	0.5536	0.3205

Table 7
Closed-loop safety and stability metrics based on CA50 (°CA).

	Var.	$\mathbb{P}(\geq 13)$	$\mathbb{P}(\leq 2)$	Maximum
SMPC	5.24	3.9 %	0.3 %	18.85
MPC	6.64	5.5 %	0.2 %	25.55
	Minimum	**Ext. events (≥ 20)**	**Ext. events (≤ 0)**	
SMPC	−4.45	0	2	
MPC	−0.25	5	1	

2.5 to 4 bar, indicating that the SMPC controller allows stable operation at lower IMEP values compared to the nominal MPC.

It is worth mentioning that the stable operating ranges depicted in Fig. 8 were ultimately established through repeated experimentation. When the controllers operate outside these stable ranges shown in Fig. 8, a large amount of irregular combustion and even misfire could occur. This extension highlights the potential advantages of the proposed NoM-DiR-SMPC compared to MPC.

Table 6 summarizes the experimental closed-loop performance quantitatively. Compared with the learned MPC controller, the learned SMPC controller reduces the IMEP tracking RMSE from 0.2066 bar to 0.1351 bar. At the same time, it lowers the average fuel injection duration, while using slightly more ethanol injection, which is associated with improved combustion stabilization.

Table 7 summarizes the closed-loop safety and stability metrics in terms of CA50 for the learned controllers. Compared with the learned nominal MPC, the learned SMPC achieves a lower variance (5.24 vs. 6.64). In addition, the violation probability of the upper CA50 bound (13 °CA) is reduced from 5.5 % to 3.9 %, while the lower bound violation remains at a comparable level. The maximum CA50 value is reduced from 25.55 °CA to 18.85 °CA. In terms of extreme events, no occurrences of CA50 exceeding 20 °CA are observed with the learned SMPC under the considered test condition, whereas the learned nominal MPC exhibits multiple such events. Events with CA50 below 0 °CA are rare for both controllers, with only minor differences observed.

Fig. 9 shows the trajectories of CA50 and DPmax over 1000 engine cycles under both MPC and NoM-DiR-SMPC controllers. The reference value for CA50 is set to 7 °CA, and the constraint threshold for DPmax

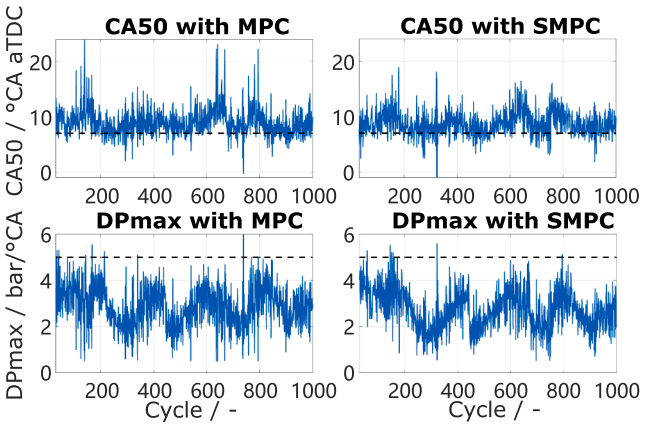


Fig. 9. Comparison of CA50 and DPmax trajectories over 1000 engine cycles under MPC and NoM-DiR-SMPC strategies. The reference of CA50 shown as a dashed line is set to 7 °CA, and the boundary of DPmax presented by a dashed line is set to 5 bar/°CA.

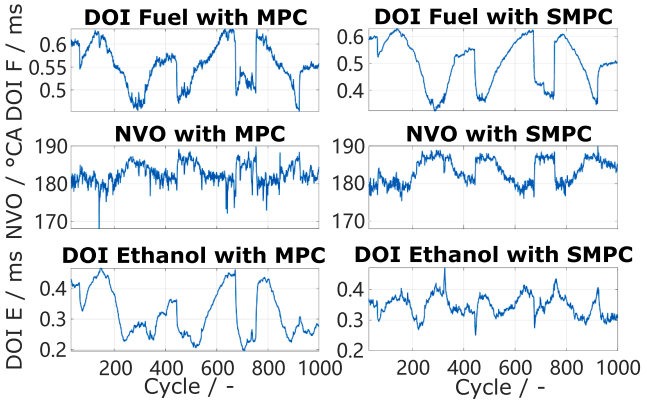


Fig. 10. Control inputs for DOI Fuel, NVO, and DOI Ethanol over 1000 engine cycles under MPC and NoM-DiR-SMPC strategies.

is 5 bar/°CA. Fig. 10 shows the profiles of the control inputs generated by each controller during operation.

The results in Fig. 9 show that, compared to MPC, SMPC leads to fewer extreme deviations in CA50, indicating improved robustness to process uncertainties mproved robustness. This observation correlates with the IMEP results, as incomplete combustions (indicated by low IMEP) are generally associated with larger CA50 values. Due to the coupling between cycles, a late combustion in one cycle may lead to premature combustion in the subsequent cycle, manifested as an early combustion phasing. In terms of DPmax comparison, both controllers perform quite well, maintaining values within the 5 bar/°CA threshold throughout the test cycles.

Overall, the experimental results indicate that SMPC provides improved combustion stability and supports a wider operational range than MPC, while maintaining comparable DPmax performance. Additionally, during the SMPC test cycles, a single incomplete combustion event was observed at approximately the 320th cycle. The controller promptly responded by reducing NVO with less residual gas and increasing ethanol injection time to enhance cooling, as demonstrated in Fig. 10.

To further evaluate the computational efficiency of the considered controllers, the computation time is profiled over 200 consecutive engine cycles in simulation, and the results are summarized in Table 8.

For clarity, the reported computation time of the optimization-based controllers includes the complete online control pipeline at each cycle: model linearization, feedback gain and terminal weight update, constraint update, and the solution of the optimization problem using the embedded HPIPM solver (Frison & Diehl, 2020).

Table 8
Comparison of computation time of controllers.

Controllers	Avg. time (ms)	Max. time (ms)
learned MPC	0.07	0.10
learned NoM-DiR-SMPC	0.07	0.11
nominal MPC	8.9	11.2
SM-SMPC	9.7	12.0
NoM-DiR-SMPC + tightening	16.1	20.3
LP tightening	8.2	9.1

Regarding real-time feasibility, it should be noted that in practical HCCI engine control, the available computation time is highly constrained. In our test bench, at 1500 min^{-1}, the available time window from measurement update to control action is approximately 6.7 ms, which further reduces to about 1 ms at 3500 min^{-1}. In addition, ECU task scheduling typically introduces additional delays, further shrinking the effective computation time (Nuss, 2024).

5. Discussion

By analyzing both simulation and experimental results, we aim to understand how the proposed NoM-DiR-SMPC controller enhances robustness and tracking accuracy compared to traditional MPC and second-moment-based SMPC.

As shown in Fig. 5, both SMPC methods achieve improved CA50 stability compared with classical MPC. By explicitly accounting for uncertainty propagation over the prediction horizon, they effectively prevent CA50 constraint violations and further reduce the variability range of CA50. Notably, SM-SMPC with the second-order moment exhibits fewer constraint violations than the proposed NoM-DiR-SMPC. This can be attributed to the fact that SM-SMPC employs a more conservative ambiguity set, defined solely by the first and second moments of the uncertainty distribution. In contrast, NoM-DiR-SMPC considers higher-order moment information, leading to a smaller ambiguity set that is effectively a subset of the second-order moment set. As a result, while NoM-DiR-SMPC achieves a more accurate representation of the uncertainty, SM-SMPC adopts a more cautious strategy regarding constraint handling, which translates into slightly reduced CA50 variation.

Unexpectedly, the proportion of constraint violations generally remains below 1 %, which does not align with our predefined allowable violation probability ϵ of 10 %. This discrepancy likely arises because the engine operates under stable conditions in most cases, and CA50 remains within the constraints despite the influence of disturbances.

Although the SM-SMPC achieves a smaller CA50 variation, its overly conservative ambiguity set noticeably reduces control flexibility, which in turn leads to poor IMEP tracking performance, as demonstrated in Fig. 6. At low IMEP levels, the imposed constraints become so restrictive that the controller fails to follow the reference accurately. In practical applications, especially for systems like HCCI that are inherently highly stochastic, such a conservative approach renders SM-SMPC infeasible. This limitation is the primary reason why we excluded SM-SMPC from the subsequent experimental comparisons.

Based on the simulation comparisons, NoM-DiR-SMPC demonstrates a well-balanced trade-off between robustness (CA50 stability) and performance (IMEP reference tracking). By leveraging a more precise ambiguity set, it offers a higher degree of feasible control flexibility.

In the learning phases, the results in Fig. 7 demonstrate that the imitation learning framework is effective in replicating the MPC behavior. The close alignment between predicted and true inputs, particularly in the case of fuel and NVO, shows the strong learning capability of the framework. Although the ethanol injection control shows slightly more deviation, especially at higher values, the error is still within a reasonable range. Overall, the imitation learning strategy used here al-

lows for a highly accurate, real-time implementable approximation of the original MPC and SMPC policies. This noticeably reduces online computational effort without compromising control quality.

It should be emphasized that the experimental comparison is performed with learned controllers. Therefore, the closed-loop behavior observed in the engine experiments reflects not only the differences between the underlying control strategies, but also the approximation errors introduced by imitation learning. In real-time engine experiments, this influence cannot be completely excluded, and its effect on the final performance cannot be fully isolated. For this reason, the simulation study plays an important role in the present work, as it allows the intrinsic differences among the control strategies to be examined more directly, without the additional approximation error introduced by learning. In addition, the fidelity of the learned policies is quantitatively evaluated in Fig. 7 and Table 4. The effect of policy approximation is reduced to a limited level, although it cannot be regarded as negligible. Under this interpretation, the experimental results should be understood primarily as evidence of the real-time applicability of the learned control policies, while the simulation results more clearly reveal the performance differences of the underlying control formulations.

The results in Fig. 8 suggest that the NoM-DiR-SMPC controller achieves better combustion stability than MPC, especially under low-load conditions. The observed IMEP drops under MPC reflect late or incomplete combustion cycles, which are effectively avoided by NoM-DiR-SMPC. This highlights the benefits of explicitly incorporating uncertainty handling into the controller design. Notably, NoM-DiR-SMPC enables a broader stable operating range, extending the low-load limit down to 2.5 bar without inducing misfire, whereas the MPC approach maintains stability only above 3 bar. The broader stable range under NoM-DiR-SMPC is particularly valuable for practical HCCI operation.

The CA50 trajectories shown in Fig. 9 reveal that NoM-DiR-SMPC reduces extreme fluctuations more effectively than MPC, indicating improved robustness to process uncertainties in combustion timing. This observation aligns with the earlier IMEP results, where MPC exhibits undesirable fluctuations. Moreover, because HCCI combustion is strongly coupled across cycles, late combustion in one cycle may lead to early combustion in the next. The reduced variability observed under NoM-DiR-SMPC implies enhanced mitigation of such disturbance propagation across cycles.

In addition to the improved tracking performance, Table 6 shows that the SMPC controller reduces the average fuel injection duration, while slightly increasing the ethanol injection duration. Under the fixed injection conditions used in this study, the injected fuel mass is directly related to the injection duration. Therefore, the reduced average fuel injection duration indicates a lower average fuel consumption. The increased ethanol injection under SMPC can be interpreted as a trade-off for improved combustion stability and robustness. By allocating more actuation effort to ethanol injection, the controller is able to better regulate the combustion process under uncertainty.

Table 7 evaluates the performance of the learned controllers from a safety and reliability perspective rather than purely tracking performance. In particular, the reduction of extreme CA50 excursions and constraint violation indicates that the learned SMPC is better able to handle unfavorable combustion variations. From a physical standpoint, large CA50 values correspond to significantly delayed combustion, which is typically associated with inefficient heat release, while very small CA50 values indicate excessively advanced combustion. Both types of events are undesirable in practice, as they may lead to unstable operation or increased mechanical stress. The observed reduction in such extreme events suggests that the uncertainty-aware formulation underlying the SMPC is effectively preserved in the learned policy. Under the tested conditions, the learned SMPC achieves a more consistent and well-contained combustion behavior, indicating improved practical reliability.

In these experimental results, NoM-DiR-SMPC tends to adopt more conservative control strategies due to its proactive consideration of uncertainty propagation during prediction. In contrast, MPC, in pursuit of optimizing the performance index, often selects overly aggressive control actions, disregarding the potential impact of uncertainties on process stability. For systems such as HCCI that are highly sensitive to stochastic uncertainty, NoM-DiR-SMPC demonstrates strong potential for practical application.

From the perspective of computational efficiency, Table 8 indicates that the average online computation time of the optimization-based controllers exceeds the available 6.7 ms time window of our test bench at $1500\,\mathrm{min}^{-1}$. This challenge becomes even more critical at higher engine speeds, where the available feedback window is further reduced. Therefore, while the optimization-based controllers provide meaningful benchmarks for performance evaluation, their direct real-time implementation remains difficult under practical constraints. This result further highlights the importance of the proposed learning-based approach, which retains the performance benefits of NoM-DiR-SMPC while substantially reducing the online computation time.

Additionally, the learning method resolves the real-time feasibility issues arising from the optimization-based controller. The benefits of achieving real-time performance extend beyond the scope of these tests. Specifically, the test cycles conducted in this study were performed at a constant engine speed of $1500\,\mathrm{min}^{-1}$. The feasible computation time for optimization-based controllers at higher-speed operation is much more limited. In contrast, learning-based models demonstrate greater potential for scalability and adaptability.

However, our method has certain limitations. Firstly, we did not incorporate filtering, meaning that state feedback is subject to noise. Secondly, the neural network-based controllers are highly sensitive to input variations, resulting in high-frequency changes in control inputs as shown in Fig. 10. Thirdly, in the considered HCCI setting, the current controller design does not explicitly optimize fuel consumption or efficiency. As such, the proposed method should be interpreted primarily as a framework for improving combustion stability, tracking performance, and constraint handling, rather than for achieving fuel-optimal operation. Fourthly, the experimental validation in this work is conducted on a single-cylinder engine platform at a fixed engine speed. Accordingly, the reported robustness improvement and the observed extension of the stable low-load operating region should be understood as findings under the tested conditions, rather than as general conclusions for all HCCI operating scenarios. Additional validation under different engine speeds, thermal conditions, and wider operating conditions is required to further assess the generality of the proposed method.

6. Conclusion

This work presented the nonlinear moment-based distributionally robust stochastic MPC (NoM-DiR-SMPC) for HCCI engines and the imitation learning method to transform the SMPC policy into an explicit controller model. Results indicated that incorporating higher-order moment information into the ambiguity set enabled the proposed approach to achieve a better balance between robustness and performance than traditional second-order moment-based methods.

Results from simulations and experiments demonstrated that NoM-DiR-SMPC improved combustion stability, reduced constraint violations, and extended the stable low-load operating range compared to conventional MPC, while retaining real-time applicability under the tested operating conditions. These findings highlight the potential of the proposed methods for high-variability control applications such as HCCI combustion systems.

Future work will focus on integrating more reliable state estimation mechanisms and exploring interpretable policy architectures, such as neuro-fuzzy systems, to further enhance robustness and safety. Additionally, incorporating economic objectives like fuel efficiency into the control formulation could unlock broader practical impact. Furthermore,

extending the experimental validation to multiple operating conditions will be an important direction to assess the generality of the proposed framework.

CRediT authorship contribution statement

Xu Chen: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization; **Julian Bedei:** Writing – review & editing, Validation, Software, Resources, Investigation, Data curation; **Jakob Andert:** Writing – review & editing, Resources, Investigation, Funding acquisition, Data curation, Conceptualization; **Lorenz Dörschel:** Writing – review & editing, Methodology, Investigation; **Heike Vallery:** Writing – review & editing, Supervision, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be consideredas potential competing interests: Xu Chen reports financial support was provided by German Research Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This study was part of Research Unit 2401 with DFG project number 277012063, titled “Optimization-Based Multiscale Control of Low-Temperature Combustion Engines”, and was supported by the German Research Foundation (DFG). We sincerely appreciate their invaluable support.

Appendix A. Duality justification for the min-max reformulation

This appendix provides the technical justification for the min-max interchange used in the reformulation. We first recall the following duality result, adapted from Shapiro and Kleywegt (2002), and then verify its assumptions for the ambiguity set considered in this work.

Theorem 1. *Let $X \subseteq \mathbb{R}^n$ be a closed convex decision set, and let $\mathfrak{P}$ be a nonempty family of probability measures on the measurable space $(\Omega, \mathcal{F})$. Consider the optimization problem*

$$\inf_{u\in X} \sup_{\mathbb{P}\in\mathfrak{P}} \int_\Omega \ell(u,\omega)\,\mathbb{P}(d\omega), \tag{A.1}$$

where $\ell : \mathbb{R}^n \times \Omega \to \mathbb{R}$ is a real-valued function. Suppose that an optimal solution $u^\star$ of (A.1) exists, and that the following assumptions hold.

1. *The family $\mathfrak{P}$ is tight and closed with respect to weak convergence.*
2. *There exists a convex open set $U \subseteq \mathbb{R}^n$ with $X \subseteq U$ such that, for every $u \in U$ and every $\mathbb{P} \in \mathfrak{P}$, the mapping $\omega \mapsto \ell(u,\omega)$ is measurable and integrable, and*

   $$\sup_{\mathbb{P}\in\mathfrak{P}} \int_\Omega \ell(u,\omega)\,\mathbb{P}(d\omega) < \infty.$$

   Moreover, for each fixed $\omega \in \Omega$, the function $u \mapsto \ell(u,\omega)$ is convex on U.
3. *For every u in a neighborhood of $u^\star$, the function $\omega \mapsto \ell(u,\omega)$ is bounded and upper semicontinuous on Ω, while $\omega \mapsto \ell(u^\star,\omega)$ is bounded and continuous on Ω.*

Then there exists a measure $\mathbb{P}^\star \in \mathfrak{P}$ such that $(u^\star, \mathbb{P}^\star)$ is a saddle point of (A.1). In particular, the primal min–max problem admits the same optimal value as its associated dual problem.

To justify the inf-sup interchange used in the reformulation, we first introduce a unified representation for the functions involved in the optimization. In the present work, the functions appearing in the CVaR-based reformulation and the tightening problems are piecewise affine with respect to the decision variable and the uncertainty. Such functions can be expressed in a max-affine form. Therefore, it is sufficient to establish the inf-sup interchange for this class of functions. Specifically, consider the function

$$h(u,w) := \max_{j\in\mathcal{J}} \left\{ a_j^\top u + b_j^\top w + c_j \right\}, \tag{A.2}$$

where $u \in \mathbb{R}^{n_u}$ is the decision variable, $w \in \mathbb{R}^{n_w}$ is the uncertainty realization, $\mathcal{J}$ is a finite index set, and $a_j \in \mathbb{R}^{n_u}$, $b_j \in \mathbb{R}^{n_w}$, and $c_j \in \mathbb{R}$ are given coefficients.

First, we verify the first assumption in our case holds. Since the lifted ambiguity set in (12) is constructed from the same support information and the same nonlinear moment constraints through the lifting variables, it is equivalent to $\mathcal{M}$ in (11) in the sense that every admissible distribution on the lifted set corresponds to an admissible distribution on the original set. Therefore, it is sufficient to verify tightness and weak closedness for the original ambiguity set $\mathcal{M}$ in (11).

Lemma 2. *Let the ambiguity set be defined as*

$$\mathcal{M} := \left\{ \mathbb{P} \in \mathcal{P}(\mathcal{W}) :\ \mathbb{E}_\mathbb{P}[g_{rl}(w)] \le \gamma_{rl},\ \forall r, \forall l \right\}, \tag{A.3}$$

where $\mathcal{W} \subseteq \mathbb{R}^{n_w}$ is the compact support set and $g_{rl} : \mathcal{W} \to \mathbb{R}$ are the nonlinear moment functions. Then $\mathcal{M}$ is tight and closed with respect to weak convergence.

Proof. We first show that $\mathcal{M}$ is tight. By construction, every probability measure $\mathbb{P} \in \mathcal{M}$ is supported on the common compact set $\mathcal{W}$. Hence, for any $\delta > 0$, choosing the compact set $K = \mathcal{W}$ yields

$$\mathbb{P}(K) = \mathbb{P}(\mathcal{W}) = 1 \ge 1 - \delta, \qquad \forall \mathbb{P} \in \mathcal{M}.$$

Therefore, the family $\mathcal{M}$ is tight.

Next, we prove that $\mathcal{M}$ is closed under weak convergence. Let $\{\mathbb{P}_n\} \subseteq \mathcal{M}$ be a sequence such that $\mathbb{P}_n \Rightarrow \mathbb{P}$ weakly for some probability measure $\mathbb{P}$.

We show that $\mathbb{P} \in \mathcal{M}$. Since each $\mathbb{P}_n$ is supported on the closed set $\mathcal{W}$, we have

$$\mathbb{P}_n(\mathcal{W}) = 1, \qquad \forall n.$$

By the Portmanteau theorem, for the closed set $\mathcal{W}$,

$$\mathbb{P}(\mathcal{W}) \ge \limsup_{n\to\infty} \mathbb{P}_n(\mathcal{W}) = 1.$$

Hence $\mathbb{P}(\mathcal{W}) = 1$, which implies that $\mathbb{P}$ is also supported on $\mathcal{W}$.

It remains to verify the nonlinear moment constraints. For any r and l, the function $g_{rl}(w)$ is piecewise affine and continuous on $\mathcal{W}$. Since $\mathcal{W}$ is compact, g_{rl} is also bounded on $\mathcal{W}$. Therefore, $g_{rl} \in C_b(\mathcal{W})$, and weak convergence implies

$$\mathbb{E}_{\mathbb{P}_n}[g_{rl}(w)] \to \mathbb{E}_\mathbb{P}[g_{rl}(w)].$$

Because $\mathbb{P}_n \in \mathcal{M}$, we have

$$\mathbb{E}_{\mathbb{P}_n}[g_{rl}(w)] \le \gamma_{rl}, \qquad \forall n.$$

Then, taking the limits on both side gives

$$\mathbb{E}_\mathbb{P}[g_{rl}(w)] \le \gamma_{rl}, \qquad \forall r, \forall l.$$

Thus $\mathbb{P}$ satisfies all defining constraints of $\mathcal{M}$, and hence $\mathbb{P} \in \mathcal{M}$. Therefore, $\mathcal{M}$ is weakly closed. □

Having established in Lemma 2 that the ambiguity set is tight and weakly closed, we next verify the second assumption of Theorem 1. To this end, it remains to show that the max-affine integrand is measurable and integrable with respect to every admissible probability measure, that its worst-case expectation is finite, and that it is convex with respect to the decision variable. The following lemma provides these properties.

Lemma 3. *Let*

$$h(u,w) := \max_{j\in\mathcal{J}} \{ a_j^\top u + b_j^\top w + c_j \}, \tag{A.4}$$

where $\mathcal{J}$ is a finite index set, $u \in U \subseteq \mathbb{R}^{n_u}$, $w \in \mathcal{W} \subseteq \mathbb{R}^{n_w}$, and $a_j \in \mathbb{R}^{n_u}$, $b_j \in \mathbb{R}^{n_w}$, $c_j \in \mathbb{R}$ are given coefficients. Assume that U is convex and that $\mathcal{W}$ is compact. Then the following statements hold.

For every fixed $u \in U$, the mapping $w \mapsto h(u, w)$ is measurable and bounded on $\mathcal{W}$. Consequently, $h(u, \cdot)$ is integrable with respect to every probability measure $\mathbb{P} \in \mathfrak{P}$ supported on $\mathcal{W}$.

For every fixed $u \in U$, $\sup_{\mathbb{P}\in\mathfrak{P}} \mathbb{E}_{\mathbb{P}}[h(u, w)] < \infty$.

For every fixed $w \in \mathcal{W}$, the function $u \mapsto h(u, w)$ is convex on U.

Therefore, the function $h(u, w)$ satisfies the second assumption of Theorem 1.

Proof. For each $j \in \mathcal{J}$, the function

$$h_j(u, w) := a_j^\top u + b_j^\top w + c_j$$

is affine in (u, w). Since $\mathcal{J}$ is finite, the function $h(u, w)$ in (A.4) is the pointwise maximum of finitely many affine functions.

We first consider the dependence on w. For any fixed $u \in U$, each $h_j(u, \cdot)$ is affine in w, hence continuous and measurable on $\mathcal{W}$. Since the maximum of finitely many measurable functions is measurable, $h(u, \cdot)$ is measurable on $\mathcal{W}$. Moreover, because $\mathcal{W}$ is compact and each $h_j(u, \cdot)$ is continuous, every $h_j(u, \cdot)$ is bounded on $\mathcal{W}$. The finite maximum of bounded functions is again bounded, which shows that $h(u, \cdot)$ is bounded on $\mathcal{W}$. Therefore, for every probability measure $\mathbb{P} \in \mathfrak{P}$ supported on $\mathcal{W}$, the expectation $\mathbb{E}_{\mathbb{P}}[h(u, w)]$ is well defined and finite.

This directly implies $\sup_{\mathbb{P}\in\mathfrak{P}} \mathbb{E}_{\mathbb{P}}[h(u, w)] < \infty$, since each admissible measure is supported on the same compact set $\mathcal{W}$ and $h(u, \cdot)$ is bounded on $\mathcal{W}$.

Next, for any fixed $w \in \mathcal{W}$, each $h_j(\cdot, w)$ is affine in u. Since the pointwise maximum of finitely many affine functions is convex, it follows that $h(\cdot, w)$ is convex on U.

Hence all requirements in the second assumption of Theorem 1 are satisfied. □

We now check the remaining regularity requirement in Theorem 1, namely the boundedness and upper semicontinuity of the integrand with respect to the uncertainty variable. For the max-affine representation considered here, this property follows directly from continuity on the compact support set.

Lemma 4. *Let*

$$h(u, w) := \max_{j\in\mathcal{J}} \{a_j^\top u + b_j^\top w + c_j\}, \tag{A.5}$$

where $\mathcal{J}$ is a finite index set, $u \in U \subseteq \mathbb{R}^{n_u}$, $w \in \mathcal{W} \subseteq \mathbb{R}^{n_w}$, and $a_j \in \mathbb{R}^{n_u}$, $b_j \in \mathbb{R}^{n_w}$, $c_j \in \mathbb{R}$ are given coefficients. Assume that $\mathcal{W}$ is compact, and let $u^\star$ be an optimal solution of (A.1). Then, for every u in a neighborhood of $u^\star$, the mapping $w \mapsto h(u, w)$ is bounded and continuous on $\mathcal{W}$. In particular, $w \mapsto h(u, w)$ is bounded and upper semicontinuous on $\mathcal{W}$ for every such u, and $w \mapsto h(u^\star, w)$ is bounded and continuous on $\mathcal{W}$.

Therefore, the function $h(u, w)$ satisfies the third assumption of Theorem 1.

Proof. Fix any u in a neighborhood of $u^\star$. For each $j \in \mathcal{J}$, define

$$h_j(u, w) := a_j^\top u + b_j^\top w + c_j.$$

For fixed u, the function $h_j(u, \cdot)$ is affine in w, and hence continuous on $\mathcal{W}$. Since $\mathcal{J}$ is finite, the function $h(u, w) = \max_{j\in\mathcal{J}} h_j(u, w)$ is the pointwise maximum of finitely many continuous functions of w. Therefore, $h(u, \cdot)$ is continuous on $\mathcal{W}$.

Because $\mathcal{W}$ is compact, every continuous function on $\mathcal{W}$ is bounded. Hence $h(u, \cdot)$ is bounded on $\mathcal{W}$ for every u in a neighborhood of $u^\star$. Since continuity implies upper semicontinuity, it follows that $h(u, \cdot)$ is bounded and upper semicontinuous on $\mathcal{W}$ for all such u.

In particular, by taking $u = u^\star$, we conclude that $h(u^\star, \cdot)$ is bounded and continuous on $\mathcal{W}$.

Thus the third assumption of Theorem 1 is satisfied. □

We are now in a position to conclude the duality result. All conditions of Theorem 1 are satisfied for the ambiguity set and the max-affine representation considered in this work. It follows that the inf-sup interchange holds without duality gap. In particular,

$$\inf_{u\in U} \sup_{\mathbb{P}\in\mathfrak{P}} \mathbb{E}_{\mathbb{P}}[h(u, w)] = \sup_{\mathbb{P}\in\mathfrak{P}} \inf_{u\in U} \mathbb{E}_{\mathbb{P}}[h(u, w)]. \tag{A.6}$$

This justifies the reformulation steps employed in the main text.

Appendix B. Stability analysis

In this appendix, we provide a brief analysis of the practical stability of the proposed MPC framework. Due to the presence of persistent additive disturbances, asymptotic convergence of the closed-loop system to the origin cannot generally be guaranteed. Therefore, stability is characterized in a practical sense through the boundedness of the long-term average quadratic cost.

Assumption 1. Assume that there exists a terminal set $\mathcal{X}_f$, a linear terminal feedback gain K_k, and a terminal weighting matrix $Q_T \succ 0$ such that the following conditions hold.

First, the terminal set is compatible with the terminal controller and the disturbance set, namely

$$\begin{aligned} &(A_k + B_k K_k)\mathcal{X}_f \oplus \mathcal{W} \subseteq \mathcal{X}_f, \\ &K_k \mathcal{X}_f \subseteq \mathcal{U}, \text{ and } \mathcal{X}_f \subseteq \mathcal{X}. \end{aligned} \tag{B.1}$$

Second, the terminal cost is chosen as

$$V_f(x) = \|x\|_{Q_T}^2, \tag{B.2}$$

where Q_T satisfies the discrete-time Lyapunov equation

$$(A_k + B_k K_k)^\top Q_T (A_k + B_k K_k) - Q_T = -Q - K_k^\top R K_k, \tag{B.3}$$

with $Q \succ 0$ and $R \succ 0$. Finally, the additive disturbance w_k is assumed to have zero mean and covariance Σ_w.

Proposition 3. *Suppose Assumption 1 holds and let $J^*(x_k)$ denote the optimal value function of MPC problem at time step k in (6). Then the closed-loop system is practically stable in the sense that the long-term average quadratic cost is bounded according to*

$$\lim_{T\to\infty} \frac{1}{T} \sum_{k=0}^{T} \mathbb{E}\Big[\|x_k\|_Q^2 + \|u_k\|_R^2 \mid x_0\Big] \le \operatorname{tr}(\Sigma_w Q_T). \tag{B.4}$$

Therefore, although the state and input do not asymptotically converge to the origin due to the persistent stochastic disturbance, the closed-loop trajectories remain bounded in an average quadratic sense.

Proof. Consider the optimal input sequence at time k,

$$\mathbf{u}_k^* = \{u_{0|k}^*, u_{1|k}^*, \dots, u_{N-1|k}^*\},$$

with the corresponding predicted state sequence

$$\mathbf{x}_k^* = \{x_{0|k}^*, x_{1|k}^*, \dots, x_{N|k}^*\},$$

where $x_{0|k}^* = x_k$. The optimal value function at time k is

$$J^*(x_k) = \mathbb{E}\left[\sum_{i=0}^{N-1} \left(\|x_{i|k}^*\|_Q^2 + \|u_{i|k}^*\|_R^2\right) + \|x_{N|k}^*\|_{Q_T}^2 \mid x_k\right]. \tag{B.5}$$

At time $k+1$, we construct a feasible candidate sequence by shifting the optimal solution at time k and appending the terminal control law $u = K_k x$, namely,

$$\tilde{\mathbf{u}}_{k+1} = \{u_{1|k}^*, u_{2|k}^*, \dots, u_{N-1|k}^*, K_k \tilde{x}_{N-1|k+1}\}.$$

Let the corresponding candidate state sequence at time $k+1$ be

$$\tilde{\mathbf{x}}_{k+1} = \{\tilde{x}_{0|k+1}, \tilde{x}_{1|k+1}, \dots, \tilde{x}_{N|k+1}\},$$

with $\tilde{x}_{0|k+1} = x_{k+1}$. Denote $V^*(x_{k+1})$ is the optimal value at $k+1$ conditioned on x_k, and it is upper-bounded by the cost of this feasible candidate sequence, i.e.,

$$V^*(x_{k+1}) \le \tilde{V}(x_{k+1}), \tag{B.6}$$

where

$$\tilde{J}(x_{k+1}) = \mathbb{E}\Big[\sum_{i=0}^{N-2} \Big(\|\tilde{x}_{i|k+1}\|_Q^2 + \|\tilde{u}_{i|k+1}\|_R^2 \Big) + \|\tilde{x}_{N-1|k+1}\|_Q^2 + \|K_k \tilde{x}_{N-1|k+1}\|_R^2 + \|\tilde{x}_{N|k+1}\|_{Q_T}^2 \;\Big|\; x_k \Big]. \tag{B.7}$$

By construction of the shifted sequence, the first $N-1$ stage terms coincide with the tail of the optimal cost at time k. Therefore,

$$\tilde{V}(x_{k+1}) - J^*(x_k) = -\|x_k\|_Q^2 - \|u_k\|_R^2 + \mathbb{E}\Big[\|\tilde{x}_{N-1|k+1}\|_Q^2 + \|K_k \tilde{x}_{N-1|k+1}\|_R^2 + \|\tilde{x}_{N|k+1}\|_{Q_T}^2 - \|x^*_{N|k}\|_{Q_T}^2 \;\Big|\; x_k \Big]. \tag{B.8}$$

Next, under the terminal control law, the terminal state evolves as

$$\tilde{x}_{N|k+1} = (A_k + B_k K_k) x^*_{N|k} + w_k, \tag{B.9}$$

Taking expectation conditioned on x_k, and using the zero-mean assumption on w_k, the cross term vanishes, so that

$$\begin{aligned} \mathbb{E}\Big[\|\tilde{x}_{N|k+1}\|_{Q_T}^2 \mid x_k \Big] &= \mathbb{E}\Big[\|(A + BK_k) x^*_{N|k}\|_{Q_T}^2 \mid x_k \Big] \\ &\quad + \mathbb{E}\big[w_k^\top Q_T w_k \big] \\ &= \mathbb{E}\Big[\|(A + BK_k) x^*_{N|k}\|_{Q_T}^2 \mid x_k \Big] \\ &\quad + \operatorname{tr}\big(\Sigma_w Q_T\big). \end{aligned} \tag{B.10}$$

Substituting this into (B.8) and using the Lyapunov equation gives

$$\begin{aligned} &\mathbb{E}\Big[\|\tilde{x}_{N-1|k+1}\|_Q^2 + \|K_k \tilde{x}_{N-1|k+1}\|_R^2 + \|\tilde{x}_{N|k+1}\|_{Q_T}^2 \\ &\quad - \|x^*_{N|k}\|_{Q_T}^2 \;\Big|\; x_k \Big] \\ &= \mathbb{E}\Big[x^{*\top}_{N|k} (Q + K_k^\top R K_k + (A_k + B_k K_k)^\top Q_T (A_k + B_k K_k) \\ &\quad - Q_T) x^*_{N|k} \;\Big|\; x_k \Big] + \operatorname{tr}(\Sigma_w Q_T) \\ &= \operatorname{tr}(\Sigma_w Q_T). \end{aligned} \tag{B.11}$$

Therefore,

$$\mathbb{E}\big[\tilde{J}(x_{k+1}) - J^*(x_k) \mid x_k \big] \le \operatorname{tr}(\Sigma_w Q_T) - \|x_k\|_Q^2 - \|u_k\|_R^2. \tag{B.12}$$

Since $J^*(x_{k+1}) \le \tilde{J}(x_{k+1})$, it follows that

$$\mathbb{E}\big[J^*(x_{k+1}) - J^*(x_k) \mid x_k \big] \le \operatorname{tr}(\Sigma_w Q_T) - \|x_k\|_Q^2 - \|u_k\|_R^2. \tag{B.13}$$

This inequality shows that the expected increase of the optimal cost is upper-bounded by a constant term caused by the disturbance, minus the instantaneous quadratic stage cost.

Next, take the expectation of both sides conditioned on the initial state x_0. Then

$$\begin{aligned} \mathbb{E}\big[J^*(x_{k+1}) \mid x_0 \big] - \mathbb{E}\big[J^*(x_k) \mid x_0 \big] &\le \operatorname{tr}(\Sigma_w Q_T) \\ &\quad - \mathbb{E}\Big[\|x_k\|_Q^2 + \|u_k\|_R^2 \mid x_0 \Big]. \end{aligned} \tag{B.14}$$

Summing the above inequality from $k = 0$ to $k = T$ yields

$$\begin{aligned} \mathbb{E}\big[J^*(x_{T+1}) \mid x_0 \big] - J^*(x_0) &\le (T+1) \operatorname{tr}(\Sigma_w Q_T) \\ &\quad - \sum_{k=0}^{T} \mathbb{E}\Big[\|x_k\|_Q^2 + \|u_k\|_R^2 \mid x_0 \Big]. \end{aligned} \tag{B.15}$$

Rearranging gives

$$\begin{aligned} \sum_{k=0}^{T} \mathbb{E}\Big[\|x_k\|_Q^2 + \|u_k\|_R^2 \mid x_0 \Big] &\le (T+1) \operatorname{tr}(\Sigma_w Q_T) \\ &\quad + J^*(x_0) - \mathbb{E}\big[J^*(x_{T+1}) \mid x_0 \big]. \end{aligned} \tag{B.16}$$

Dividing both sides by T, we obtain

$$\begin{aligned} \frac{1}{T} \sum_{k=0}^{T} \mathbb{E}\Big[\|x_k\|_Q^2 + \|u_k\|_R^2 \mid x_0 \Big] &\le \frac{T+1}{T} \operatorname{tr}(\Sigma_w Q_T) \\ &\quad + \frac{J^*(x_0) - \mathbb{E}\big[J^*(x_{T+1}) \mid x_0 \big]}{T}. \end{aligned} \tag{B.17}$$

Since the optimal value function is nonnegative and $J^*(x_0)$ is finite, we can obtain

$$\lim_{T \to \infty} \frac{1}{T} \sum_{k=0}^{T} \mathbb{E}\Big[\|x_k\|_Q^2 + \|u_k\|_R^2 \mid x_0 \Big] \le \operatorname{tr}(\Sigma_w Q_T). \tag{B.18}$$

This proves that the closed-loop system is practically stable in the sense of bounded long-term average quadratic cost. □

References

Albin, T., Zweigel, R., Heßeler, F., & Abel, D. (2013). A 2-stage MPC approach for the cycle-to-cycle dynamics of GCAI (gasoline controlled auto ignition). *IFAC Proceedings Volumes*, *46*(21), 183–188.

Albin, T., Zweigel, R., Heßeler, F., Morcinkowski, B., Brassat, A., & Abel, D. (2012). Controlling GCAI (gasoline controlled auto ignition) in an extended operating map. *IFAC Proceedings Volumes*, *45*(30), 466–473.

Batina, I. (2004). Model predictive control for stochastic systems by randomized algorithms. Ph.D. thesis, Technische Universiteit Eindhoven, Eindhoven, The Netherlands. ISBN: 90-386-0812-8.

Bedei, J., Koch, L., Badalian, K., Winkler, A., Schaber, P., & Andert, J. (2025). Safe reinforcement learning for real-world engine control. arXiv:2501.16613.

Bedei, J., Oberlies, M., Schaber, P., Gordon, D., Nuss, E., Li, L., & Andert, J. (2023). Dynamic measurement with in-cycle process excitation of HCCI combustion: The key to handle complexity of data-driven control? *International Journal of Engine Research*, *24*(3), 1155–1174.

Bemporad, A., Borrelli, F., Morari, M. et al. (2002). Model predictive control based on linear programming - the explicit solution. *IEEE Transactions on Automatic Control*, *47*(12), 1974–1985.

Bemporad, A., & Morari, M. (2007). Robust model predictive control: A survey. In *Robustness in identification and control* (pp. 207–226). Springer.

Bengtsson, J., Strandh, P., Johansson, R., Tunestål, P., & Johansson, B. (2006a). Hybrid control of homogeneous charge compression ignition (HCCI) engine dynamics. *International Journal of Control*, *79*(05), 422–448.

Bengtsson, J., Strandh, P., Johansson, R., Tunestal, P., & Johansson, B. (2006b). Model predictive control of homogeneous charge compression ignition (HCCI) engine dynamics. In *2006 IEEE conference on computer aided control system design* (pp. 1675–1680). IEEE.

Bertsekas, D. (2012). Dynamic programming and optimal control: Volume I (vol. 4). Athena scientific.

Blackmore, L., Ono, M., Bektassov, A., & Williams, B. C. (2010). A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Transactions on Robotics*, *26*(3), 502–517.

Calafiore, G. C., & Ghaoui, L. E. (2006). On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, *130*, 1–22.

Chen, H., Paoletti, N., Smolka, S. A., & Lin, S. (2021). MPC-guided imitation learning of Bayesian neural network policies for the artificial pancreas. In *2021 60th IEEE conference on decision and control (CDC)* (pp. 2525–2532). IEEE.

Chen, X., Basler, M., Stemmler, M., & Abel, D. (2023). Gasoline controlled auto-ignition with learning-based uncertainty using stochastic model predictive control. In *2023 62nd IEEE conference on decision and control (CDC)* (pp. 6328–6335). IEEE.

De Schutter, J., Nuss, E., Baumgärtner, K., Abel, D., & Diehl, M. (2022). Nonlinear MPC and MHE of a gasoline controlled auto-ignition engine based on reduced-order differential equation models. In *2022 IEEE 61st conference on decision and control (CDC)* (pp. 2411–2416). IEEE.

Delage, E., & Ye, Y. (2010). Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, *58*(3), 595–612.

Ebrahimi, K., & Koch, C. B. (2018). Real-time control of HCCI engine using model predictive control. In *2018 Annual american control conference (ACC)* (pp. 1622–1628). IEEE.

Frison, G., & Diehl, M. (2020). HPIPM: A high-performance quadratic programming framework for model predictive control. *IFAC-PapersOnLine*, *53*(2), 6563–6569.

Gao, R., & Kleywegt, A. (2023). Distributionally robust stochastic optimization with Wasserstein distance. *Mathematics of Operations Research*, *48*(2), 603–655.

Ghezzi, A., Hoffman, J., Frey, J., Boedecker, J., & Diehl, M. (2023). Imitation learning from nonlinear MPC via the exact q-loss and its Gauss-Newton approximation. In *2023 62nd IEEE conference on decision and control (CDC)* (pp. 4766–4771). IEEE.

Gordon, D. C., Winkler, A., Bedei, J., Schaber, P., Pischinger, S., Andert, J., & Koch, C. R. (2024). Introducing a deep neural network-based model predictive control framework for rapid controller implementation. In *2024 American control conference (ACC)* (pp. 5232–5237). IEEE.

Gros, S., & Zanon, M. (2019). Data-driven economic NMPC using reinforcement learning. *IEEE Transactions on Automatic Control*, *65*(2), 636–648.

Gros, S., & Zanon, M. (2021). Reinforcement learning based on MPC and the stochastic policy gradient method. In *2021 American control conference (ACC)* (pp. 1947–1952). IEEE.

Hussein, A., Gaber, M. M., Elyan, E., & Jayne, C. (2017). Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, *50*(2), 1–35.

Janakiraman, V. M., Nguyen, X., & Assanis, D. (2015). Nonlinear model predictive control of a gasoline HCCI engine using extreme learning machines. arXiv:1501.03969.

Johansen, T. A. (2004). Approximate explicit receding horizon control of constrained nonlinear systems. *Automatica*, *40*(2), 293–300.

Koopmans, L., Ogink, R., & Denbratt, I. (2003). Direct gasoline injection in the negative valve overlap of a homogeneous charge compression ignition engine. *SAE Transactions*, *112*, 1365–1376.

Kouvaritakis, B., Cannon, M., Raković, S. V., & Cheng, Q. (2010). Explicit use of probabilistic distributions in linear predictive control. *Automatica*, *46*(10), 1719–1724.

Kulzer, A., Christ, A., Rauscher, M., Sauer, C., Würfel, G., & Blank, T. (2006). Thermodynamic analysis and benchmark of various gasoline combustion concepts. Technical Report SAE Technical Paper.

Lall, S., & Glover, K. (1994). A game theoretic approach to moving horizon control. In D. W. Clarke (Ed.), *Advances in Model-Based Predictive Control* (pp. 131–144). Oxford University Press, Oxford, UK.

Lu, S., Lee, J. H., & You, F. (2020). Soft-constrained model predictive control based on data-driven distributionally robust optimization. *AIChE Journal*, *66*(10), e16546.

Magni, L., Nijmeijer, H., & Van Der Schaft, A. J. (2001). A receding–horizon approach to the nonlinear H_∞ control problem. *Automatica*, *37*(3), 429–435.

Mesbah, A., Kolmanovsky, I. V., & Di Cairano, S. (2019). Stochastic model predictive control. *Handbook of Model Predictive Control*, (pp. 75–97).

Mesbah, A., Streif, S., Findeisen, R., & Braatz, R. D. (2014). Stochastic nonlinear model predictive control with probabilistic constraints. In *2014 American control conference* (pp. 2413–2419). IEEE.

Nuss, E. (2024). Modeling and nonlinear model predictive control for gasoline-controlled auto-ignition. Ph.D. thesis. Dissertation, Rheinisch-Westfälische Technische Hochschule Aachen, 2024.

Nuss, E., Ritter, D., Wick, M., Andert, J., Abel, D., & Albin, T. (2019). Reduced order modeling for multi-scale control of low temperature combustion engines. In *Active flow and combustion control 2018: Papers contributed to the conference "active flow and combustion control 2018", September 19–21, 2018, Berlin, Germany* (pp. 167–181). Springer.

Nuss, E., Wick, M., Andert, J., De Schutter, J., Diehl, M., Abel, D., & Albin, T. (2019). Nonlinear model predictive control of a discrete-cycle gasoline-controlled auto ignition engine model: Simulative analysis. *International Journal of Engine Research*, *20*(10), 1025–1036.

Oakley, A., Zhao, H., Ladommatos, N., & Ma, T. (2001). Experimental studies on controlled auto-ignition (CAI) combustion of gasoline in a 4-stroke engine. *SAE Transactions*, *110*, 1062–1075.

Oldewurtel, F., Jones, C. N., & Morari, M. (2008). A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback. In *2008 47th IEEE conference on decision and control* (pp. 4731–4736). IEEE.

Popescu, I. (2007). Robust mean-covariance solutions for stochastic optimization. *Operations Research*, *55*(1), 98–112.

Rahimian, H., & Mehrotra, S. (2019). Distributionally robust optimization: A review. arXiv:1908.05659.

Reske, A., Carius, J., Ma, Y., Farshidian, F., & Hutter, M. (2021). Imitation learning from MPC for quadrupedal multi-gait control. In *2021 IEEE International conference on robotics and automation (ICRA)* (pp. 5014–5020). IEEE.

Ritter, D., Andert, J., Abel, D., & Albin, T. (2018). Model-based control of gasoline-controlled auto-ignition. *International Journal of Engine Research*, *19*(2), 189–201.

Rockafellar, R. T., Uryasev, S. et al. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*, 21–42.

Saxena, S., & Bedoya, I. D. (2013). Fundamental phenomena affecting low temperature combustion and HCCI engines, high load limits and strategies for extending these limits. *Progress in Energy and Combustion Science*, *39*(5), 457–488.

Scarf, H. E., Arrow, K., & Karlin, S. (1957). A min-max solution of an inventory problem. Rand Corporation Santa Monica.

Shang, C., & You, F. (2018). Distributionally robust optimization for planning and scheduling under uncertainty. *Computers & Chemical Engineering*, *110*, 53–68.

Shapiro, A., & Kleywegt, A. (2002). Minimax analysis of stochastic problems. *Optimization Methods and Software*, *17*(3), 523–542.

Torabi, F., Warnell, G., & Stone, P. (2018). Behavioral cloning from observation. arXiv:1805.01954.

van Hessem, D. H., & Bosgra, O. H. (2003). A full solution to the constrained stochastic closed-loop MPC problem via state and innovations feedback and its receding horizon implementation. In *42nd IEEE international conference on decision and control (IEEE cat. no. 03CH37475)* (pp. 929–934). IEEE (*vol. 1*).

Wang, Z., Glynn, P. W., & Ye, Y. (2016). Likelihood robust optimization for data-driven problems. *Computational Management Science*, *13*, 241–261.

Wei, H., Zhu, T., Shu, G., Tan, L., & Wang, Y. (2012). Gasoline engine exhaust gas recirculation–a review. *Applied Energy*, *99*, 534–544.

Wiesemann, W., Kuhn, D., & Sim, M. (2014). Distributionally robust convex optimization. *Operations Research*, *62*(6), 1358–1376.

Zhao, H. (2007). HCCI and CAI engines for the automotive industry. Elsevier.

Zymler, S., Kuhn, D., & Rustem, B. (2013). Distributionally robust joint chance constraints with second-order moment information. *Mathematical Programming*, *137*, 167–198.