Control Engineering Practice 170 (2026) 106834

Contents lists available at ScienceDirect

Control Engineering Practice

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

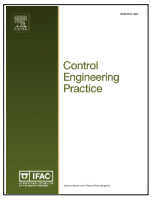

Experimental demonstration of safe and automated in-cylinder pressure shaping using constrained extremum seeking

Mats Versmissen [a,*], Maarten Vlaswinkel [a], Frank Willems [a,b]

[a] Eindhoven University of Technology, Department of Mechanical Engineering, Control Systems Technology, Eindhoven, 5600 MB, The Netherlands
[b] TNO Mobility & built environment, Powertrains Department, Helmond, 5700 AT, The Netherlands

ARTICLE INFO

Keywords:
Constrained extremum seeking
Principal component decomposition
Control calibration
Dual-fuel engines

ABSTRACT

To support decarbonization of transport, research is focusing on advanced and highly efficient combustion concepts running on low-carbon fuels. Optimizing the performance of these complex engines over a wide range of operating conditions results in exploded control calibration efforts for traditional calibration methods. In this paper, an automated engine calibration framework based on constrained Extremum Seeking (ES) control is proposed to effectively reduce calibration times. Contrary to previous research that optimizes an efficiency metric, the proposed ES algorithm directly shapes the entire in-cylinder pressure trace by decomposing it into Principal Components (PCs) and controlling the associated weights towards their optimal reference, defined by an Ideal Thermodynamic cycle. This method finds an optimal trade-off between controllability of the PC weights and optimality, while explicitly addressing combustion safety constraints using a novel gradient based projection method for the PC weights. The proposed ES controller was successfully implemented for fuel path calibration on a single cylinder engine running in dual-fuel (diesel-gasoline) mode. Besides demonstrating optimality and convergence of the ES concept within 4 min, the handling of an arbitrary constraint on maximal in-cylinder pressure is validated. Despite the under-actuated nature of the control problem for the PC weights, the resulting ES optimal operating point is close to the actual maximal thermal efficiency point. This work highlights the practical viability of automated calibration methods through a model-free in-cylinder pressure shaping approach.

1. Introduction

Meeting the objective of a net zero CO_2 emission scenario by 2050 requires a stepwise reduction in global transport emissions, targeting a 35% reduction for long-haul transport by 2030, while accounting for the increasing demand in this sector (IEA, 2023). Electrification is the primary solution for passenger vehicles, whereas low-carbon fuels can accelerate decarbonization in the heavy-duty sector, including long-haul road transport, maritime, aviation, and mining. This is especially true as challenges regarding electric charging infrastructure persist. In addition to low-carbon fuels, low temperature combustion (LTC) concepts promise higher thermal efficiencies and lower engine-out emissions compared to traditional compression-ignition (CI) and spark-ignition (SI) engines (Agarwal et al., 2017). These LTC concepts can enhance current de-carbonization strategies, particularly when used in combination with clean alternative fuels such as hydrogen or natural gas (Karimi et al., 2021).

Pre-mixed combustion concepts rely on controlled auto-ignition and are very sensitive to operating conditions, such as differences in ambient conditions or varying fuel types or qualities. Additionally, the overall engine complexity increases and the requirement for safe and stable operation under a wide operating range must be guaranteed. Therefore, the time and costs associated with robust engine calibration explodes when mapping the complete operating range. Auto-calibration can greatly reduce the manual effort required, minimizing the time spent on the testbench and adapting to real-world conditions using on-line learning. Willems, 2018 proposes a cylinder pressure-based control approach that utilizes in-cylinder pressure data for real-time combustion control. This approach will gain traction as the commercial viability of in-cylinder pressure sensors increases.

Extremum Seeking (ES) is a promising approach to automated engine calibration and has demonstrated effective results in the automotive field for fuel efficiency optimization. Popovic et al. (2003) optimizes the Brake Specific Fuel Consumption (BSFC) of a variable cam timing SI engine while using combustion stability as a constraint. This stability region is mapped beforehand, failing to capture varying operating conditions. Demonstrated on a heavy-duty CI engine, van der Weijst et al. (2019) minimizes BSFC by adapting feedforward references while

* Corresponding author.
E-mail addresses: m.versmissen@tue.nl (M. Versmissen), m.g.vlaswinkel@tue.nl (M. Vlaswinkel), f.p.t.willems@tue.nl (F. Willems).

https://doi.org/10.1016/j.conengprac.2026.106834
Received 2 November 2025; Received in revised form 16 January 2026; Accepted 3 February 2026
Available online 11 February 2026

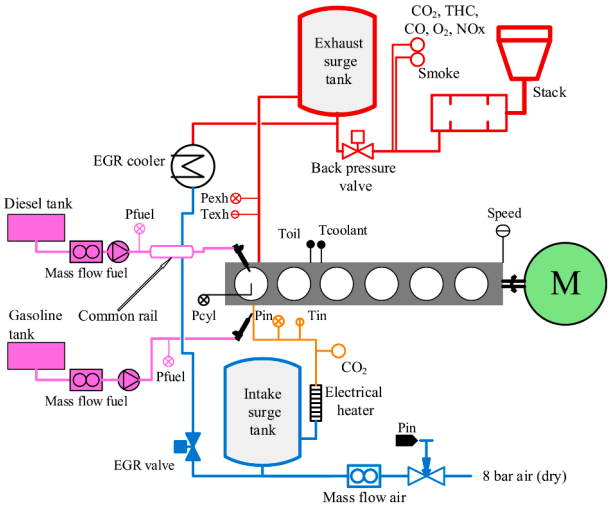

Fig. 1. Experimental SCE equipped with custom port-fuel injection system. This setup is similar as the one used in Vlaswinkel and Willems (2023).

ensuring that constraints on engine out NO_x, maximum in-cylinder pressure, and Exhaust Gas Recirculation (EGR) valve position are satisfied. ES has been successfully applied to a Homogeneous Charge Compression Ignition (HCCI) engine, where the gains of a PID controller were optimized using ES, and it was also used to optimize the injection timing (Killingsworth et al., 2009). The last two methods do not incorporate combustion safety parameters as constraints for the ES algorithm.

In this paper, the constrained ES framework for engine control calibration is extended for in-cylinder pressure shaping using the Principal Component (PC)-based method introduced in Vlaswinkel and Willems (2023). Following this approach, we not only maximize the Gross Indicated Efficiency (GIE), but also explicitly deal with safety constraints (i.e. peak in-cylinder pressure and maximal in-cylinder pressure rise rate) using fuel path actuators. The main contributions of this work are summarized as follows; i) shaping of the entire in-cylinder pressure trace by decomposition of the in-cylinder pressure trace into multiple Principal Components (PCs) and controlling of these PC weights towards their optimal values, defined by the Ideal Thermodynamic cycle (ITC); ii) explicit handling of combustion safety constraints by a novel PC weights-based method that effectively handles these constraints through gradient projection onto the constrained multi-dimensional hyperplanes; iii) successful demonstration of the proposed ES algorithm on a single cylinder engine set-up operating in dual-fuel (diesel-gasoline) mode.

This paper is organized as follows. First, an overview of the experimental testbench is provided. Next, the automated engine calibration approach is presented. Then, the developed ES algorithm is presented and validated on the engine testbench, and lastly, the experimental results are discussed.

2. System description

This section provides an overview of the single cylinder engine (SCE) setup from Zero Emission Lab at the Eindhoven University of Technology. The SCE was used to demonstrate the potential of the constrained Extremum Seeking framework.

2.1. Single cylinder setup

The SCE setup is a modified heavy-duty PACCAR MX13 engine. Only the first cylinder supports fuel injection and is equipped with a common rail diesel injection (DI) system and a custom port-fuel injection (PFI) system for gasoline (See Fig. 1).

To monitor the in-cylinder pressure (p_{cyl}) during engine operation, a pressure transducer (6125C) and a charge amplifier (5018) from Kistler

Table 1
SCE engine specifications.

Parameter	Value or Type	Unit
Displacement volume V_d	2.15	[L]
Compression Ratio (r_c)	17.2:1	[–]
Intake Valve Closed	173	[CADbTDC]
Intake Valve Open	−357	[CADbTDC]
Exhaust Valve Closed	359	[CADbTDC]
Exhaust Valve Open	−146	[CADbTDC]
EGR Valve	0	[%open]
Engine Speed	1200	[RPM]
DI Fuel	Diesel	[–]
DI Rail Pressure	1200	[bar]
PFI Fuel	Gasoline	[–]
PFI Rail Pressure	5	[bar]

are used. The SCE is attached to an electric machine that generates the required torque to maintain a fixed engine speed. The crankshaft is equipped with a crank encoder (RSI 503) from Leine Linde, which has a resolution of 0.2 CAD. The main engine specifications of this setup are shown in Table 1 and more details can be found in Willems (2020).

2.2. Real-time combustion control system

The original engine control system is replaced with two separate rapid prototyping systems. A Compact RIO (National Instruments Corp.) sets both the intake manifold temperature T_{im} through an inlet heater; and the intake manifold pressure p_{im} via a pneumatic intake valve, to a fixed actuator position. Common rail pressure control is handled by the Compact RIO as well. In addition to this, a Speedgoat Real-Time Target Machine (Speedgoat GmbH, 2018) is integrated for on-line optimization and next-cycle fuel path control in Simulink (R2021b).

The Speedgoat consists of a Kintex 7 Field Programmable Gate Array (FPGA), running at 20 ns task rate, and an Intel Core i7 3.5 GHz CPU with four cores. The FPGA handles fast parallel processing of the in-cylinder pressure and encoder signals and computes the Internal Mean Effective Pressure (IMEP) used in a next-cycle control architecture. Safety monitoring and injector driver actuation logic runs on the FPGA as well. The next-cycle controller and the Extremum Seeking controller are implemented on the CPU in Real-Time (RT) and use feedback from the FPGA. From a requested mass flow and injection timing, the Speedgoat generates injection commands that specify the start of actuation (SOA) and duration of actuation (DOA) and communicates these to the injector drivers:

1. A RapidPro Power module (dSpace GmbH, 2017) controls the diesel injector to meet timing and duration requirements set by the Speedgoat.
2. A custom driver controls the port fuel injector, ensuring that the injection timing and duration meet the requirements set by the Speedgoat.

Communication between the FPGA and CPU is realized by a PCIe interface and combustion data is visualized on the CPU at a frequency of 100 Hz. This setup is used for automated steady-state engine calibration, which is discussed in the following sections.

3. Engine control problem

3.1. Engine calibration objectives

An Ideal Thermodynamic Cycle (ITC) describes the optimal engine cycle under perfect conditions, such as zero heat loss and an infinitely short combustion duration (Heywood, 1988). While achieving such an ideal cycle in an engine is not feasible, the primary goal of engine calibration is to find the required control inputs such that the actual combustion process meets the power demand and closely approximates this

ITC, thereby maximizing the Gross Indicated Efficiency (GIE):

$$\mathrm{GIE}(u) = \frac{\mathrm{IMEP}_g(u)}{Q_{fuel}}$$

where u are the controlled inputs and IMEP_g as the Gross Indicated Mean Effective Pressure:

$$\mathrm{IMEP}_g(u) = \frac{1}{V_d}\int_{\Theta} p_{cyl}(\theta, u)\, dV(\theta)$$

with θ the crank angle, $\Theta = [-180, 180]$ CAD and $V(\theta)$, the in-cylinder volume. Additionally, the calibration process should comply with emissions and combustion stability targets, as well as combustion safety constraints, which include upper limits on maximum pressure ($p_{\max}$) and maximum pressure rise rate ($dp_{\max}$).

For dual-fuel engines, three controlled variables which relate to the IMEP_g and GIE are available for fuel path calibration. It is assumed that the air path remains constant, as the SCE is not equipped with a turbocharger. First, there is the total amount of fuel power injected during a single combustion cycle:

$$Q_{fuel} = m_{\mathrm{DI}}\mathrm{LHV}_{\mathrm{DI}} + m_{\mathrm{PFI}}\mathrm{LHV}_{\mathrm{PFI}}$$

with $m_{\mathrm{DI}}, m_{\mathrm{PFI}} \in [\frac{mg}{inj}]$, the injected fuel mass, and $\mathrm{LHV}_{\mathrm{DI}}, \mathrm{LHV}_{\mathrm{PFI}} \in [\frac{J}{mg}]$, the lower heating values for high and low-reactivity fuels, respectively. Next, there is the energy-based blend ratio BR:

$$BR = \frac{m_{\mathrm{PFI}}\mathrm{LHV}_{\mathrm{PFI}}}{Q_{fuel}}$$

as a measure of fuel reactivity and $\mathrm{SOI}_{\mathrm{DI}} \in [\mathrm{CADbTDC}]$, the injection timing for high reactivity fuel is also considered. The controlled variables in this work are:

$$u = [Q_{fuel}, BR, \mathrm{SOI}_{\mathrm{DI}}]^\top \in \mathbb{R}^3 \tag{1}$$

3.2. In-cylinder pressure shaping

Vlaswinkel and Willems (2023) rewrote this calibration task, using an Otto cycle as the optimal reference for the in-cylinder pressure trace $p_{cyl}(\theta, u)$. The goal is to minimize J:

$$J = \left(\int_{\Theta} (p_{cyl}(\theta, u) - p_{\mathrm{ITC}}(\theta, Q_{fuel}))\, dV(\theta)\right)^2 \tag{2}$$

representing the difference between the measured pressure trace $p_{cyl}(\theta, u)$ and the ideal reference in-cylinder pressure $p_{\mathrm{ITC}}(\theta, Q_{fuel})$, which aligns with a maximization of the GIE.

This method is based on the idea that an arbitrary in-cylinder pressure trace can be decomposed into multiple Principal Components (PC) as follows:

$$p_{cyl}(\theta, u) = w(u)^\top f(\theta) + f_{mp}(\theta, p_{im}) \tag{3}$$

where $w(u) \in \mathbb{R}^{n_{\mathrm{PC}}}$ the weights associated with each PC as function of the controlled fueling parameters, n_{PC} the number of PCs, and $f(\theta) \in \mathbb{R}^{n_{\mathrm{PC}}}$ the static PC traces. $f_{mp}(\theta, p_{im})$ is the underlying motored pressure, modeled as an adiabatic compression and is a function of the intake manifold pressure p_{im} (Vlaswinkel and Willems, 2024). Using this Principal Component Decomposition (PCD), Eq. (2) can be written as:

$$J = (w(u) - w_{\mathrm{ITC}}(Q_{fuel}))^\top Z_1 (w(u) - w_{\mathrm{ITC}}(Q_{fuel})) \tag{4}$$

which is a composite cost function $J(w(u)) : \mathbb{R}^3 \to \mathbb{R}^{n_{\mathrm{PC}}} \to \mathbb{R}$, with an intermediate mapping function $w(u) : \mathbb{R}^3 \to \mathbb{R}^{n_{\mathrm{PC}}}$, from the fuel inputs to the PC weights and Z_1 given as:

$$Z_1 = \iint_{\Theta} f(\theta_1) f^\top(\theta_2)\, dV(\theta_1)\, dV(\theta_2)$$

3.3. Optimal control problem

The combustion safety constraints, $p_{\max}$ and $dp_{\max}$ can be expressed as linear inequality constraints using the decomposition from (3). The resulting calibration task can be expressed as follows:

$$u^* = \min_{u \in \mathcal{U}} \quad J(w(u)) \tag{5a}$$

$$\text{s. t.} \quad \mathrm{IMEP}_g(u) = \mathrm{IMEP}_{\mathrm{ref}} \tag{5b}$$

$$w(u)^\top f(\theta) + f_{\mathrm{mp}}(\theta) \le p_{max,ub} \tag{5c}$$

$$w(u)^\top \frac{df(\theta)}{d\theta} + \frac{df_{\mathrm{mp}}(\theta)}{d\theta} \le dp_{max,ub} \tag{5d}$$

where $\mathcal{U}$ is the feasible search space for the fuel path actuators from (1), and $p_{max,ub}$ and $dp_{max,ub}$ represent the maximum pressure and the maximum pressure rise rate, respectively. The following section proposes an Extremum Seeking framework capable of minimizing the given cost function while satisfying the aforementioned constraints.

4. Constrained extremum seeking control

To optimize the fuel path of the dual-fuel engine, a dither-based constrained Extremum Seeking controller (ESC) is proposed, as shown in Fig. 2. ESC is a model-free gradient descent optimization algorithm that does not require prior knowledge of the input-output relationship. The only requirements for the implementation of the algorithm are; an initial safe setpoint can be chosen, a (local) minimum must exist, and there is the ability to measure the cost function (Scheinker, 2024).

4.1. Next-cycle IMEP controller

We focus on optimizing the fuel path for a specified driver demand ($\mathrm{IMEP}_{\mathrm{ref}}$). To satisfy this constraint (5b), a next-cycle closed-loop controller for Q_{fuel} is implemented as follows:

$$Q_{fuel,k+1} = Q_{fuel,k} + K_{\mathrm{IMEP}}\,(\mathrm{IMEP}_{\mathrm{ref}} - \mathrm{IMEP}_{g,k})$$

where 'k' denotes the current combustion cycle and K_{IMEP} the control gain. This next-cycle controller is inherently fast, responding in approximately 10 combustion cycles to local changes in driver demand.

4.2. Extremum seeking optimizer

For the ESC, we define $\tilde{u}_0 = [BR_0, \mathrm{SOI}_{\mathrm{DI},0}]^\top \subset u$ as the reference inputs to be optimized by the algorithm. The ES optimizer is given in a next-cycle manner as:

$$\tilde{u}_{0,k+1} = \tilde{u}_{0,k} + K_{ES}\left.\frac{\partial J}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}}$$

where $K_{ES} = diag([K_{BR}, K_{\mathrm{SOI}}])$ is the control gain matrix and $\left.\frac{\partial J}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}} \in \mathbb{R}^2$ the control gradient that drives the references towards a minimal cost J. The main components of the ES framework are gradient estimation and constraint handling, which are discussed in further detail.

4.3. Gradient estimation

Classic dither-based ES relies on accurately estimating the gradient of cost J with respect to the system input references $\tilde{u}_0$ by perturbing these static references with a sinusoidal dither signal:

$$\tilde{u}_k = \tilde{u}_{0,k} + d_k$$

where

$$d_k = \begin{bmatrix} a_{BR}\sin(\omega_{BR} t_k) \\ a_{\mathrm{SOI}}\sin(\omega_{\mathrm{SOI}} t_k) \end{bmatrix}$$

with a_{BR}, a_{SOI} the dither amplitudes and $\omega_{BR}, \omega_{\mathrm{SOI}}$ the dither frequencies for the BR and $\mathrm{SOI}_{\mathrm{DI}}$ inputs, respectively. As shown in (4), a composite cost function was obtained $J(w(\tilde{u})) : \mathbb{R}^2 \to \mathbb{R}^{n_{\mathrm{PC}}} \to \mathbb{R}$, with an intermediate mapping function: $w(\tilde{u}) : \mathbb{R}^2 \to \mathbb{R}^{n_{\mathrm{PC}}}$. This intermediate mapping can be leveraged to satisfy the combustion safety constraints (5c)

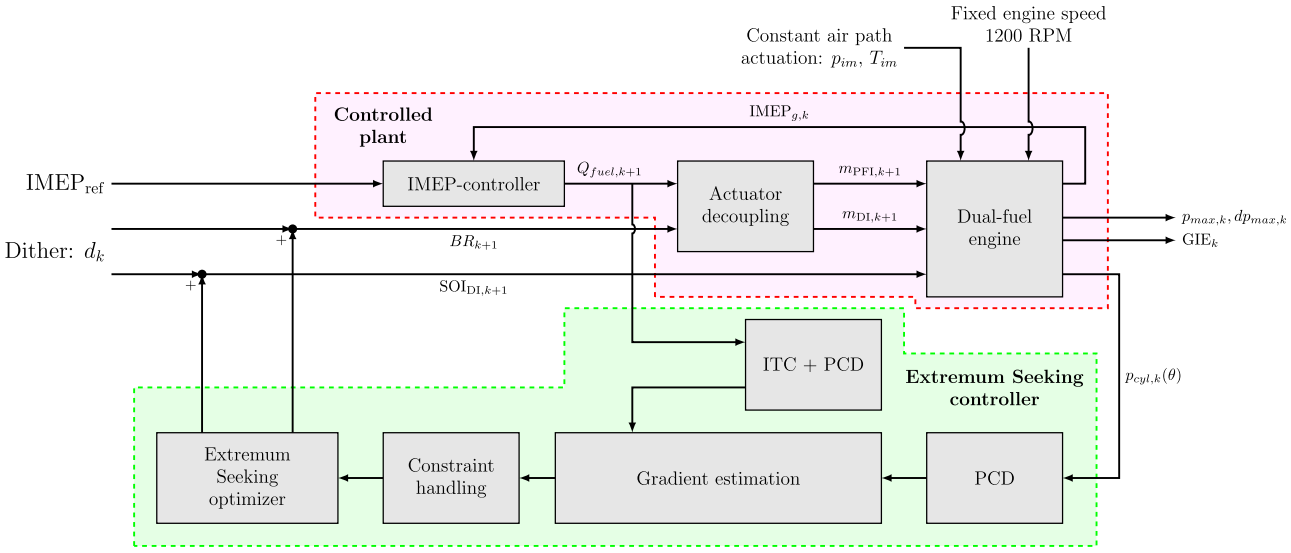

Fig. 2. Control architecture overview. The Controlled plant represents the engine with the coupled control variables (Q_{fuel}, $BR \rightarrow m_{\text{DI}}, m_{\text{PFI}}$). The Extremum Seeking controller is highlighted and the individual components structure is shown.

and (5d), as will be seen later. Now, consider the composite cost function[1]:

$$J(w(\tilde{u})) = (w(\tilde{u}) - w_{\text{ITC}})^\top Z_1 (w(\tilde{u}) - w_{\text{ITC}})$$

where the optimal control gradient $\frac{\partial J}{\partial \tilde{u}}$, evaluated around $\tilde{u}_{0,k}$, can be obtained by applying the chain rule as follows:

$$\left.\frac{\partial J}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}} = \left.\frac{\partial J}{\partial w}\right|^\top_{\tilde{u}_{0,k}} \left.\frac{\partial w}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}} \in \mathbb{R}^2 \tag{6}$$

4.3.1. Algebraic gradient

The first part from the optimal control gradient,

$$\left.\frac{\partial J}{\partial w}\right|_{\tilde{u}_{0,k}} = 2(w(\tilde{u}_{0,k}) - w_{\text{ITC},k})^\top Z_1 \in \mathbb{R}^{n_{\text{PC}}} \tag{7}$$

where $w(\tilde{u}_{0,k})$ the PC weight for a fixed reference input $\tilde{u}_0$ and represents the algebraic gradient that defines the optimal direction for the PC weights to shift in order to minimize J. The second part,

$$\left.\frac{\partial w}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}} \in \mathbb{R}^{n_{\text{PC}} \times 2}$$

represents the PC mapping gradient around $\tilde{u}_0$, which defines the controllability from the dithered inputs for the individual PC weights.

4.3.2. Kalman filter for PC mapping gradient

The PC weight $w(\tilde{u}_{0,k})$ and the mapping gradients around $\tilde{u}_0$ are estimated using a Kalman filter. The Kalman filter (KF) is a computationally efficient state estimator that explicitly handles cycle-to-cycle combustion data (streams of data). Therefore, the filter can be implemented on the Speedgoat Real-Time Target Machine without any implications and provide a new state estimate every combustion cycle. Welch and Bishop, 2006. The KF also performs well when the system suffers from additive Gaussian noise, which these dual-fuel combustion concepts inherently suffer from Vlaswinkel and Willems, 2024. In this study, the PC weights are approximated by a first order (N = 1) Taylor expansion, similar to van der Weijst et al., 2017. Here, the approach is demonstrated for a single PC weight.

[1] Dependency of w_{ITC} on Q_{fuel} is omitted for more compact writing

Consider the first order Taylor expansion around $\tilde{u}_{0,k}$ of the i^{th}-PC weight:

$$w_i(\tilde{u}_{0,k}) \approx \hat{w}_i(\tilde{u}_{0,k}) + \left.\frac{\partial w_i}{\partial BR}\right|_{BR_{0,k}} a_{BR}\cos(\omega_{BR} t_k) + \left.\frac{\partial w_i}{\partial \text{SOI}_{\text{DI}}}\right|_{\text{SOI}_{\text{DI},0,k}} a_{\text{SOI}}\cos(\omega_{\text{SOI}} t_k) + R_N$$

where $\hat{w}_i(\tilde{u}_{0,k})$ the zero-frequency component, $\frac{\partial w_i}{\partial BR, \text{SOI}_{\text{DI}}}$ the partial derivatives of w_i with respect to the dithered inputs and R_N the remainder term as a function of all higher order derivatives ($N \geq 2$). Assuming that R_N is sufficiently small and the ES controller is relatively slow compared to the combustion dynamics, a local linear approximation can be obtained:

$$x_{i,k} = x_{i,k-1}$$
$$w_i(\tilde{u}_{0,k}) = H x_{i,k}$$

with

$$x_{i,k} = [\hat{w}_i(\tilde{u}_{0,k}), \frac{\partial w_i}{\partial BR}, \frac{\partial w_i}{\partial \text{SOI}_{\text{DI}}}]^\top$$

the state vector and

$$H = [1, a_{BR}\cos(\omega_{BR} t_k), a_{\text{SOI}}\cos(\omega_{\text{SOI}} t_k)]$$

the measurement transformation vector. Using this state-space model, a KF is used to obtain an estimate of the of the state vector $x_{i,k}$. The time update equations for the KF are:

$$x_{i,k|k-1} = x_{i,k-1|k-1}$$
$$P_{k|k-1} = P_{k-1|k-1} + Q$$

with Q an estimate for the process noise covariance and P the estimate error covariance. Every cycle, a measurement is obtained, allowing the states to be recursively updated as follows:

$$G_k = \frac{P_{k|k-1} H^\top}{H P_{k|k-1} H^\top + R}$$
$$x_{i,k|k} = x_{i,k|k-1} + G_k(w_i(\tilde{u}_{0,k}) - H x_{i,k|k-1})$$
$$P_{k|k} = (I_3 - G_k H) P_{k|k-1}$$

with G_k the Kalman gain, I_3 an identity matrix and R the measurement noise covariance estimate. In practice, there are n_{PC} weights used in the decomposition of the in-cylinder pressure trace $p_{cyl}(\theta)$. Therefore, this

filtering process is repeated n_{PC}-times in parallel. The resulting local estimates are stored in a vector

$$w(\tilde{u}_{0,k}) = \left[x^1_{1,k}, x^1_{2,k}, \dots, x^1_{n_{\text{PC}},k}\right]^\top \in \mathbb{R}^{n_{\text{PC}}}$$

and the PC mapping gradient in a matrix:

$$\left.\frac{\partial w}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}} = \begin{bmatrix} x^2_{1,k} & x^3_{1,k} \\ x^2_{2,k} & x^3_{2,k} \\ \vdots & \vdots \\ x^2_{n_{\text{PC}},k} & x^3_{n_{\text{PC}},k} \end{bmatrix} \in \mathbb{R}^{n_{\text{PC}}\times 2}$$

where the superscript 'n' denotes the n^{th}-element of the state estimation vector x_k. Using these estimations, the control gradient can be constructed as shown in (6), minimizing the composite cost function.

4.4. Constraint handling

In addition to the IMEP_{ref} constraint (5b), the combustion safety constraints (5c–5d) are addressed by the ESC. Decomposing the in-cylinder pressure trace into different PCs, both safety parameters can be formulated as linear inequality constraints as functions of $\hat{w}(\tilde{u}_{0,k})$. These safety limits can be satisfied by projecting the algebraic gradient onto the multi-dimensional hyperplane defined by an active constraint.

Every combustion cycle, the maximum pressure estimate is obtained using the PC weights from the KF:

$$\hat{p}_{max} = \max\left(\hat{w}(\tilde{u}_{0,k})^\top f(\theta) + f_{mp}(\theta)\right) \tag{8}$$

Next, the crank angle at which the pressure is maximum ($\theta_{p_{\max}}$) is found:

$$\theta_{p_{max}} = \arg\max_{\theta\in\Theta}\left(\hat{w}(\tilde{u}_{0,k})^\top f(\theta) + f_{mp}(\theta)\right)$$

and can be used to construct the normal vector of the constrained hyperplane: $\vec{n}_{p_{max}} = f(\theta_{p_{max}})$. Similarly, the estimate for the maximum pressure rise rate is considered:

$$\hat{dp}_{max} = \max\left(\hat{w}(\tilde{u}_{0,k})^\top \frac{df(\theta)}{d\theta} + \frac{df_{mp}(\theta)}{d\theta}\right)$$

along with the crank angle at which this maximum is found ($\theta_{dp_{\max}}$):

$$\theta_{dp_{max}} = \arg\max_{\theta\in\Theta}\left(\hat{w}(\tilde{u}_{0,k})^\top \frac{df(\theta)}{d\theta} + \frac{df_{mp}(\theta)}{d\theta}\right)$$

Using $\theta_{dp_{\max}}$, the normal vector of the constrained hyperplane: $\vec{n}_{dp_{max}} = \frac{df(\theta_{dp_{max}})}{d\theta}$ can be obtained. These estimates, both for the maximum pressure and the maximum pressure rise rate, are used to evaluate the combustion safety constraints in real-time during optimization.

For a single active constraint, such as $\hat{p}_{\max} = p_{max,ub}$ or $\hat{dp}_{\max} = dp_{max,ub}$, it is possible to comply with the constraint while further minimizing the cost function. This is achieved by projecting the algebraic gradient from (7) onto the constrained hyperplane defined by the active constraint. The gradient projection method is demonstrated for an active maximum pressure constraint.

Consider the normal on the hyperplane for the maximum pressure ($p_{\max}$) constraint: $\vec{n}_{p_{\max}} \in \mathbb{R}^{n_{\text{PC}}}$ and a random vector $\vec{m} \in \mathbb{R}^{n_{\text{PC}}}$ of equal length. Vector $\vec{v}_{p_{\max}} \perp \vec{n}_{p_{\max}}$ can be found as:

$$\vec{v}_{p_{max}} = \vec{m} - \text{proj}_{\vec{n}_{p_{max}}}(\vec{m})$$

where $\text{proj}_{\vec{n}_{p_{\max}}}(\vec{m})$ denotes the projection of vector $\vec{m}$ onto $\vec{n}_{p_{\max}}$ and is given as:

$$\text{proj}_{\vec{n}_{p_{max}}}(\vec{m}) = \frac{\vec{n}_{p_{max}} \cdot \vec{m}}{\|\vec{n}_{p_{max}}\|_2^2}\vec{n}_{p_{max}}$$

Using this projection, $\vec{v}_{p_{\max}}$ is tangent with the hyperplane representing the maximum pressure constraint. Similarly, the algebraic gradient can be projected onto the hyperplane using $\vec{v}_{p_{\max}}$:

$$\text{proj}_{\vec{v}_{p_{max}}}\left(\left.\frac{\partial J}{\partial w}\right|_{\tilde{u}_{0,k}}\right) = \frac{\vec{v}_{p_{max}} \cdot \left.\frac{\partial J}{\partial w}\right|_{\tilde{u}_{0,k}}}{\|\vec{v}_{p_{max}}\|_2^2}\vec{v}_{p_{max}}$$

Table 2
Experimental operating conditions.

Parameter	Value	Unit
IMEP_{ref}	4	[bar]
$\mathcal{U}_{BR}$	[0.3 - 0.6]	[-]
$\mathcal{U}_{\text{SOI}}$	[3 - 11]	[CADbTDC]
SOI_{PFI}	360	[CADbTDC]
LHV_{DI}	42.6	[J/mg]
LHV_{PFI}	43.4	[J/mg]
p_{im}	1.39 ± 0.08	[bar]
T_{im}	21.8	[degC]
$T_{coolant}$	70	[degC]
T_{oil}	90	[degC]

$$= \left(\left.\frac{\partial J}{\partial w}\right|_{\tilde{u}_{0,k}}\right)_{mod}$$

which results in a modified algebraic gradient that minimizes the cost function while satisfying the active constraint bound. The control gradient from (6) can be written as:

$$\left.\frac{\partial J}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}} = \left(\left.\frac{\partial J}{\partial w}\right|^\top_{\tilde{u}_{0,k}}\right)_{mod} \left.\frac{\partial w}{\partial \tilde{u}}\right|_{\tilde{u}_{0,k}}$$

This gradient projection method is capable of dealing with only one active constraint. The actuator references $\tilde{u}_0$ are saturated to remain within their feasible region $\mathcal{U}$.

5. Experimental results

The ES algorithm is demonstrated on the SCE from Section 2

5.1. Selection of principal components to include in the optimization

The in-cylinder pressure trace is decomposed into eight Principal Components (PCs), which was shown to give the highest accuracy with minimal number of PCs. The Principal Components used in this study are shown in Fig. 3 and are computed from the same dataset as in Vlaswinkel and Willems, 2024. To decide which Principal Component weights are included into the optimization, fixed $\tilde{u}_0$ reference experiments were performed to determine the gradient estimation quality, which depends on the Kalman filtering process. The KF measurement noise estimate R was measured on the SCE and the process noise estimate Q was tuned manually. These experimental results around a fixed reference input: $\text{SOI}_{\text{DI},0,k=0} = 4$[CADbTDC] and $BR_{0,k=0} = 0.425$[-] are shown in Fig. 4. From these figures, it can be seen that the gradients of the first two PC weights converge much slower compared to other PC weights. The other gradient estimates converged much faster for the defined filtering parameters. It was decided that the quality of the gradient estimation for the first two Principal Components was insufficient using the set Kalman filter parameters. Therefore, they were omitted from the optimization.

The remaining PCs (3 – 8) are used to construct the optimal control gradient such that $n_{\text{PC}} = 6$. The relevant experimental operating conditions are shown in Table 2.

5.2. Fuel efficiency optimization

The real-time optimization is first demonstrated using initial conditions: $\text{SOI}_{\text{DI},0,k=0} = 4$[CADbTDC] and $BR_{0,k=0} = 0.425$[-]. The ES and Kalman filter parameters are presented in Table 3. The dither frequencies were selected much slower to enable the IMEP-controller to statically decouple IMEP_{ref} from $\tilde{u}_0$. The amplitudes were determined to enable the Kalman filter to accurately identify the partial derivatives with respect to the inputs, approximately 20% of the actuator range $\mathcal{U}$ sufficed. The control gains were selected arbitrarily to balance fast convergence and stability during optimization. The Kalman filter

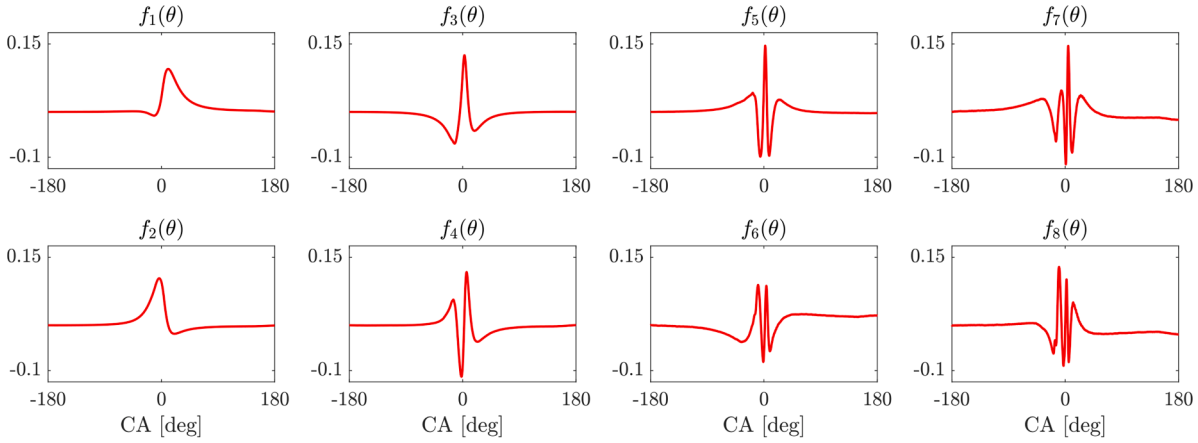

Fig. 3. The Principal Component traces used to decompose the measured in-cylinder pressure traces. The Principal Components are computed from the same dataset as in Vlaswinkel and Willems, 2024.

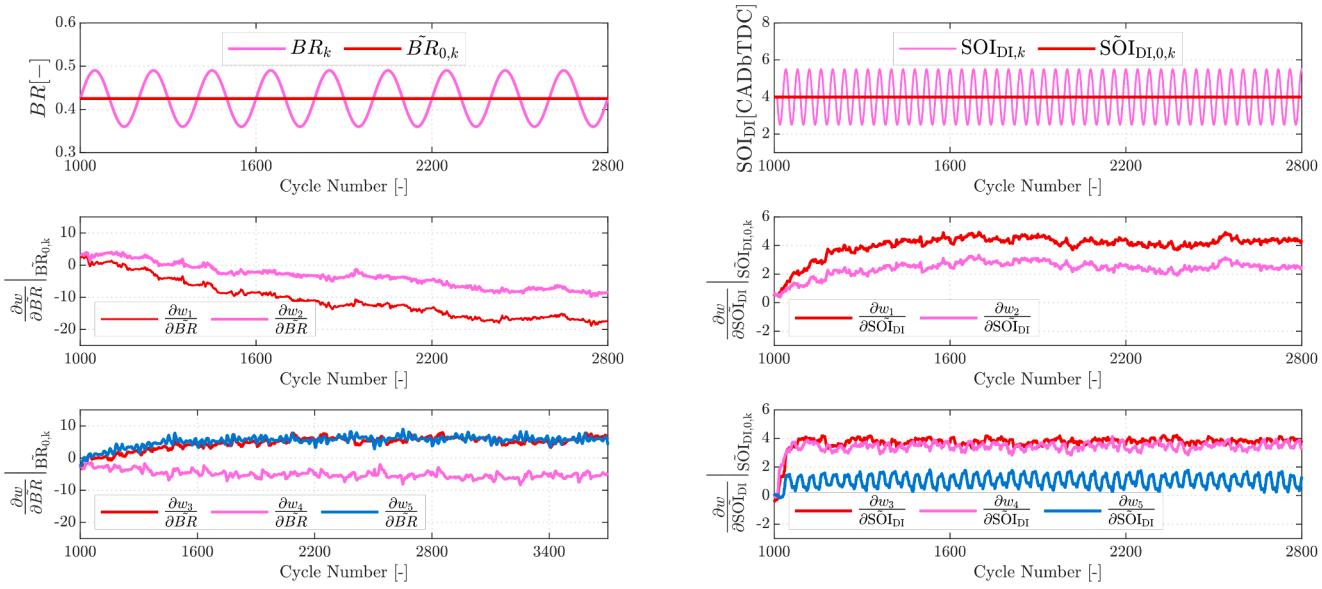

Fig. 4. fixed $\tilde{u}_0$ reference experiments to determine the gradient estimation quality of the individual Principal Component weights. Similar results as PC numbers 3–5 were observed for PC numbers 6–8.

parameter R was estimated from static measurements around $\tilde{u}_{0,k=0}$ and parameter Q was tuned manually. The ES controller is activated after 1000 engine cycles, ensuring that the KF has converged to an accurate state estimate (≈ 500 cycles). Convergence of the ES algorithm is achieved when both of the fuel path references $\tilde{u}_0$ are saturated at the input boundary $\mathcal{U}$ or the deviation of $\tilde{u}_0$ is not larger than 0.1% over the last 300 cycles.

The performance of the method is evaluated using a relative efficiency increase (δGIE). This change in efficiency is always measured starting from the GIE at the initial conditions ($\mathrm{GIE}_{k=0}$), representing the baseline and can be defined as:

$$\delta\mathrm{GIE} = \mathrm{GIE}_k - \mathrm{GIE}_{k=0}$$

where GIE_k the efficiency evaluated for each combustion cycle. The results of the ES algorithm are illustrated in Fig. 5, where convergence has been achieved after 2444 engine cycles(≈ 4 minutes). The next-cycle IMEP-controller ensures stable load tracking during optimization ($\mathrm{IMEP}_g = 4 \pm 0.25$[bar]) and decouples the driver demand ($\mathrm{IMEP}_{\mathrm{ref}}$) from the dithered inputs ($\tilde{BR}_0, \tilde{\mathrm{SOI}}_{\mathrm{DI},0}$). During optimization, the cost J decreases and accordingly, an increase in Gross Indicated Efficiency is observed (Vlaswinkel and Willems 2023). The new setpoint after 2444 cycles is given as: $\mathrm{SOI}_{\mathrm{DI},0} = 7.60$[CADbTDC] and $BR_0 = 0.3$[-] and results in a 3.74% increase in efficiency. This convergence time of approximately 4 min is a significant improvement compared to a traditional Design of Experiment (DoE) calibration. To highlight the possible improvements, consider a DoE where two fuel path actuators need to

Table 3
Extremum Seeking and Kalman filter parameters.

Parameter	Value	Unit
n_{PC}	6	[-]
ω_{BR}	0.05	[Hz]
ω_{SOI}	0.25	[Hz]
a_{BR}	0.065	[-]
a_{SOI}	1.5	[CAD]
K_{BR}	4.5	[-]
K_{SOI}	5500	[-]
K_{IMEP}	95	$[\frac{J}{\Delta\mathrm{bar}}]$
Q	$[diag([1, 10, 1]) \otimes 0.0025 I_{n_{\mathrm{PC}}}]$	[-, -, $\frac{1}{\mathrm{CAD}^2}$]
R	[1.35,0.31,0.25,0.3,0.075,0.15]	[-]

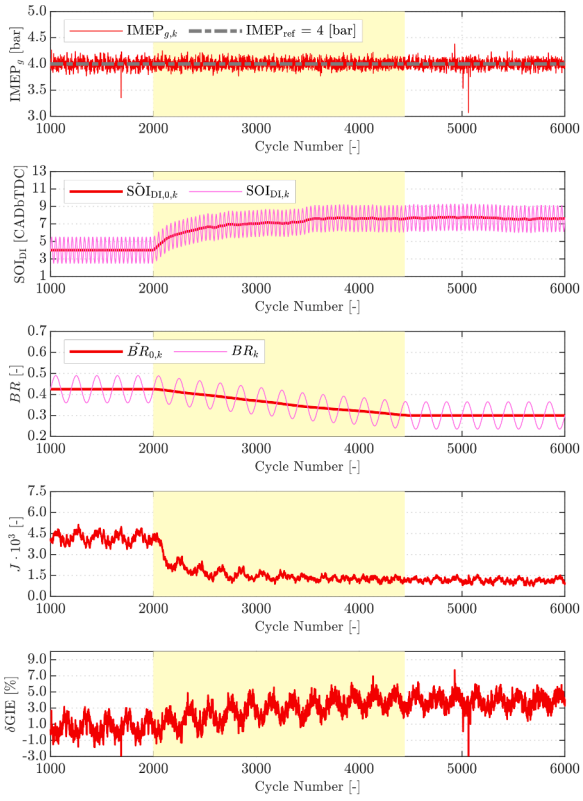

Fig. 5. Extremum Seeking optimization results for $SOI_{0,DI,k=0} = 4$[CADbTDC] and $BR_{0,k=0} = 0.425$[-]. ES convergence is highlighted in the yellow area. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

be calibrated. A 5×5 grid is spanned across the operating space and each grid point takes 1 min to stabilize and acquire data. This results in 25 min of measurement data to find the optimal references. This highlights the effectiveness of the automated Extremum Seeking calibration framework.

5.3. Convergence

To illustrate the convergence properties of the algorithm, see Fig. 6, which shows the control gradient $\frac{\partial J}{\partial \tilde{SOI}_{DI}}$, alongside the algebraic gradients $\frac{\partial J}{\partial w_i}\Big|_{\tilde{SOI}_{DI,0,k}}$ and mapping gradients $\frac{\partial w_i}{\partial \tilde{SOI}_{DI}}\Big|_{\tilde{SOI}_{DI,0,k}}$ for the first three controlled PC weights ($i = 3, 4, 5$). As mentioned, the SOI_{DI} converges when the control gradient ≈ 0, shown in the first plot. The second plot shows the algebraic gradient, which relates to the local sensitivity of the PC weights $\rightarrow J$. In theory, w_4 has the largest potential for further minimizing the cost J, as its magnitude is the greatest. In practice, exploiting this potential depends on the controllability from $\tilde{SOI}_{DI,0} \rightarrow w_4$. This controllability is defined by the PC mapping gradients, which represents the sensitivity of changes in $\tilde{u}_0$ with respect to the PC weights. These PC mapping gradients are shown with respect to the dithered input SOI_{DI} in the bottom figure. It illustrates that w_4 becomes less controllable from SOI_{DI} compared to w_3 and w_5 during optimization, as the mapping gradient magnitude decreases over time. The control gradient $\frac{\partial J}{\partial \tilde{SOI}_{DI}}$ is a compromise between optimality via the algebraic gradient and controllability of the PC weights via the mapping gradients. Each individual PC weight influences the overall control gradient and contributes to the convergence of the algorithm.

This experiment was conducted under different initial conditions to study the convergence properties of the algorithm. The results of these runs are summarized in Table 4 and are plotted on a normalized efficiency map, which is measured over the complete actuator range ($\mathcal{U}$). This map was generated by evaluating the steady-state GIE at a fixed BR and SOI_{DI} for 2500 consecutive cycles. In Figs. 7 and 9, the measurement points are highlighted as yellow dots. The efficiency map was normalized according to:

$$\delta\text{GIE} = \text{GIE}(\mathcal{U}) - \min(\text{GIE}(\mathcal{U}))$$

The results from Table 4 are visualized in Fig. 7.

The maximum GIE within the feasible actuator range is not found, likely because the control gradient construction omits w_1 and w_2. As shown in Fig. 8, which plots the weighted squared errors $(\hat{w}_i - w_{\text{ITC},i})^2 \int_\Theta f_i(\theta)^2$ for $i = 1, 2, \ldots, 8$, the errors for w_1 and w_2 increase overall during optimization, yielding a suboptimal solution for the GIE. For w_7, the error initially rises but later drops below its initial setpoint. This behavior highlights a trade-off in the control gradient: the algorithm temporarily increases the error for w_7, as this action is outweighed by the reductions in errors for the other PC weights.

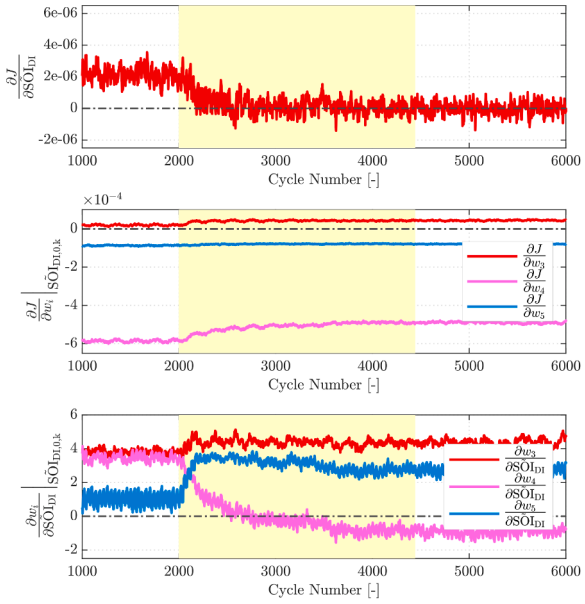

Fig. 6. Control gradient composed of algebraic gradients and PC mapping gradients for $\tilde{SOI}_{DI}$.

Table 4
Extremum Seeking control starting from various initial conditions.

Initial conditions		Final conditions		
$SOI_{DI,0}$ [CADbTDC]	BR_0 [−]	$SOI_{DI,0}$ [CADbTDC]	BR_0 [−]	δGIE δ%
4.0	0.425	7.60	0.3	3.74
9.0	0.5	7.64	0.3	2.54
4.0	0.325	7.83	0.3	2.21
4.0	0.55	7.65	0.3	6.11

5.4. Constraint handling for maximum in-cylinder pressure

To demonstrate the algorithm's constraint handling, an arbitrary maximum pressure limit is introduced. The maximum in-cylinder pressure is constrained on its mean-value estimate, according to (8). The upper bound ($p_{max,ub}$) is set at 78 [bar] and the initial conditions are chosen as: $SOI_{DI,0,k=0} = 4$[CADbTDC] and $BR_{0,k=0} = 0.5$[-].

The results are shown in Fig. 9, where the GIE increased by 3.1%.

The maximum pressure during optimization is illustrated in Fig. 10, where the mean-value estimate $\hat{p}_{\max}$ is centered around 78 [bar], with

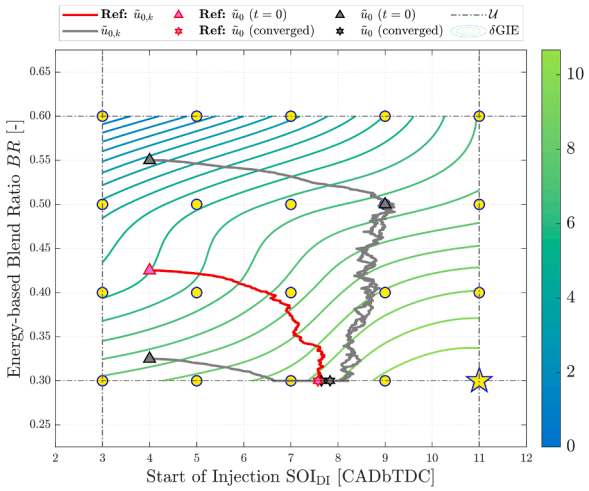

Fig. 7. Extremum Seeking optimization results for different initial conditions. The reference case is highlighted. The contour plot is shows $\delta\mathrm{GIE} = \mathrm{GIE}(\mathcal{U}) - \min(\mathrm{GIE}(\mathcal{U}))$ and the maximum GIE is shown in the bottom right corner (★). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

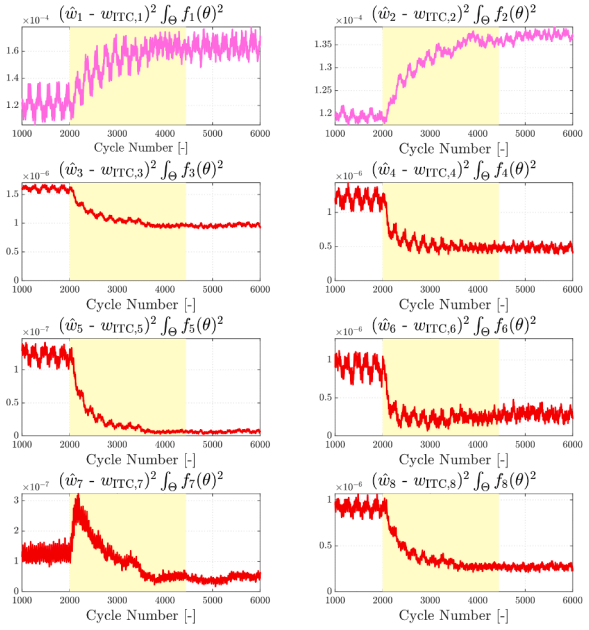

Fig. 8. Squared weighted PC weights error traces during optimization. Controlled weights are shown in red and the uncontrolled weights are pink. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

a maximum constraint violation of 0.45 [bar] during optimization. The maximum pressure oscillation (pink shaded area) is caused by the pneumatic intake valve, which is not actively controlled. This highlights the robustness for the constraint handling, as the intake manifold pressure varies between 1.31 [bar] and 1.47 [bar], influencing $p_{\max}$. The new set point is found only when the aforementioned convergence criteria are relaxed. This is required because the ES algorithm compensates for the varying intake pressure. The new reference control settings are given as: $\mathrm{SOI}_{\mathrm{DI},0,k=0} = 6.42$[CADbTDC] and $BR_{0,k=0} = 0.3$[-].

To constrain the actual maximum pressure and not the mean-value, an additional safety margin (S_{margin}) has to be included to compensate for the changing intake pressure, the influence of the dither signals and the process noise.

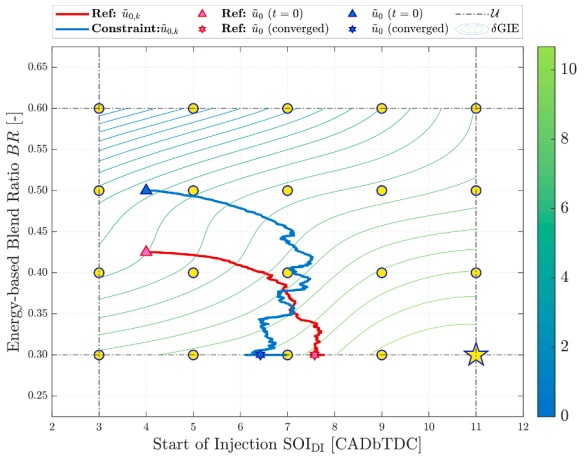

Fig. 9. Extremum Seeking optimization results with artificial maximum pressure constraint.

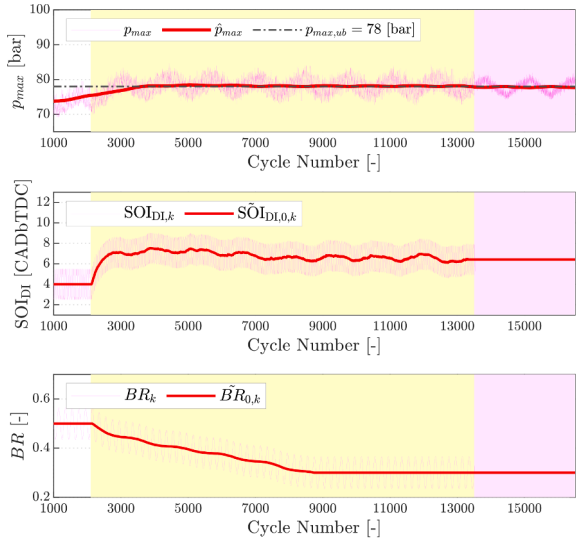

Fig. 10. Maximum in-cylinder pressure constrained on its mean-value estimate $\hat{p}_{\max}$ at 78 [bar].

6. Conclusion

This work introduces a novel auto-calibration method based on constrained Extremum Seeking and demonstrates its effectiveness on a single-cylinder engine testbench. The algorithm is computationally efficient as it runs on a cycle-to-cycle basis using a Speedgoat Real-Time Target Machine. The algorithm directly optimizes in-cylinder pressure by shaping it as close as possible towards an Ideal Thermodynamic cycle, thereby increasing the Gross Indicated Efficiency of the engine. The approach balances controllability and optimality of the Principal Component weights for the under-actuated system. Experimental results confirm that the method converges to an optimal cost defined by J and explicitly deals with safety constraints.

Future work will involve identifying an optimal selection of Principal Component weights for each input actuator, thereby optimizing both controllability and optimality for the Principal Components. The Principal Component traces were computed from a Reactivity Controlled Compression Ignition (RCCI) dataset. Additional research is needed to study the influence of this transfer of PC traces to different combustion concepts. Next, the possibility of increasing the number of actuators will be studied to increase control authority of the in-cylinder

pressure trace. The maximum pressure rise rate constraint will be experimentally demonstrated on the SCE. Lastly, combustion stability will be included as a constraint, which requires an extension of the current framework, as cyclic variability hasn't been treated yet.

CRediT authorship contribution statement

Mats Versmissen: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Maarten Vlaswinkel:** Supervision, Methodology, Investigation, Conceptualization; **Frank Willems:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Funding acquisition.

Funding

This work is financially supported by the Dutch GroenvermogenNL program.

Declaration of competing interests

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Mats Versmissen reports financial support was provided by Dutch GroenvermogenNL program. Mats Versmissen reports a relationship with Eindhoven University of Technology that includes: employment. 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

The authors would like to thank Michel Cuijpers, Theo de Groot, and Bart van Pinxten from the Zero Emission Lab for their technical support during experiments.

References

Agarwal, A. K., Singh, A. P., & Maurya, R. K. (2017). Evolution, challenges and path forward for low temperature combustion engines. *Progress in Energy and Combustion Science*, *61*, 1–56. https://doi.org/10.1016/j.pecs.2017.02.001

Heywood, J. B. (1988). Internal combustion engine fundamentals. https://api.semanticscholar.org/CorpusID:110590482.

IEA (2023). Transport. https://www.iea.org/energy-system/transport. Accessed: 2025-10-16.

Karimi, M., Wang, X., Hamilton, J., Negnevitsky, M., & Lyden, S. (2021). Status, challenges and opportunities of dual fuel hybrid approaches-a review. *International Journal of Hydrogen Energy*, *46*(70), 34924–34957. https://doi.org/10.1016/j.ijhydene.2021.08.008

Killingsworth, N. J., Aceves, S. M., Flowers, D. L., Espinosa-Loza, F., & Krstic, M. (2009). HCCI Engine combustion-timing control: Optimizing gains and fuel consumption via extremum seeking. *IEEE Transactions on Control Systems Technology*, *17*(6), 1350–1361. https://doi.org/10.1109/TCST.2008.2008097

Popovic, D., Jankovic, M., Magner, S., & Teel, A. (2003). Extremum seeking methods for optimization of variable cam timing engine operation. In *Proceedings of the 2003 American control conference, 2003.* (pp. 3136–3141). (*vol. 4*). https://doi.org/10.1109/ACC.2003.1244011

Scheinker, A. (2024). 100 Years of extremum seeking: A survey. *Automatica*, *161*, 111481. https://doi.org/10.1016/j.automatica.2023.111481

Vlaswinkel, M., & Willems, F. (2023). Cylinder pressure feedback control for ideal thermodynamic cycle tracking: Towards self-learning engines. *IFAC-PapersOnLine*, *56*(2), 8260–8265. https://doi.org/10.1016/j.ifacol.2023.10.1011

Vlaswinkel, M., & Willems, F. (2024). Data-based in-cylinder pressure model with cyclic variations for combustion control: An RCCI engine application. *Energies*, *17*, 1881. https://doi.org/10.3390/en17081881

van der Weijst, R., van Keulen, T., & Willems, F. (2019). Constrained multivariable extremum-seeking for online fuel-efficiency optimization of diesel engines. *Control Engineering Practice*, *87*, 133–144. https://doi.org/10.1016/j.conengprac.2019.03.008

van der Weijst, R., van Keulen, T. A. C., & Willems, F. P. T. (2017). A generalized framework for perturbation-based derivative estimation in multivariable extremum-seeking. *IFAC-PapersOnLine*, *50*(1), 3148–3153. https://doi.org/10.1016/j.ifacol.2017.08.326

Welch, G., & Bishop, G. (2006). An introduction to the Kalman filter. Technical Report TR 95-041. University of North Carolina at Chapel Hill, Department of Computer Science Chapel Hill, NC 27599-3175.

Willems, F. (2018). Is cylinder pressure-based control required to meet future HD legislation? *IFAC-PapersOnLine*, *51*(31), 111–118. https://doi.org/10.1016/j.ifacol.2018.10.021

Willems, R. (2020). Designed experiments for efficient engines. Phd thesis 1 (research tu/e / graduation tu/e) Mechanical Engineering. Dissertation.