



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Maass, A;Manzie, C;Shames, I;Nakada, H

Title:

Zeroth-Order Optimization on Subsets of Symmetric Matrices With Application to MPC Tuning

Date:

2021-11-08

Citation:

Maass, A., Manzie, C., Shames, I. & Nakada, H. (2021). Zeroth-Order Optimization on Subsets of Symmetric Matrices With Application to MPC Tuning. IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY, PP (4), <https://doi.org/10.1109/TCST.2021.3123605>.

Persistent Link:

<https://hdl.handle.net/11343/297893>

Zeroth-order optimisation on subsets of symmetric matrices with application to MPC tuning

Alejandro I. Maass, Chris Manzie, Iman Shames, and Hayato Nakada

Abstract—This paper provides a zeroth-order optimisation framework for non-smooth and possibly non-convex cost functions with matrix parameters that are real and symmetric. We provide complexity bounds on the number of iterations required to ensure a given accuracy level for both the convex and non-convex case. The derived complexity bounds for the convex case are less conservative than available bounds in the literature since we exploit the symmetric structure of the underlying matrix space. Moreover, the non-convex complexity bounds are novel for the class of optimisation problems we consider. The utility of the framework is evident in the suite of applications that use symmetric matrices as tuning parameters. Of primary interest here is the challenge of tuning the gain matrices in model predictive controllers, as this is a challenge known to be inhibiting industrial implementation of these architectures. To demonstrate the framework we consider the problem of MIMO diesel air-path control, and consider implementing the framework iteratively “in-the-loop” to reduce tracking error on the output channels. Both simulations and experimental results are included to illustrate the effectiveness of the proposed framework over different engine drive cycles.

Index Terms—Numerical optimisation, gradient-free, complexity bounds, model predictive control, diesel air-path

I. INTRODUCTION

OPTIMISATION with matrix-valued decision variables is a problem that appears in a wide variety of applications. Particularly, this approach can be found in non-negative matrix factorization [1], low-rank matrix optimisation [2], [3], signal processing with independent component analysis [4], energy efficient wireless systems [5], graph weight allocation [6], [7], and controller tuning such as model predictive controllers [8],

This work was supported in part by the Australian Research Council through linkage project LP160100650, and by Toyota Motor Corporation, Japan.

A.I. Maass and C. Manzie are with the Department of Electrical and Electronic Engineering, and the Melbourne Information, Decision and Autonomous Systems (MIDAS) Laboratory, The University of Melbourne, Melbourne, VIC 3010, Australia (email: alejandro.maass@unimelb.edu.au; manziec@unimelb.edu.au).

I. Shames is with the College of Engineering & Computer Science, The Australian National University, Australia (email: iman.shames@anu.edu.au)

H. Nakada is with the Advanced Unit Management System Development Division, Toyota Motor Corporation, Japan (email: hayato.nakada@mail.toyota.co.jp)

linear quadratic regulators [9], and PIDs for MIMO systems [10]. In most of the aforementioned applications, either explicit expressions of the cost function are not available, or derivatives are difficult or infeasible to obtain. Consequently, we restrict our attention to this class of problems in which the cost function is treated as a black-box. In this case, the cost function directional derivatives can still be approximated by means of cost function values using finite-differences for instance. Of main interest in this paper is the challenge of tuning the gain matrices in model predictive controllers (MPC), which have the property of being symmetric and positive (semi) definite. Efficient calibration of MPCs is often a difficult task given the large number of tuning parameters and their non-intuitive correlation with the output response.

Existing literature regarding MPC tuning can be broadly divided in heuristic guidelines [11], analytical studies [12]–[14], and meta-heuristic algorithms [8], [15]–[21]. Both heuristic guidelines and analytical studies hold for specific scenarios and MPC formulations, thus trial-and-error approaches are unavoidably adopted in practice. Regarding meta-heuristic methods, the works [15], [18]–[21] focus on set-point tuning based on step response characteristics (e.g. overshoot, settling and rise times, and steady-state errors). The authors in [16] deal with trajectory tuning of MPC in which the trajectories are generated by first-order transfer function models. General trajectory tuning is considered in [17] and [8] via particle swarm optimisation (PSO) and gradient-based algorithms, respectively. The number of particles in PSO algorithms depends on the specific problem and it usually oscillates around 20 to 50 particles. This translates into running 20 to 50 closed-loop experiments per iteration, which is not practical for applications where the plant is in the loop. Gradient-based methods [8] and Bayesian optimisation (BO) methods [22], [23] can also be used as an alternative with less experiments per iteration, but existent results are presented only for vectors of parameters (i.e. diagonal tuning matrices). Lastly, inverse optimal control techniques such as [24] could be used for MPC tuning, see e.g. [25]. However, [25] does not provide any theoretical guarantees and it relies on a good choice of initial weighting matrices and heuristic guidelines. We note there is currently a lack of a general framework with convergence guarantees that can directly deal with, not only diagonal matrices, but also general subsets of symmetric matrices in which the algorithm preserves the structure of the tuning

matrix at every iteration.

In response to the above discussion, we propose a zeroth-order optimisation framework for non-smooth and possibly non-convex cost functions with matrix parameters that are real and symmetric. We adopt iterates that use zeroth-order matrix oracles to approximate the directional derivative in a randomly chosen direction. This random direction is taken from the Gaussian orthogonal ensemble (GOE), which is the set of matrices with independent normally distributed entries, see e.g. [26]–[28]. The algorithm we use is the natural extension of the zeroth-order random search algorithm from [29] but tailored to deal with cost functions with matrix parameters. For this optimisation framework, we provide convergence guarantees in terms of complexity bounds for both convex and non-convex cost functions. In [29], the counterpart of our optimisation problem is studied, in which functions with vector parameters are used, and the zeroth-order oracle samples from the normal distribution. The authors in [29] present complexity bounds for both convex and non-convex functions. The bounds for the convex case can be applied to our matrix setup if one vectorises both the parameter and random direction matrices. However, this does not carry the matrix structure and, in fact, these bounds are significantly more conservative than the ones we propose. Moreover, the non-convex bounds proposed in [29] are only applicable to unconstrained problems and thus cannot be used in our constrained setting. Consequently, we develop novel bounds for the non-convex case, in which the optimisation parameters are constrained to subsets of symmetric matrices. The optimisation framework presented in this paper, when applied to the context of MPC tuning, offers a more general approach with respect to available literature, and it also comes with convergence guarantees. Particularly, it can deal with MPC tuning over trajectories in a general setting as opposed to [15], [16], [18]–[20], and it provides MPC weighting matrices that satisfy the required constraints of being symmetric and positive (semi) definite, and thus provide more degrees of freedom than the usual diagonal choices in [8], [17], [22], [23].

The applicability of the proposed approach is illustrated via both simulations and experiments in the problem of diesel air-path control. The MPC parameters are iteratively tuned within the closed-loop setting with the goal of improving the overall tracking performance over an engine drive cycle. For the simulation study, we compare the performance of our proposed framework with other gradient-free algorithms available in the literature such as dividing rectangles (DIRECT), PSO, and BO. In the experimental testing, we tune MPC controllers in a diesel engine test bench over two segments of the new European driving cycle (NEDC), and one segment of the worldwide harmonised light duty test cycle (WLTC).

A summary of our main contributions can be found below.

- 1) We present theoretical guarantees for a general class of optimisation problems with matrix-valued decision variables on subsets of real and symmetric matrices. The adopted iterates are tailored to this setting in the sense that they produce matrices that belong to the adequate space through projection. When the cost function is convex, the derived complexity bounds are demonstrably less

conservative than the ones in [29] (vector-valued decision variables), since we exploit the symmetric structure of the underlying matrix space. In the non-convex case, we derive new complexity bounds that were not available in the literature for this context.

- 2) We illustrate the applicability of our framework by using it to tune the matrix parameters in MPC. This extends current literature such as [15], [16], [18]–[21].
- 3) We highlight the efficacy of the approach by a simulation study in which we compare it to other available gradient-free algorithms in the literature (DIRECT, PSO, and BO).
- 4) We provide various experimental evaluations in a diesel engine test bench showing significant improvement of MPC tracking performance with only a few iterations of the proposed optimisation approach.

Organisation: Section II introduces the optimisation framework. The complexity bounds are derived in Section III. In Section IV, the proposed optimisation framework is applied to MPC tuning in the context of air-path control. In Section V, we compare different gradient-free algorithms available in the literature to our framework. Experimental testing is presented in Section VI for different engine drive cycles. Lastly, conclusions are drawn in Section VII.

Notation: Denote by $\mathbb{R}^{m \times n}$ the set of real matrices of dimension $m \times n$, and \mathbb{S}^n the set of real symmetric matrices of dimension $n \times n$. Let $\mathbb{N} \triangleq \{1, 2, \dots\}$ and $\mathbb{N}_0 \triangleq \mathbb{N} \cup \{0\}$. Given a matrix M , M^\top denotes its transpose. We use $\text{diag}\{M_1, \dots, M_n\}$ to denote the standard block diagonal matrix. The identity and zero matrices of dimension $m \times n$ are denoted by $I_{m \times n}$ and $0_{m \times n}$, respectively. $\nabla f(X)$ denotes the matrix gradient for a scalar function f with matrix parameter X , see e.g. [30], [31]. For matrices $M \in \mathbb{R}^{n \times n}$ and $N \in \mathbb{R}^{n \times n}$, the Frobenius inner product is defined as $\langle M, N \rangle_F \triangleq \text{Tr}\{M^\top N\}$, which induces the norm $\|M\|_F \triangleq \sqrt{\langle M, M \rangle_F}$. For a random variable $x \in \mathbb{R}$, we write $x \sim \mathcal{N}(0, \sigma^2)$ to say that x is normally distributed with zero mean and variance σ^2 .

II. OPTIMISATION FRAMEWORK

We consider the following class of optimisation problems

$$\min_{X \in \mathbb{Q}^n} f(X), \quad (1)$$

where X is an $n \times n$ matrix parameter, $\mathbb{Q}^n \subset \mathbb{R}^{n \times n}$ is a closed convex set of admissible parameters, and f is a non-smooth and possibly non-convex cost function. We build upon the random search ideas of [29], but tailored to the matrix space \mathbb{Q}^n . That is, we are interested in solving (1) using iterates of the form

$$X_{k+1} = \pi_{\mathbb{Q}^n}\{X_k - h_k \mathcal{O}_\mu(X_k, U_k)\}, \quad (2)$$

where $\pi_{\mathbb{Q}^n}$ denotes the Euclidean projection onto the closed convex set \mathbb{Q}^n , h_k is a positive scalar known as the step size, and \mathcal{O}_μ denotes the so-called *zeroth-order random oracle* which is defined as

$$\mathcal{O}_\mu(X, U) \triangleq \frac{f(X + \mu U) - f(X)}{\mu} U, \quad (3)$$

Zeroth-Order Random Matrix Search (ZO-RMS)

- 1: Choose $X_0 \in \mathbb{Q}^n$, $\mu > 0$ and $\{h_k\}_{k \in \mathbb{N}_0}$.
 - 2: **for** $k = \{0, \dots, N\}$, $N \in \mathbb{N}$ **do**
 - 3: Generate $U_k \in \mathbb{G}^n$
 - 4: $\mathcal{O}_\mu(X_k, U_k) \leftarrow \frac{1}{\mu}(f(X_k + \mu U_k) - f(X_k))U_k$
 - 5: $X_{k+1} \leftarrow \pi_{\mathbb{Q}^n}\{X_k - h_k \mathcal{O}_\mu(X_k, U_k)\}$
 - 6: **end for**
 - 7: Return $\hat{X}_N \triangleq \arg \min_X [f(X) : X \in \{X_0, \dots, X_N\}]$.
-

where $\mu > 0$ denotes the oracle's precision, and U is a random symmetric matrix that belongs to the Gaussian orthogonal ensemble (GOE) as per Definition 1 below.

Definition 1: The GOE, denoted by \mathbb{G}^n , is the set of random symmetric matrices $U \in \mathbb{S}^n$ with i.i.d. entries such that $[U]_{ii} \sim \mathcal{N}(0, 1)$, and $[U]_{ij} \sim \mathcal{N}(0, \frac{1}{2})$ independent of $[U]_{ii}$ for $i < j$, see e.g. [26]–[28].

The overall method to solve (1) is represented by the pseudo-code above, which we name *zeroth-order random matrix search* (ZO-RMS) algorithm, where N denotes the number of iterations. Note that the oracle, per iteration, only requires the cost function value at two points instead of first-order or second-order derivative information, and it computes an estimate of the directional derivative in a randomly chosen direction.

We emphasise that the ZO-RMS algorithm is the natural extension of the random vector search algorithm proposed in [29] but tailored to matrix decision variables. This extension is essential to our setting since there is currently an absence of a general framework for the constrained class of problems (1), i.e. adequate algorithm and convergence guarantees. In [29], the authors provide gradient-free algorithms to solve the optimisation problem $\min_{x \in \mathbb{Q}^n \subset \mathbb{R}^n} f(x)$ for both convex and non-convex cost functions, which is the counterpart of (1) for vector optimisation variables. A zeroth-order random oracle with a multivariate normally distributed direction u is used. In this paper, we extend the setting in [29] to optimisation problems over a matrix space using random search matrices U from the GOE, which are the natural counterpart to the normally distributed vector u . Note that the complexity bounds provided in [29] for the convex case are applicable to our setting if we consider $x = \text{vec}\{X\}$ and $u = \text{vec}\{U\}$. However, the matrix structure of X and U is immediately lost. Consequently, the first theoretical goal of this work is to obtain less conservative complexity bounds for the ZO-RMS algorithm when f in (1) is convex, by exploiting the structure of the underlying matrix space. With respect to the non-convex case, we note that the complexity bounds available in [29] are not applicable to our case since they consider an unconstrained non-convex problem in \mathbb{R}^n as opposed to a constrained one such as (1). Therefore, our second theoretical goal is to develop new complexity bounds for the ZO-RMS algorithm when solving constrained non-convex optimisation problems in a matrix space, i.e. problem (1) with f non-convex.

As mentioned in the introduction, the proposed optimisation framework presented here is general and fits many

applications. Therefore, the theoretical results presented in the following section will hold for any family of problems that fit the framework. Particularly, in this paper we will illustrate how to apply this framework to tune MPC controllers, since it is an important application for which we can provide interesting contributions validated by simulations and experiments.

III. COMPLEXITY BOUNDS

In this section, we study the performance of the ZO-RMS algorithm in terms of complexity bounds that guarantees a given level of accuracy for the algorithm. We provide complexity bounds for both the convex and non-convex case.

Let us first introduce some important definitions. Definition 1 implies that the probability distribution $\mathbb{P}(dU)$ in the GOE \mathbb{G}^n is $\mathbb{P}(dU) = \frac{1}{\kappa} e^{-\frac{1}{2}\|U\|_F^2} dU$, where dU is the Lebesgue measure on $\mathbb{S}^n \cong \mathbb{R}^{n(n+1)/2}$, i.e. $dU \triangleq \prod_{i=1}^n dU_{ii} \prod_{i < j} dU_{ij}$, and κ is the normalising constant defined as $\kappa \triangleq \int_{\mathbb{G}^n} e^{-\frac{1}{2}\|U\|_F^2} dU = 2^{\frac{n}{2}} \pi^{\frac{n^2+n}{4}}$. Therefore, we can define the expectation of a function $q : \mathbb{G}^n \rightarrow \mathbb{R}$ as $\mathbb{E}_U\{q(U)\} \triangleq \int_{\mathbb{G}^n} q(U) \mathbb{P}(dU) = \frac{1}{\kappa} \int_{\mathbb{G}^n} q(U) e^{-\frac{1}{2}\|U\|_F^2} dU$, and the following moments of interest in \mathbb{G}^n , $m_p \triangleq \mathbb{E}_U\{\|U\|_F^p\} = \frac{1}{\kappa} \int_{\mathbb{G}^n} \|U\|_F^p e^{-\frac{1}{2}\|U\|_F^2} dU$, for $p \in \mathbb{N}_0$. We state the following lemma for some moments of interest.

Lemma 1: For every $n \geq 1$, we have that $m_1 \leq \sqrt{\frac{n^2+n}{2}}$, $m_2 = \frac{n^2+n}{2}$, and $m_4 = \frac{1}{4}(n^4 + 2n^3 + 5n^2 + 4n)$.

Proof: Define $\psi(p) \triangleq \ln(m_p)$, and note that this function is convex in p . Let us write $p = (1 - \alpha) \times 0 + \alpha \times 2$ (i.e. $\alpha = p/2$). For $p \in [0, 2]$ we have $\alpha \in [0, 1]$. Therefore, since $\psi(p)$ is convex, we have that $\psi(p) = \psi(\alpha \times 2 + (1 - \alpha) \times 0) \leq \alpha \psi(2) + (1 - \alpha) \psi(0)$. Particularly, $\psi(0) = \ln(m_0) = 0$, and $\psi(2) = \ln(m_2)$, then $\psi(p) \leq \alpha \ln(m_2) = \frac{p}{2} \ln(m_2)$, and thus $m_p \leq m_2^{p/2}$. Particularly, $m_1 \leq \sqrt{m_2}$. On the other hand, from [32] we get that $m_2 = (n^2 + n)/2$ and $m_4 = \frac{1}{4}(n^4 + 2n^3 + 5n^2 + 4n)$, which concludes the proof. ■

Our analysis relies in the following assumption.

Assumption 1:

- (a) The set of admissible parameters is a subset of the set of real symmetric matrices of dimension $n \times n$, i.e. $\mathbb{Q}^n \subset \mathbb{S}^n$.
- (b) The function $f : \mathbb{Q}^n \rightarrow \mathbb{R}$ is Lipschitz continuous, i.e. $\exists L_0(f) > 0$ s.t. $|f(X) - f(Y)| \leq L_0(f) \|X - Y\|_F$ holds for all $X, Y \in \mathbb{Q}^n$.

A. Convex case

Let $X^* \in \mathbb{Q}^n$ be a stationary point of (1), and $f^* \triangleq f(X^*)$. In addition, define $U_{0:k} \triangleq \text{diag}\{U_0, \dots, U_k\}$, a random matrix composed by i.i.d. variables $\{U_k\}_{k \in \mathbb{N}_0}$ associated with each iteration of the scheme.

We are now in a position to state the main result.

Theorem 1: Consider problem (1) with f convex. Suppose Assumption 1 holds, and let the sequence $\{X_k\}_{k \in \{0, \dots, N\}}$ be generated by the ZO-RMS iterates (2). Then, for any $N \geq 0$

$$\frac{1}{S_N} \sum_{k=0}^N h_k (\phi_k - f^*) \leq \mu L_0(f) \sqrt{\frac{n^2+n}{2}} + \frac{1}{S_N} \left(\frac{1}{2} \|X_0 - X^*\|_F^2 + \frac{1}{8} (n^4 + 2n^3 + 5n^2 + 4n) L_0^2(f) \sum_{k=0}^N h_k^2 \right), \quad (4)$$

where $S_N \triangleq \sum_{k=0}^N h_k$, $\phi_k \triangleq \mathbb{E}_{U_{0:k-1}}\{f(X_k)\}$ for $k \in \mathbb{N}$, and $\phi_0 \triangleq f(X_0)$.

Proof: See Appendix II. \blacksquare

A corollary from Theorem 1 can be obtained which provides expressions for μ , h_k , and N that ensure a given level of accuracy for the ZO-RMS algorithm.

Corollary 1: Let $\epsilon > 0$ be given. If μ and h_k are chosen such that

$$\mu \leq \frac{\epsilon}{L_0(f)\sqrt{2(n^2+n)}}, \quad (5a)$$

$$h_k = \frac{2\bar{r}}{L_0(f)\sqrt{n^4+2n^3+5n^2+4n\sqrt{N+1}}}, \quad (5b)$$

for $k \in \{0, \dots, N\}$, then, $\mathbb{E}_{U_{0:N-1}}\{f(\hat{X}_N)\} - f^* \leq \epsilon$ is guaranteed by the ZO-RMS algorithm in

$$N \geq \frac{L_0^2(f)\bar{r}^2}{\epsilon^2}(n^4+2n^3+5n^2+4n) \quad (6)$$

iterations, where \bar{r} is such that $\|X_0 - X^*\|_F \leq \bar{r}$.

Proof: Note that

$$\begin{aligned} & \mathbb{E}_{U_{0:N-1}}\{f(\hat{X}_N)\} - f^* \\ & \leq \mathbb{E}_{U_{0:N-1}}\left\{\frac{1}{S_N}\sum_{k=0}^N h_k(f(X_k) - f^*)\right\} \\ & \stackrel{(4)}{\leq} \mu L_0(f)\sqrt{\frac{n^2+n}{2}} \\ & + \frac{1}{S_N}\left(\frac{R^2}{2} + \frac{1}{8}(n^4+2n^3+5n^2+4n)L_0^2(f)\sum_{k=0}^N h_k^2\right). \end{aligned}$$

Suppose the number of steps N is fixed, we can choose μ and h_k as in (5) and further obtain the following bound

$$\begin{aligned} \mathbb{E}_{U_{0:N-1}}\{f(\hat{X}_N)\} - f^* & \leq \frac{\epsilon}{2} + \frac{\bar{r}^2}{S_N} \\ & = \frac{\epsilon}{2} + \frac{\bar{r}L_0(f)\sqrt{n^4+2n^3+5n^2+4n}}{2\sqrt{N+1}} \end{aligned}$$

Therefore, in order to satisfy $\mathbb{E}_{U_{0:N-1}}\{f(\hat{X}_N)\} - f^* \leq \epsilon$, we need $N \geq \frac{L_0^2(f)\bar{r}^2}{\epsilon^2}(n^4+2n^3+5n^2+4n) - 1$, concluding the proof. \blacksquare

The complexity bound (6) in Corollary 1 corresponds to the number of iterations in which the ZO-RMS algorithm guarantees a given accuracy of ϵ , provided the step size h_k and parameter μ are chosen as per (5). Certainly, the expressions for h_k , μ , and N depend on the Lipschitz constant $L_0(f)$ and \bar{r} which may be hard to obtain explicitly depending on the application. However, these can be numerically bounded as we explain further below in Section VI-A. Since the oracle (3) is random, the guarantees of the ZO-RMS algorithm hold in expectation. In essence, Corollary 1 provides sufficient conditions on the oracle's precision μ and step size h_k such that we can get sufficiently close to the optimal cost function f^* in N iterations, in average.

We now compare our complexity bound (6) to the bounds in [29] which are applicable to (1) (for convex f) after grouping the elements of the tuning matrix X into a single vector. Particularly, since X is symmetric, we group the lower triangular elements into a vector of size $n(n+1)/2$ (or often

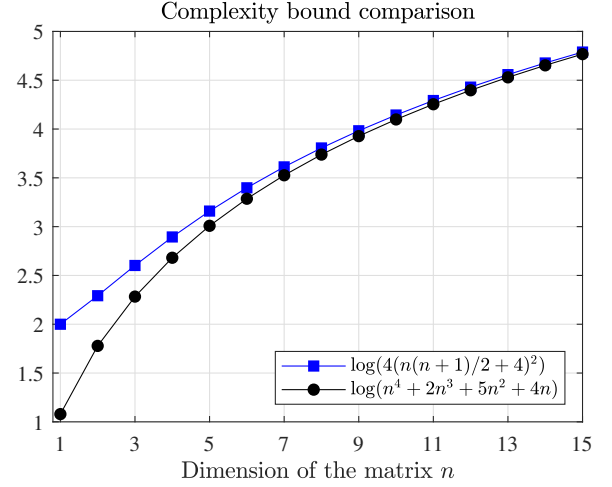


Fig. 1: Comparison of the complexity bounds (6) and (7), see black and blue lines, respectively.

called *half vectorization*), therefore the results in [29] apply and the following complexity bound holds (cf. Theorem 6 in [29])

$$N_{[29]} \geq \frac{4L_0^2(f)\bar{r}^2}{\epsilon^2}(n(n+1)/2+4)^2. \quad (7)$$

Fig. 1 depicts (6) and (7) for different values of n . We can see that our bound (6) is less conservative than (7) for all $n \in \mathbb{N}$, and tends to (7) as $n \rightarrow \infty$. We emphasise that our optimisation framework is directly tailored to matrix parameters, and the reduction in conservatism with respect to a vector approach such as [29] follows from exploiting the special structure of the underlying matrix space, namely the symmetry, and the use of iterates that are tailored to it. The reduction in conservativeness is significant, for instance, for 3×3 matrices (i.e. $n = 3$), $N_{[29]}/N_{\text{Cor.1}} = 2.08$. This translates in a sufficient condition that guarantees a level of accuracy in twice less iterations.

B. Non-convex case

Compared to the convex optimisation problem considered in Section III-A, the analysis for constrained non-convex problems is more challenging. Particularly, constrained convex problems typically focus on the optimality gap $\mathbb{E}\{f(x)\} - f^*$ to measure the convergence rate (as we do in Corollary 1). On the other hand, Nesterov in [29] provided complexity bounds for unconstrained non-convex problems using zeroth order iterates, where the optimisation variable is a vector in a subset of \mathbb{R}^n . In these unconstrained non-convex problems, the object $\mathbb{E}\{\|\nabla f(x)\|^2\}$ is the typical measure for stationarity. We emphasise that the bounds for unconstrained non-convex problems in [29] are not applicable to our constrained problem (1), since we need to search over the matrix space \mathbb{Q}^n and use iterates of the form (2) that project onto \mathbb{Q}^n . For constrained non-convex problems like ours, a fitting alternative to $\mathbb{E}\{\|\nabla f(x)\|^2\}$ is to consider the so-called *gradient mapping*,

see e.g. [33], [34], which is defined as follows

$$P_{\mathbb{Q}^n}(X, \mathcal{O}_\mu, h) \triangleq \frac{1}{h} \left[X - \pi_{\mathbb{Q}^n} \{ X - h \mathcal{O}_\mu(X, U) \} \right],$$

where \mathcal{O}_μ is the random zeroth order oracle in (3), and recall that $\pi_{\mathbb{Q}^n}$ denotes the Euclidean projection onto the closed convex set \mathbb{Q}^n . The natural interpretation of $P_{\mathbb{Q}^n}$ is that it represents the projected gradient, which offers a feasible update from the previous point X . The main goal in this section is to provide complexity bounds for the ZO-RMS algorithm (2) in terms of bounding $\mathbb{E} \{ \|P_{\mathbb{Q}^n}(X, \mathcal{O}_\mu, h)\|_F^2 \}$ when solving the constrained optimisation problem (1) with f non-convex.

Before stating our results, we impose an additional assumption.

Assumption 2: $\mathbb{E} \{ \|\mathcal{O}_\mu(X, U) - \nabla f_\mu(X)\|_F^2 \} \leq \sigma^2$, $\sigma > 0$, where f_μ is the Gaussian approximation of f defined in (16).

Assumption 2 essentially bounds the variance of the random oracle \mathcal{O}_μ . This assumption is often adopted in non-convex problems, see e.g. [34]. If constructing σ explicitly is of interest, we can construct it as follows. We note that \mathcal{O}_μ is an unbiased estimator of $\nabla f_\mu(X)$ (cf. (18)), i.e. $\mathbb{E} \{ \mathcal{O}_\mu(X, U) \} = \nabla f_\mu(X)$, then $\mathbb{E} \{ \|\mathcal{O}_\mu(X, U) - \nabla f_\mu(X)\|_F^2 \} \leq \mathbb{E} \{ \|\mathcal{O}_\mu(X, U)\|_F^2 \} \stackrel{\text{Thm.3}}{\leq} (1/4)L_0^2(f)(n^4 + 2n^3 + 5n^2 + 4n) \triangleq \sigma^2$.

The main result for the non-convex case is stated below. It can be seen as the non-convex counterpart of Theorem 1.

Theorem 2: Consider problem (1) with f non-convex. Suppose Assumptions 1 and 2 hold, and let the sequence $\{X_k\}_{k \in \{0, \dots, N\}}$ be generated by the ZO-RMS iterates (2). Then, for any $N \geq 0$,

$$\begin{aligned} & \frac{1}{S_N} \sum_{k=0}^N h_k \mathbb{E}_{U_{0:k}} \{ \|P_{\mathbb{Q}^n}(X_k, \mathcal{O}_\mu(X_k, U_k), h_k)\|_F^2 \} \\ & \leq \frac{1}{S_N} \left[f_\mu(X_0) - f^* + C(\mu) \sum_{k=0}^N h_k^2 \right] + \sigma^2, \end{aligned} \quad (8)$$

where $C(\mu) \triangleq \frac{L_0^3(f)}{4\mu} \sqrt{\frac{n^2+n}{2}} (n^4 + 2n^3 + 5n^2 + 4n)$.

Proof: See Appendix III. ■

The next corollary provides a choice of μ , h_k , and N that ensure a given level of accuracy for the ZO-RMS algorithm up to an error of order σ^2 .

Corollary 2: Let $\delta > 0$ and $\varepsilon > 0$ be given, and $h_k = h$ for all $k \in \{0, \dots, N\}$. If μ and h are chosen as

$$\mu = \frac{\varepsilon}{L_0(f)[(n^2+n)/2]^{1/2}}, \quad (9)$$

$$h = \left[\frac{8\varepsilon\bar{r}}{(N+1)L_0^3(f)(n^2+n)(n^4+2n^3+5n^2+4n)} \right]^{1/2}, \quad (10)$$

then $\mathbb{E} \{ \|P_{\mathbb{Q}^n}(X_D, \mathcal{O}_\mu(X_D, U_D), h)\|_F^2 \} \leq \delta + \sigma^2$ is guaranteed by the ZO-RMS algorithm in

$$N \geq \frac{L_0^5(f)\bar{r}(n^2+n)(n^4+2n^3+5n^2+4n)}{2\varepsilon\delta^2} \quad (11)$$

iterations, where

$$D \triangleq \arg \min_{k \in \{0, \dots, N\}} \|P_{\mathbb{Q}^n}(X_k, \mathcal{O}_\mu(X_k, U_k), h)\|_F.$$

Proof: Recall that f_μ is a smooth approximation of f , and the gap in this approximation can be upper bounded by $|f_\mu(X) - f(X)| \leq \mu L_0(f) \sqrt{\frac{n^2+n}{2}}$, see Lemma 2-(III) in the appendix. Then, to bound this gap by ε we need to choose μ as per (9). For this choice of μ , which we denote $\bar{\mu}$, $C(\bar{\mu}) = \frac{L_0^4(f)}{4\varepsilon} \left(\frac{n^2+n}{2} \right) (n^4 + 2n^3 + 5n^2 + 4n)$. For a constant step size, the right-hand size of (8) becomes

$$\frac{f_{\bar{\mu}}(X_0) - f^*}{(N+1)h} + C(\bar{\mu})h + \sigma^2 \leq \frac{L_0(f)\bar{r}}{(N+1)h} + C(\bar{\mu})h + \sigma^2.$$

Let $\rho(h) \triangleq \frac{L_0(f)\bar{r}}{(N+1)h} + C(\bar{\mu})h$. Minimising $\rho(h)$ in h leads to $h^* = \sqrt{L_0(f)\bar{r}/[C(\bar{\mu})(N+1)]}$, which corresponds to (10). For this choice of step size, we get $\rho(h^*) = \sqrt{4L_0(f)\bar{r}C(\bar{\mu})}/(N+1)$. Note that $\mathbb{E} \{ \|P_{\mathbb{Q}^n}(X_D, \mathcal{O}_\mu(X_D, U_D), h)\|_F^2 \} \leq \frac{1}{N+1} \sum_{k=0}^N \mathbb{E}_{U_{0:k}} \{ \|P_{\mathbb{Q}^n}(X_k, \mathcal{O}_\mu(X_k, U_k), h)\|_F^2 \} \leq \rho(h^*) + \sigma^2$. Then, we can guarantee that $\rho(h^*) \leq \delta$ in $4L_0(f)\bar{r}C(\bar{\mu})/\delta^2$ iterations, which corresponds to (11), completing the proof. ■

Our analysis shows that ZO-RMS for constrained non-convex problems can suffer an additional error of order σ^2 which does not appear in the convex case, see Corollary 1. This effect is consistent with the literature on constrained non-convex problems, where this effect has been reported for different algorithms, see e.g. [33]–[35].

IV. APPLICATION TO MPC TUNING: AIR-PATH CONTROL IN DIESEL ENGINES

In this section, we illustrate how to apply the proposed optimisation framework to tune MPC controllers in the context of diesel air-path control.

A. Diesel engine air-path model

A schematic representation of the diesel air-path is shown in Fig. 2. The dynamics of a diesel engine air-path are highly non-linear, see e.g. [36], and can be captured in the general form

$$x_{k+1} = \mathbf{f}(x_k, u_k, \theta), \quad (12a)$$

$$y_k = \mathbf{g}(x_k, u_k, \theta), \quad (12b)$$

where $x_k \in \mathbb{R}^{n_x}$, $y_k \in \mathbb{R}^{n_y}$, and $u_k \in \mathbb{R}^{n_u}$ are the state, output, and input of the system respectively, at time instant $k \in \mathbb{N}_0$. The engine operational space is typically parametrised by $\theta \triangleq (\omega_e, \dot{m}_f) \in \Theta \subseteq \mathbb{R}^2$, where ω_e denotes the engine speed, \dot{m}_f denotes the fuel rate, and Θ is the engine operating space defined as $\Theta \triangleq \{ \theta \in \mathbb{R}^2 : \omega_{e,\min} \leq \omega_e \leq \omega_{e,\max}, \dot{m}_{f,\min} \leq \dot{m}_f \leq \dot{m}_{f,\max} \}$, for some $\omega_{e,\min}, \omega_{e,\max}, \dot{m}_{f,\min}, \dot{m}_{f,\max} \in \mathbb{R}$. Given these highly non-linear dynamics, a common approach to control-oriented modelling of the air-path is to generate linear models trimmed at various operating points θ . Particularly, we follow [21] and use twelve models to approximate the operating space uniformly. That is, the engine operating space Θ is divided into twelve regions as per Fig. 3,

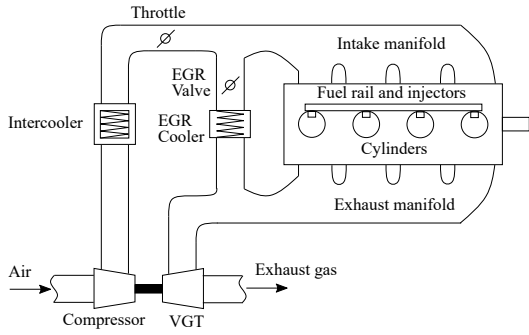


Fig. 2: A diesel engine air-path schematic.

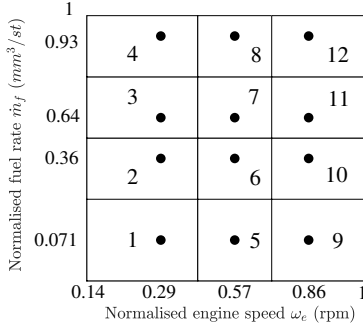


Fig. 3: Engine operational space divisions and the corresponding linearisation points.

with a linear model representing the engine dynamics in each region. For commercial in confidence purposes, we only show normalised axes in every plot. We emphasise that these regions are chosen to provide adequate coverage over the range of operating points encountered along a drive cycle. The control-oriented model state consists in the intake manifold pressure p_{im} (also known as boost pressure), the exhaust manifold pressure p_{em} , the compressor flow W_{comp} , and the EGR rate y_{EGR} . The input consists of the throttle position u_{thr} , EGR valve position u_{EGR} , and VGT valve position u_{VGT} . Lastly, the measured output is $(p_{\text{im}}, y_{\text{EGR}})$.

For a given operating point $(\omega_e^\sigma, \dot{m}_f^\sigma)$, $\sigma \in \{1, \dots, 12\}$, the engine control unit (ECU) applies certain steady-state actuator values. We denote by $\bar{u}^\sigma \in \mathbb{R}^3$, $\bar{x}^\sigma \in \mathbb{R}^4$, and $\bar{y}^\sigma \in \mathbb{R}^2$ the steady-state values of the input, state, and output respectively at each operating condition $(\omega_e^\sigma, \dot{m}_f^\sigma)$. Then, by following the system identification procedure detailed in [37], a linear representation of the non-linear diesel air-path (12) trimmed around the grid point $(\omega_e^\sigma, \dot{m}_f^\sigma)$ is given by

$$\tilde{x}_{k+1} = A^\sigma \tilde{x}_k + B^\sigma \tilde{u}_k, \quad (13a)$$

$$\tilde{y}_k = C^\sigma \tilde{x}_k + D^\sigma \tilde{u}_k, \quad (13b)$$

where $\sigma \in \{1, \dots, 12\}$, $\tilde{x}_k = x_k - \bar{x}^\sigma \in \mathbb{R}^4$ is the perturbed state around the corresponding steady-state \bar{x}^σ , $\tilde{u}_k = u_k - \bar{u}^\sigma \in \mathbb{R}^3$ is the perturbed input around the corresponding steady-state input \bar{u}^σ , and $\tilde{y}_k = y_k - \bar{y}^\sigma \in \mathbb{R}^2$ is the perturbed output around the corresponding steady-state output \bar{y}^σ .

B. MPC formulation

For each operating point $(\omega_e^\sigma, \dot{m}_f^\sigma)$, $\sigma \in \{1, \dots, 12\}$, and corresponding model (13) associated to that region, we formulate an MPC with augmented integrator (see e.g. [38], [39]). To this end, define the augmented state $\mathbf{x}_k = (\tilde{x}_k, e_k)$, where e_k is the integrator state with dynamics $e_{k+1} = -C^\sigma \tilde{x}_k + e_k$. Define the cost function as

$$J(\mathbf{x}_k, \mathbf{u}) \triangleq \mathbf{x}_{k+H}^\top P^\sigma \mathbf{x}_{k+H} + \sum_{i=0}^{H-1} (\mathbf{x}_{k+i}^\top Q^\sigma \mathbf{x}_{k+i} + \tilde{u}_{k+i}^\top R^\sigma \tilde{u}_{k+i}),$$

where $H \in \mathbb{N}$ is the prediction horizon, $\mathbf{u} \triangleq \{\tilde{u}_k, \tilde{u}_{k+1}, \dots, \tilde{u}_{k+H-1}\}$ is the sequence of control values applied over the horizon H , and $P^\sigma \in \mathbb{S}^6$, $Q^\sigma \in \mathbb{S}^6$, and $R^\sigma \in \mathbb{S}^3$ are real symmetric matrices containing the tuning weights. Particularly, we further assume that R^σ is positive definite, and that P^σ and Q^σ are positive semidefinite.

The corresponding MPC problem is stated below,

$$\begin{aligned} & \text{minimise} && J(\mathbf{x}_k, \mathbf{u}) \\ & \text{subject to} && \begin{cases} \mathbf{x}_{k+1} = \begin{bmatrix} A^\sigma & 0 \\ -C^\sigma & I \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} B^\sigma \\ 0 \end{bmatrix} \tilde{u}_k, \\ \tilde{x}(k) = x(k) - \bar{x}^\sigma, \\ S_x \mathbf{x}_i + S_u \tilde{u}_i \leq b_i, \forall i \in \{0, \dots, H-1\}, \\ S_H \mathbf{x}_H \leq b_H, \end{cases} \end{aligned}$$

where $S_x \in \mathbb{R}^{18 \times 6}$, $S_u \in \mathbb{R}^{18 \times 3}$, $b_i \in \mathbb{R}^{18}$, $S_H \in \mathbb{R}^{12 \times 6}$, and $b_H \in \mathbb{R}^{12}$ are given by the relevant state and input constraints. The solution to the above optimisation problem yields the optimising control sequence $\mathbf{u}^* = \{\tilde{u}_k^*, \tilde{u}_{k+1}^*, \dots, \tilde{u}_{k+H-1}^*\}$. The first element of the sequence \mathbf{u}^* is applied to the system and the whole process is repeated as k is incremented.

For transient operations between operating points, we use a switching LTI-MPC architecture so that the controllers are switched based on the current operating condition. As in [40], the switching LTI-MPC strategy selects the MPC controller at the nearest operating point to the current operating condition.

Under the above scenario, our goal is to use the ZO-RMS algorithm to tune the MPC weighting matrices $\{P^\sigma, Q^\sigma, R^\sigma\}$ in order to achieve a satisfactory tracking performance for a given prediction horizon $H \in \mathbb{N}$. Since there are twelve available controllers to tune, we could potentially use the ZO-RMS algorithm to tune all of them. However, since experiments are costly, we will focus on tuning only the controllers that have poor performance given an initial choice of tuning parameters. That is, we utilise an initial calibration denoted by $\{P_0^\sigma, Q_0^\sigma, R_0^\sigma\}$ for every controller $\sigma \in \{1, \dots, 12\}$, which we call *baseline controller*. The choice of $\{P_0^\sigma, Q_0^\sigma, R_0^\sigma\}$ may be based on model dynamics, experience, or using ad-hoc guidelines as per [11]. Then, by looking at a drive cycle response with the baseline controller, we detect the controllers that have poor performance. Let us denote by σ^* a controller with unsatisfactory performance and that we attempt to tune with the ZO-RMS algorithm. Therefore, in this context, we want to solve

$$\min_{P^{\sigma^*}, Q^{\sigma^*} \in \mathbb{S}_+^6, R^{\sigma^*} \in \mathbb{S}_+^3} g(y(P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*})), \quad (14)$$

where g is the tracking performance of the MPC defined in (15) below, which depends on a closed-loop response of interest y , and $\mathbb{S}_+^4 \subset \mathbb{S}^4$ and $\mathbb{S}_{++}^3 \subset \mathbb{S}^3$ denote the set of symmetric positive semi-definite matrices and the set of symmetric positive definite matrices, respectively.

As mentioned above, we use the tracking error as the measure of performance for the MPC controller, that is,

$$g(y(P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*})) \triangleq \frac{1}{\sqrt{M}} \sqrt{\sum_{k=0}^{M-1} \|y_k(P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*}) - y_k^{\text{ref}}\|_2^2}, \quad (15)$$

where $M \in \mathbb{N}$ is the experiment length, y_k^{ref} is the vector containing the boost pressure and EGR rate references, respectively, and $y_k(P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*})$ is the process measured output when the σ^* -th controller is using tuning parameters $P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*}$, and the rest of controllers are using the baseline parameters $\{P_0, Q_0, R_0\}$. We emphasise that for the simulations and experiments below we input normalised signals in (15) so that they are weighted equally.

C. Implementing the ZO-RMS algorithm

We now show how to implement the ZO-RMS algorithm for the MPC tuning problem (14). First note that (14) fits the general optimisation problem (1) with $X = \text{diag}\{P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*}\}$, $f(X) \triangleq g(y(X))$, and

$$\mathbb{Q}^n \triangleq \{X \in \mathbb{R}^{11 \times 11} : X = \text{diag}\{P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*}\}, \\ P^{\sigma^*}, Q^{\sigma^*} \in \mathbb{S}_+^4, R^{\sigma^*} \in \mathbb{S}_{++}^3\}.$$

Therefore, we can apply ZO-RMS to solve (14). The overall implementation of ZO-RMS in this context is depicted in Fig. 4. the ZO-RMS algorithm iteratively updates the MPC parameters $\{P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*}\}$ to minimise $f(P^{\sigma^*}, Q^{\sigma^*}, R^{\sigma^*})$, which is calculated from closed-loop response experiments carried out between the MPC and diesel engine.

It remains to show how to compute the oracle and projection in (2). Particularly, the oracle in this context is computed as per (3) with $U = \text{diag}\{U^P, U^Q, U^R\}$, where $U^P, U^Q \in \mathbb{G}^4$ and $U^R \in \mathbb{G}^3$ (see Definition 1). We emphasise that, at each iteration step, the oracle is computed once the entire closed-loop response is available. Particularly, two experiments are required to compute the oracle, one with parameters $\{P_k^{\sigma^*} + \mu U_k^P, Q_k^{\sigma^*} + \mu U_k^Q, R_k^{\sigma^*} + \mu U_k^R\}$ and one with $\{P_k^{\sigma^*}, Q_k^{\sigma^*}, R_k^{\sigma^*}\}$. Once the two closed-loop responses are finished, then ZO-RMS computes the oracle and next update $\{P_{k+1}^{\sigma^*}, Q_{k+1}^{\sigma^*}, R_{k+1}^{\sigma^*}\}$. Two new closed-loop experiments are then performed with these new controller parameters and the process continues iteratively in this fashion.

Since X is block diagonal, the projection onto \mathbb{Q}^n is computed as $\pi_{\mathbb{Q}^n}\{X\} = \text{diag}\{\pi_{\mathbb{S}_+^4}\{P\}, \pi_{\mathbb{S}_+^4}\{Q\}, \pi_{\mathbb{S}_{++}^3}\{R\}\}$, where $\pi_{\mathbb{S}_+^4}$ and $\pi_{\mathbb{S}_{++}^3}$ denote the Euclidean projection onto \mathbb{S}_+^4 and \mathbb{S}_{++}^3 , respectively. To compute them, we follow the well known results in [41]. That is, let $K = V\Lambda V^\top$ be the eigenvalue decomposition of a matrix K . Then, $\pi_{\mathbb{S}_+}\{K\} \triangleq V \max\{0, \Lambda\} V^\top$ and $\pi_{\mathbb{S}_{++}}\{K\} \triangleq V \max\{d, \Lambda\} V^\top$, $d > 0$.

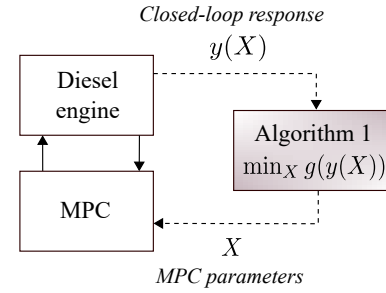


Fig. 4: Overall MPC tuning scheme.

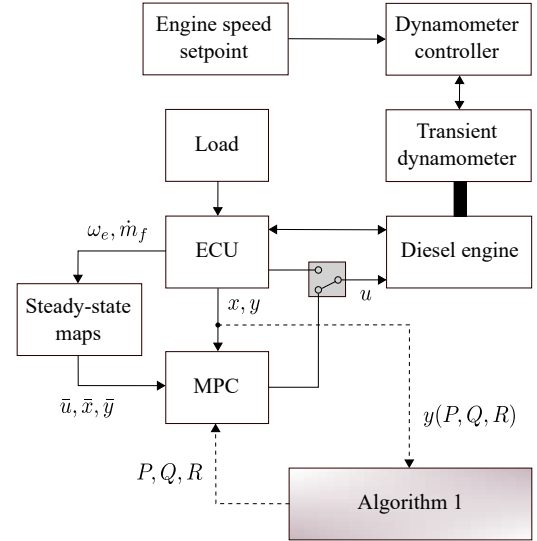


Fig. 5: Block diagram of the experimental setting.

V. SIMULATION STUDY

We now perform numerical simulations to show the advantage of our proposed MPC tuning framework with respect to other available gradient-free algorithms in the literature. Particularly, we compare our approach to the dividing rectangles (DIRECT) algorithm [42], particle swarm optimisation algorithms [17], and Bayesian optimisation algorithms [22], [23]. DIRECT is a sample-based global optimisation method for Lipschitz continuous functions defined over multidimensional domains. It partitions the space into hyperrectangles, each with a sampled point at its centre. The cost function is evaluated at these centrepieces, and then the algorithm keeps sampling and dividing into rectangles until the iteration limit has been reached. PSO algorithms solve the problem by having a population of so-called particles which move along the search space according to the updating rules of their position and velocity. The movement of the particles are guided by their own best known position as well as the entire swarm's best known position in the search space. At each iteration, the cost function is evaluated for every particle in the swarm. Lastly, BO seeks to identify the optimal tuning parameters by strategically exploring and exploiting the parameter space. Exploration aims to evaluate the objective at points in the decision space with the goal of improving the accuracy of a surrogate model of the objective function, while exploitation aims to use the surrogate model to identify decisions that

reduce (or increase) the objective function [22].

Intuition suggests that methods such as DIRECT and PSO would perform many cost function evaluations in order to find the optimal value, since they rely on the number of rectangles/centre points (DIRECT) and particles (PSO). We demonstrate that this is indeed true for the MPC tuning problem described in Section IV.

We perform simulations using the tuning scheme from Fig. 4 in which the diesel engine block is simulated via a high-fidelity model, see e.g. [21]. This model is physics-based with the form (12), in which the parameters have been obtained from system identification experiments conducted at Toyota's Higashi-Fuji Technical Centre. The specific equations and parameters are not included due to commercial in confidence purposes. We focus on tuning controller $\sigma^* = 6$ (see Fig. 3), over a segment of the NEDC, which is called the urban drive cycle (UDC). We focus on tuning diagonal matrices $\{P^6, Q^6, R^6\}$ with positive elements so that we can compare to vector-valued methods like DIRECT, PSO, and BO. We emphasise another contribution of the proposed optimisation framework is that it allows the direct tuning of the weighting matrices so that they satisfy the required constraints of symmetricity and positive (semi) definiteness. For these simulations, we do not include the augmented integrator states.

The cost function is the tracking error as per (15), and to assess the complexity of each algorithm we will use the total number of cost function evaluations that takes to achieve the optimal cost function value up to a certain tolerance. The initial calibration is $P_0^6 = Q_0^6 = \text{diag}\{0.01, 0, 0.2, 0.01\}$ and $R_0^6 = 10^{-5}I_{3 \times 3}$, and the number of tuning parameters is eleven. The ZO-RMS algorithm uses step size $h_k = 5 \times 10^{-3}/\sqrt{k+1}$ and oracle's precision $\mu = 8 \times 10^{-5}$ ($L_0(f) = 6, \epsilon = 0.008, \bar{r} = 2, n = 11$). For the PSO algorithm we have picked 40 particles as suggested by the literature [17]. For DIRECT, PSO, and BO, we use the decision variable range $[0.01, 5]$.

One realisation run of all four algorithms is summarised in Table I, where f_{opt} denotes the optimal value of the cost function achieved by each algorithm. Next we show the number of cost function evaluations, followed by the total execution time of each algorithm. We can see that in order to achieve a similar value of f_{opt} (up to a 0.0002 difference), the ZO-RMS algorithm performs 20 function evaluations, compared to 66 for DIRECT, 200 for PSO, and 28 for BO. We emphasise that each function evaluation in this context corresponds to a closed-loop experiment, and thus having to perform many of these is intractable in practice. We can see that BO requires only a few more cost function evaluations than ZO-RMS. It can also be observed that DIRECT, PSO, and BO take longer to execute in comparison to ZO-RMS.

Since PSO, BO, and ZO-RMS are stochastic, we also performed a Monte Carlo simulation of 100 realisations in order to compare the three of them more accurately. These results are listed in Table II. We can see that all achieve a similar optimal value for f_{opt} in average, but it takes about 199 cost function evaluations in average for PSO to achieve this, and only 22 for BO and 20 for ZO-RMS which makes them more efficient for applications in which the plant is in the loop. Overall, BO exhibits comparable performance with respect to

TABLE I: Comparison between available algorithms.

Algorithm	f_{opt}	N° of cost fcn. evaluations	Execution time
DIRECT	0.0478	66	~0.6 hours
PSO	0.0476	200	~1.6 hours
BO	0.0477	28	~0.5 hours
ZO-RMS	0.0476	20	~0.3 hours

TABLE II: Monte Carlo simulation (100 realisations).

Algorithm	$\mathbb{E}\{f_{\text{opt}}\}$	$\text{Var}\{f_{\text{opt}}\}$	$\mathbb{E}\{N^\circ \text{ of cost fcn. evaluations}\}$
PSO	0.0472	5×10^{-8}	199
BO	0.0477	7×10^{-8}	22
ZO-RMS	0.0496	9×10^{-6}	20

ZO-RMS in terms of cost function evaluations and execution time; however, as mentioned in the introduction, these methods only work for vector decision variables and the PSD matrix structure of the MPC tuning matrices is not preserved by the algorithm. They also lack theoretical guarantees in this context.

VI. EXPERIMENTS

The experimental testing of the proposed MPC tuning framework was performed at Toyota's Higashi-Fuji Technical Center in Susono, Japan. The test bench is equipped with a diesel engine, a transient dynamometer, and a dSPACE DS1006 processor board, which is in turn interfaced to the ECU. The block diagram in Fig. 5 depicts the overall experimental setting. The ECU logs sensor data from the engine and transmits the current state information to the controller. In addition, the ECU directly controls all engine sub-systems, however, the ECU commands for the three actuators (throttle, EGR valve, and VGT) can be overridden by the MPC commands through a virtual switch in the ControlDesk interface. The proposed tuning architecture is implemented iteratively in-the-loop as per Fig. 4, in which the diesel engine block is the real engine described above. Particularly, the switched MPC is implemented in real-time on the dSPACE board to generate the required closed-loop drive cycle responses, and the computations in the ZO-RMS algorithm are performed by MATLAB, i.e. oracle, projections, and next set of parameters.

A. Choice of algorithm parameters

To run ZO-RMS we need to choose the parameters μ and h_k . Corollary 1 provides a choice for these parameters so that a given level of accuracy is achieved by the algorithm in N iterations as per (6). However, this choice depends on the Lipschitz constant $L_0(f)$, and the bound \bar{r} . These can be numerically bounded depending on the application. Since we apply the proposed optimisation framework to MPC tuning over a real engine, it is not possible to find an explicit expression for the cost function in terms of P, Q, R . Instead, a common approach is to adopt experimental optimisation methods such as the ones in [43], [44] to ensure that the Lipschitz condition is at least satisfied in the experimental data. Particularly, since the oracle \mathcal{O}_μ uses values of the cost function at two different points, we can store these values and use a consistency-check algorithm to see if the Lipschitz condition is verified (see e.g. [43, Section 4.3]). This could be




Fig. 6: Engine response over the middle UDC segment using the initial parameters $\{P_0, Q_0, R_0\}$ for all twelve controllers.

done by performing high-fidelity simulations of the closed-loop to compute the oracle at different points, and pick the largest constant that satisfies the Lipschitz condition as a first estimate. Later on in the experiments, we can further adjust this first estimate accordingly, and verify whether the bound is satisfied for the collected experimental measurements. Similarly, \bar{r} can be estimated with the stored data of X_0 and \hat{X}_N from multiple simulations so that $\|X_0 - X^*\| \leq \bar{r}$ holds for the experimental space. In this paper, we initially picked conservative estimates from simulations, and we later refine these heuristically during experimental testing. Online estimation of the Lipschitz constant $L_0(f)$ can improve the aforementioned methods and it is considered as a future direction.

In what follows, we tune three out of the twelve MPC controllers depicted in Fig. 3 over three different drive cycle segments. Particularly, we tune: controller $\sigma^* = 9$ over the middle segment of the UDC (49–117[s]), and we call it UDC2; controller $\sigma^* = 5$ over the first 100 seconds of the WLTC, which call WLTC1; and controller $\sigma^* = 6$ over the last segment of the UDC (117–195[s]), which we call UDC3.

For all the experiments below we consider an initial calibration $Q_0 = \text{diag}\{\bar{Q}_0, \bar{Q}_{e,0}\}$, $\bar{Q}_0 = \text{diag}\{0.01, 0.01, 0.2, 1\}$, $\bar{Q}_{e,0} = 0.00001 \times \text{diag}\{1, 100\}$, $P_0 = Q_0$, $R_0 = I_{3 \times 3}$ for all twelve controllers, where \bar{Q}_0 and $\bar{Q}_{e,0}$ respectively relate to the engine state and integrator state. In addition, all the drive cycle references are generated by the ECU based on a driver model. These models are responsible for converting the drive cycle vehicle speed reference into reference operating points in terms of engine speed and load.

B. Controller $\sigma^* = 9$ over UDC2

The engine output response to the UDC2 segment with the baseline controller is given in Fig. 6. The grey area illustrates when controller $\sigma^* = 9$ is active, and we consider its tracking performance as unsatisfactory (we draw the attention to this with the boxed area). We would like to improve the tracking performance for this controller by using ZO-RMS.

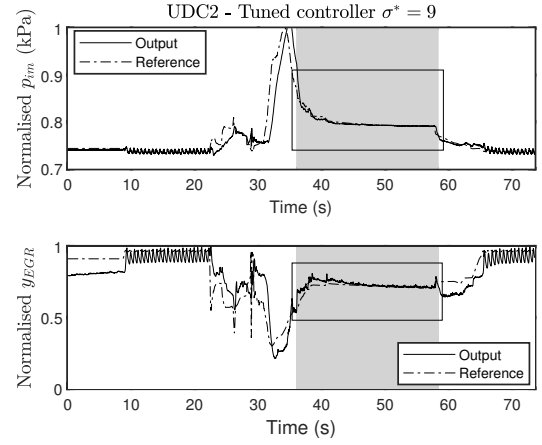


Fig. 7: Engine response over the middle UDC segment using the tuned $\{\hat{P}^9, \hat{Q}^9, \hat{R}^9\}$ for controller six, and $\{P_0, Q_0, R_0\}$ for the remaining eleven controllers.

The parameters used in this experiment are¹ $\mu = 0.003 \times 10^{-5}$, $h_k = 10^{-6}/\sqrt{k+1}$. ($L_0(f) = 2.37 \times 10^4$, $\epsilon = 0.01$, $\bar{r} = 1.1$, $n = 9$).

In this experiment, we consider $Q^9 = \text{diag}\{\bar{Q}^9, \bar{Q}_e^9\}$ and use ZO-RMS to tune $\{\bar{Q}^9, \bar{Q}_e^9, R^9\}$, whilst P^9 is constructed by solving the discrete-time algebraic Riccati equation (DARE) [45], $P^9 = \text{dare}\left(\begin{bmatrix} A^9 & 0 \\ -C^9 & I \end{bmatrix}, \begin{bmatrix} B^9 \\ 0 \end{bmatrix}, \begin{bmatrix} Q^9 & 0 \\ 0 & \bar{Q}_e^9 \end{bmatrix}, R^9\right)$ at each iteration. We emphasise that the proposed tuning framework provides enough flexibility to either tune a single matrix, or all of the MPC matrices depending on the particular application. For instance, for this drive cycle segment, tuning the integrator state matrix \bar{Q}_e helped significantly in improving tracking performance, but for the experiments further below it was not necessary to tune this matrix and it was thus fixed.

The resulting optimal tuning matrices for $N = 11$ iterations are given by

$$\hat{Q}^9 = \begin{bmatrix} 0.0611 & 0.0277 & 0.0742 & 0.1389 \\ 0.0277 & 0.1393 & -0.0554 & 0.0318 \\ 0.0742 & -0.0554 & 0.2177 & 0.0963 \\ 0.1389 & 0.0318 & 0.0963 & 1.0701 \end{bmatrix},$$

$$\hat{R}^9 = \begin{bmatrix} 1.8037 & -0.0924 & 0.4081 \\ -0.0924 & 1.5841 & 0.2252 \\ 0.4081 & 0.2252 & 1.2879 \end{bmatrix},$$

$$\hat{Q}_e^9 = \begin{bmatrix} 0.0013 & 0.0014 \\ 0.0014 & 0.0041 \end{bmatrix}.$$

The engine output response over the UDC2 using the above matrices is shown in Fig. 7. The tracking performance has significantly improved in the region where controller nine is acting. Particularly, the tuned matrices provide an improvement of performance² of 16.2% with only eleven iterations.

¹We note that for this particular experiment we used a very conservative estimate of the Lipschitz constant based on different experimental tests.

²We compute the percentage of improvement with respect to the baseline controller as: $100 \times [f(P_0, Q_0, R_0) - f(\hat{P}^{\sigma^*}, \hat{Q}^{\sigma^*}, \hat{R}^{\sigma^*})] / f(P_0, Q_0, R_0)$.




Fig. 8: Engine response over the first WLTC segment using the initial parameters $\{P_0, Q_0, R_0\}$ for all twelve controllers.

C. Controller $\sigma^* = 5$ over WLTC1

The engine output response to the WLTC1 segment with the baseline controller is given in Fig. 8. The grey area illustrates when controller $\sigma^* = 5$ is active, and we have boxed the areas with unsatisfactory performance that we would like to improve. The parameters used in this experiment are $\mu = 2.5 \times 10^{-6}$, $h_k = 10^{-6}/\sqrt{k+1}$ ($L_0(f) = 3450$, $\epsilon = 0.1$, $\bar{r} = 0.1$, $n = 7$). Similar to the previous experiment, the matrix P^5 is constructed by solving the corresponding DARE, and we consider $Q^5 = \text{diag}\{\bar{Q}^5, \bar{Q}_e^5\}$, but now the integrator state matrix \bar{Q}_e^5 is not tuned but kept equal to the initial value $\bar{Q}_{e,0} = 0.00001 \times \text{diag}\{1, 100\}$.

Consequently, the tuning parameters are $\{\bar{Q}^5, R^5\}$, and their optimal values for $N = 11$ iterations are given by

$$\hat{Q}^5 = \begin{bmatrix} 0.0101 & -0.0078 & -0.0213 & 0.0102 \\ -0.0078 & 0.0273 & 0.0061 & 0.0170 \\ -0.0213 & 0.0061 & 0.1760 & 0.0004 \\ 0.0102 & 0.0170 & 0.0004 & 0.9967 \end{bmatrix},$$

$$\hat{R}^5 = \begin{bmatrix} 0.9992 & -0.0066 & -0.0017 \\ -0.0066 & 0.9839 & 0.0016 \\ -0.0017 & 0.0016 & 0.9806 \end{bmatrix}.$$

The engine output response over the WLTC1 using the above matrices is shown in Fig. 9. The tracking performance has improved in the region where controller nine is acting. Particularly, the tuned matrices provide an improvement of performance of 22.67% with only eleven iterations, as illustrated by the boxed regions in Fig. 9. We can observe that the tracking has clearly improved in the boxed areas, but we can see that in the grey area 40-50s in Fig. 8 and 9, the EGR tracking got slightly deteriorated. However, we emphasise that the cost function being minimised is the overall tracking performance considering every grey area where controller $\sigma^* = 5$ is acting. This is considered satisfactory since the overall tracking improvement in this case was 22.67%. Different cost functions or tuning approaches could be potentially used to tackle each region individually.




Fig. 9: Engine response over the first WLTC segment using the tuned $\{\hat{P}^5, \hat{Q}^5, \hat{R}^5\}$ for controller five, and $\{P_0, Q_0, R_0\}$ for the remaining eleven controllers.

D. Controller $\sigma^* = 6$ over UDC3

The engine output response to the UDC3 segment with the baseline controller is given in Fig. 10. The grey area illustrates when controller $\sigma^* = 6$ is active, and we have boxed the areas with unsatisfactory performance. The parameters used in this experiment are $\mu = 2.5 \times 10^{-6}$, $h_k = 10^{-5}/\sqrt{k+1}$ ($L_0(f) = 2450$, $\epsilon = 0.1$, $\bar{r} = 1.2$, $n = 11$). The integrator state matrix \bar{Q}_e was not tuned but kept equal to the initial value $\bar{Q}_{e,0} = 0.00001 \times \text{diag}\{1, 100\}$. As opposed to the previous two experiments, we do include P^5 in the tuning process for this experiment. Specifically, we construct it as $P^6 = \text{diag}\{\bar{P}^6, \bar{Q}_{e,0}\}$ and tune \bar{P}^6 .

Consequently, the tuning parameters are $\{\bar{P}^6, \bar{Q}^6, R^6\}$, and their optimal values for $N = 8$ iterations are given by

$$\hat{P}^6 = \begin{bmatrix} 0.0502 & -0.0117 & -0.0398 & 0.0057 \\ -0.0117 & 0.0102 & 0.0146 & 0.0118 \\ -0.0398 & 0.0146 & 0.2003 & -0.0426 \\ 0.0057 & 0.0118 & -0.0426 & 1.0022 \end{bmatrix},$$

$$\hat{Q}^6 = \begin{bmatrix} 0.4699 & 0.0708 & 0.0192 & 0.0901 \\ 0.0708 & 0.2329 & -0.1597 & -0.1684 \\ 0.0192 & -0.1597 & 0.3459 & 0.0016 \\ 0.0901 & -0.1684 & 0.0016 & 1.1911 \end{bmatrix},$$

$$\hat{R}^6 = \begin{bmatrix} 1.1227 & -0.1419 & 0.0305 \\ -0.1419 & 1.0621 & 0.0534 \\ 0.0305 & 0.0534 & 1.2843 \end{bmatrix}.$$

The engine output response over the UDC3 using the above matrices is shown in Fig. 11. We can see that the tracking performance has improved in the region where controller nine is acting. Particularly, the tuned matrices provide an improvement of performance of 7.73% with only eight iterations.

The theoretical results in Section III are sufficient conditions on the step size h_k and oracle's precision μ that ensure a certain level of accuracy in a fixed number of iterations. We note that the complexity bounds presented in this paper hold for any problem that fits the optimisation framework surrounding (1). Depending on the application, these bounds




Fig. 10: Engine response over the third UDC segment using the initial parameters $\{P_0, Q_0, R_0\}$ for all twelve controllers.

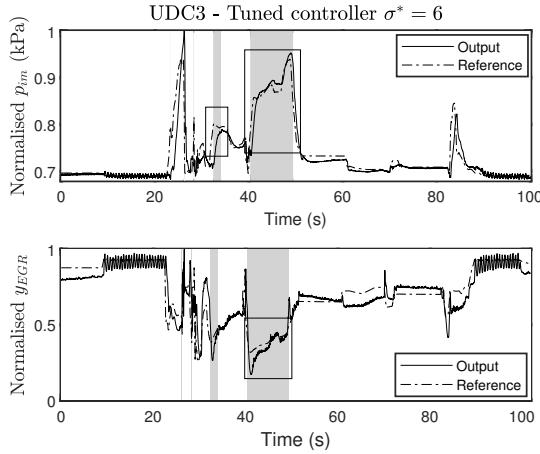


Fig. 11: Engine response over the third UDC segment using the tuned $\{\hat{P}^6, \hat{Q}^6, \hat{R}^6\}$ for controller six, and $\{P_0, Q_0, R_0\}$ for the remaining eleven controllers.

could become more or less conservative, as we observed in the air-path control case study. In fact, since these are sufficient conditions, we observed that the experimental performance of the algorithm showed tracking error improvement only in a few iterations. Nevertheless, the theoretical results serve as a design tool that guide the choice of algorithm parameters as opposed to a heuristic choice of parameters.

Remark 1: Overall, we observed that the proposed approach successfully provides improved tracking performance over three different drive cycle segments, i.e. different data sets, which sheds some light on the robustness of the method. However, an interesting future work direction is to provide a theoretical study on robustness guarantees.

VII. CONCLUSION

This paper provided an optimisation framework with theoretical guarantees for the minimisation of non-smooth and possibly non-convex cost functions with matrix parameters. We applied the proposed algorithm to tune MPCs in the context of air-path control in diesel engines, which is then

validated by experimental testing. The algorithm provides improvement of performance with a few iterations, which is demonstrated in different engine drive cycles. Therefore, it creates potential for the development of efficient tuning tools for advanced controllers (and potentially retuning online), even though the theoretical complexity bounds may be large depending on the application.

Future work includes exploring algorithms over other matrix manifolds with practical applications, testing of the proposed algorithms in stochastic drive cycles, robustness guarantees, and online estimation/adaptation of Lipschitz constant $L_0(f)$.

APPENDIX I TECHNICAL RESULTS

Appendix I provides the preliminary results needed for Appendices II and III.

Consider a function $f : \mathbb{Q}^n \rightarrow \mathbb{R}$, and let us define the Gaussian approximation of f as

$$f_\mu(X) \triangleq \frac{1}{\kappa} \int_{\mathbb{G}^n} f(X + \mu U) e^{-\frac{1}{2} \|U\|_F^2} dU. \quad (16)$$

We first introduce the following definition.

Definition 2: Let f be a convex function. A matrix G is called a subgradient of function f at $X \in \mathbb{R}^{n \times n}$ if for any $Y \in \mathbb{R}^{n \times n}$ we have $f(Y) \geq f(X) + \langle G, Y - X \rangle_F$. The set of all subgradients of f at $X \in \mathbb{R}^{n \times n}$ is called subdifferential and is denoted by $\partial f(X)$.

The following lemma holds for f_μ .

Lemma 2: (I) If f is convex and $G \in \partial f(X)$, then

$$f_\mu(X) \geq \frac{1}{\kappa} \int_{\mathbb{G}^n} (f(X) + \mu \langle G, U \rangle_F) e^{-\frac{1}{2} \|U\|_F^2} dU = f(X). \quad (17)$$

(II) The gradient matrix ∇f_μ satisfies

$$\begin{aligned} \nabla f_\mu(X) &= \frac{1}{\kappa} \int_{\mathbb{G}^n} \frac{f(X + \mu U) - f(X)}{\mu} e^{-\frac{1}{2} \|U\|_F^2} U dU. \end{aligned} \quad (18)$$

(III) Let f be Lipschitz continuous as per Assumption 1, then

$$|f_\mu(X) - f(X)| \leq \mu L_0(f) \sqrt{\frac{n^2 + n}{2}}, \quad X \in \mathbb{Q}^n. \quad (19)$$

(IV) Under Assumption 1, f_μ has Lipschitz continuous gradient, i.e. $|f_\mu(Y) - f_\mu(X) - \langle \nabla f_\mu(X), Y - X \rangle_F| \leq \frac{1}{2} L_1(f_\mu) \|X - Y\|_F^2$, for all $X, Y \in \mathbb{Q}^n$, with $L_1(f_\mu) = (2L_0(f)/\mu) \sqrt{(n^2 + n)/2}$.

Proof:

(I) We prove this statement as follows.

$$\begin{aligned} f_\mu(X) &\stackrel{(16)}{=} \frac{1}{\kappa} \int_{\mathbb{G}^n} f(X + \mu U) e^{-\frac{1}{2} \|U\|_F^2} dU \\ &\stackrel{\text{Def. 2}}{\geq} \frac{1}{\kappa} \int_{\mathbb{G}^n} (f(X) + \langle G, \mu U \rangle_F) e^{-\frac{1}{2} \|U\|_F^2} dU \\ &= f(X) + \frac{1}{\kappa} \int_{\mathbb{G}^n} \langle G, \mu U \rangle_F e^{-\frac{1}{2} \|U\|_F^2} dU \\ &= f(X) + \mu \text{Tr} \left\{ \frac{1}{\kappa} \int_{\mathbb{G}^n} G^\top U e^{-\frac{1}{2} \|U\|_F^2} dU \right\} \\ &= f(X), \end{aligned}$$

where the last equality follows from $\mathbb{E}_U\{U\} = 0_{n \times n}$.

(II) Define $Y \triangleq X + \mu U$, and note that by the change of variable formula for multivariate integrals, $dU = (1/\mu^{\bar{n}})dY$, where $\bar{n} = n(n+1)/2$. Consequently,

$$f_\mu(X) = \frac{1}{\kappa} \int_{\mathbb{G}^n} f(Y) e^{-\frac{1}{2\mu^2} \|Y-X\|_F^2} \times \frac{1}{\mu^{\bar{n}}} dY.$$

Then,

$$\begin{aligned} \nabla f_\mu(X) &= \frac{1}{\mu^{\bar{n}} \kappa} \int_{\mathbb{G}^n} f(Y) e^{-\frac{1}{2\mu^2} \|Y-X\|_F^2} \\ &\quad \times -\frac{1}{2\mu^2} \times \frac{\partial \text{Tr}\{(Y-X)^2\}}{\partial X} dY \\ &= \frac{1}{\mu^{\bar{n}+2} \kappa} \int_{\mathbb{G}^n} f(Y) e^{-\frac{1}{2\mu^2} \|Y-X\|_F^2} (Y-X) dY \\ &= \frac{1}{\mu \kappa} \int_{\mathbb{G}^n} f(X + \mu U) e^{-\frac{1}{2} \|U\|_F^2} U dU \\ &= \frac{1}{\kappa} \int_{\mathbb{G}^n} \frac{f(X + \mu U) - f(X)}{\mu} e^{-\frac{1}{2} \|U\|_F^2} U dU, \end{aligned}$$

where the last equality follows from $\mathbb{E}_U\{U\} = 0_{n \times n}$.

(III) By definition of f_μ in (16) we have

$$\begin{aligned} f_\mu(X) - f(X) &= \frac{1}{\kappa} \int_{\mathbb{G}^n} (f(X + \mu U) - f(X)) e^{-\frac{1}{2} \|U\|_F^2} dU. \end{aligned}$$

Then, by Lipschitz continuity of f (cf. Assumption 1),

$$\begin{aligned} |f_\mu(X) - f(X)| &\leq \mu L_0(f) \frac{1}{\kappa} \int_{\mathbb{G}^n} \|U\|_F e^{-\frac{1}{2} \|U\|_F^2} dU \\ &\leq \mu L_0(f) \sqrt{\frac{n^2+n}{2}}, \end{aligned}$$

where the last inequality follows from Lemma 1.

(IV) Indeed, for all $X, Y \in \mathbb{Q}^n$ we have from (18) and Assumption 1 that

$$\begin{aligned} \|\nabla f_\mu(X) - \nabla f_\mu(Y)\|_F &\leq \frac{2L_0(f)}{\kappa \mu} \int_{\mathbb{G}^n} \|U\|_F e^{-\frac{1}{2} \|U\|_F^2} dU \end{aligned}$$

and the result follows from Lemma 1. \blacksquare

For the random oracle in (3), we have the following bound.
Theorem 3: If f satisfies Assumption 1, then

$$\mathbb{E}_U\{\|\mathcal{O}_\mu(X, U)\|_F^2\} \leq \frac{1}{4} L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n). \quad (20)$$

Proof: Note that $\|\mathcal{O}_\mu(X, U)\|_F^2 = \text{Tr}\{\mathcal{O}_\mu(X, U)^\top \mathcal{O}_\mu(X, U)\} = \frac{(f(X+\mu U) - f(X))^2}{\mu^2} \|U\|_F^2$. Then, $\mathbb{E}_U\{\|\mathcal{O}_\mu(X, U)\|_F^2\} \leq \frac{1}{\mu^2} \mathbb{E}_U\{L_0^2(f) \|\mu U\|_F^2 \|U\|_F^2\} = L_0^2(f) \mathbb{E}_U\{\|U\|_F^4\}$. The proof thus follows from Lemma 1. \blacksquare

APPENDIX II PROOF OF THEOREM 1

We first note that the projection $\pi_{\mathbb{Q}^n}$ in ZO-RMS satisfies

$$\|\pi_{\mathbb{Q}^n}\{X\} - Y\|_F \leq \|X - Y\|_F, \quad (21)$$

for all $Y \in \mathbb{Q}^n$.

Now, let point X_k with $k \in \mathbb{N}$ be generated after k iterations of ZO-RMS, and define $r_k \triangleq \|X_k - X^*\|_F$. Then,

$$\begin{aligned} r_{k+1}^2 &= \|\pi_{\mathbb{Q}^n}\{X_k - h_k \mathcal{O}_\mu(X_k, U_k)\} - X^*\|_F^2 \\ &\stackrel{(21)}{\leq} \|X_k - h_k \mathcal{O}_\mu(X_k, U_k) - X^*\|_F^2 \\ &= r_k^2 - 2h_k \langle X_k - X^*, \mathcal{O}_\mu(X_k, U_k) \rangle_F \\ &\quad + h_k^2 \|\mathcal{O}_\mu(X_k, U_k)\|_F^2. \end{aligned}$$

Then, taking expectation

$$\begin{aligned} \mathbb{E}_{U_k}\{r_{k+1}^2\} &\stackrel{(20)}{\leq} r_k^2 - 2h_k \mathbb{E}_{U_k}\{\langle X_k - X^*, \mathcal{O}_\mu(X_k, U_k) \rangle_F\} \\ &\quad + \frac{1}{4} h_k^2 L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n) \\ &= r_k^2 - 2h_k \text{Tr}\{[X_k - X^*]^\top \mathbb{E}_{U_k}\{\mathcal{O}_\mu(X_k, U_k)\}\} \\ &\quad + \frac{1}{4} h_k^2 L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n). \end{aligned}$$

From (18), we get $\mathbb{E}_{U_k}\{\mathcal{O}_\mu(X_k, U_k)\} = \nabla f_\mu(X_k)$. Therefore,

$$\begin{aligned} \mathbb{E}_{U_k}\{r_{k+1}^2\} &\leq r_k^2 - 2h_k \langle X_k - X^*, \nabla f_\mu(X_k) \rangle_F \\ &\quad + \frac{1}{4} h_k^2 L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n) \\ &\stackrel{(a)}{\leq} r_k^2 - 2h_k [f(X_k) - f_\mu(X^*)] \\ &\quad + \frac{1}{4} h_k^2 L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n), \end{aligned}$$

where (a) in the last inequality is shown as follows,

$$\begin{aligned} f_\mu(X^*) &\stackrel{\text{Def.2}}{\geq} f_\mu(X_k) + \langle \nabla f_\mu(X_k), X^* - X_k \rangle_F, \\ &\stackrel{(17)}{\geq} f(X_k) - \langle X_k - X^*, \nabla f_\mu(X_k) \rangle_F. \end{aligned}$$

Taking now the expectation in $U_{0:k-1}$, we obtain

$$\begin{aligned} \mathbb{E}_{U_{0:k}}\{r_{k+1}^2\} &\leq \mathbb{E}_{U_{0:k-1}}\{r_k^2\} - 2h_k (\phi_k - f_\mu(X^*)) \\ &\quad + \frac{1}{4} h_k^2 L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n). \end{aligned}$$

Moreover, note that $f_\mu(X^*) \stackrel{(17)}{\geq} f(X^*)$, thus (19) implies that $f_\mu(X^*) \leq f^* + \mu L_0(f) \sqrt{(n^2+n)/2}$. Consequently,

$$\begin{aligned} \mathbb{E}_{U_{0:k}}\{r_{k+1}^2\} &\leq \mathbb{E}_{U_{0:k-1}}\{r_k^2\} - 2h_k (\phi_k - f^*) \\ &\quad + 2h_k \mu L_0(f) \sqrt{\frac{n^2+n}{2}} \\ &\quad + \frac{1}{4} h_k^2 L_0^2(f) (n^4 + 2n^3 + 5n^2 + 4n). \quad (22) \end{aligned}$$

We can now iterate (22) from $k=0$ to $k=N$ and get

$$\begin{aligned} \mathbb{E}_{U_{0:N}}\{r_{N+1}^2\} &\leq r_0^2 - 2 \sum_{k=0}^N h_k (\phi_k - f^*) \\ &\quad + 2S_N \mu L_0(f) \sqrt{\frac{n^2+n}{2}} \\ &\quad + \frac{1}{4} (n^4 + 2n^3 + 5n^2 + 4n) L_0^2(f) \sum_{k=0}^N h_k^2. \quad (23) \end{aligned}$$

Since $\mathbb{E}_{U_{0:N}}\{r_{N+1}^2\} \geq 0$, from (23) we immediately get (4), completing the proof. \blacksquare

APPENDIX III PROOF OF THEOREM 2

Define $g_k \triangleq P_{\mathbb{Q}^n}(X_k, \mathcal{O}_\mu(X_k, U_k), h_k)$. From Lemma 2(IV), we have

$$f_\mu(X_{k+1}) \leq f_\mu(X_k) - h_k \langle \nabla f_\mu(X_k), g_k \rangle_F + \frac{h_k^2 L_1(f_\mu)}{2} \|g_k\|_F^2.$$

From non-expansiveness of projection, i.e. (21), and from the fact that $L_1(f_\mu) = (2L_0(f)/\mu)\sqrt{(n^2+n)/2}$, we have

$$f_\mu(X_{k+1}) \leq f_\mu(X_k) - h_k \langle \nabla f_\mu(X_k), g_k \rangle_F + \frac{h_k^2 L_0(f)}{\mu} \sqrt{\frac{n^2+n}{2}} \|\mathcal{O}_\mu(X_k, U_k)\|_F^2.$$

Define $\xi_k \triangleq \mathcal{O}_\mu(X_k, U_k) - \nabla f_\mu(X_k)$, then $\langle \nabla f_\mu(X_k), g_k \rangle_F = \langle \mathcal{O}_\mu(X_k, U_k), g_k \rangle_F - \langle \xi_k, g_k \rangle_F$. Based on [34, Lemma 1] we have $\langle \mathcal{O}_\mu(X_k, U_k), g_k \rangle \geq \|g_k\|_F^2 + \frac{1}{h_k} [\mathbf{h}(X_{k+1}) - \mathbf{h}(X_k)]$, where $\mathbf{h}(X) = 0$ if $X \in \mathbb{Q}^n$ and ∞ otherwise. Since $X_0 \in \mathbb{Q}^n$ and based on (2), we know that $\mathbf{h}(X_k) = 0$ for $k \geq 0$. Therefore,

$$f_\mu(X_{k+1}) \leq f_\mu(X_k) + h_k \langle \xi_k, g_k - P_{\mathbb{Q}^n}(X_k, \nabla f_\mu(X_k), h_k) \rangle_F + h_k \langle \xi_k, P_{\mathbb{Q}^n}(X_k, \nabla f_\mu(X_k), h_k) \rangle - h_k \|g_k\|_F^2 + \frac{h_k^2 L_0(f)}{\mu} \sqrt{\frac{n^2+n}{2}} \|\mathcal{O}_\mu(X_k, U_k)\|_F^2. \quad (24)$$

From [34, Proposition 1], we can write $\langle \xi_k, g_k - P_{\mathbb{Q}^n}(X_k, \nabla f_\mu(X_k), h_k) \rangle_F \leq \|\xi_k\|_F^2$. With this fact, we take expectation in (24) and get

$$h_k \mathbb{E}_{U_k} \{\|g_k\|_F^2\} \leq f_\mu(X_k) - \mathbb{E}_{U_k} \{f_\mu(X_{k+1})\} + h_k \sigma^2 + h_k^2 C(\mu), \quad (25)$$

where we also used $\mathbb{E}_{U_k} \{\xi_k\} = 0$, Theorem 3 to compute $\mathbb{E}_{U_k} \{\|\mathcal{O}_\mu(X_k, U_k)\|_F^2\}$, and Assumption 2 to upper bound $\mathbb{E}_{U_k} \{\|\xi_k\|_F^2\} \leq \sigma^2$. Recall that $S_N \triangleq \sum_{k=0}^N h_k$ and that $f_\mu(X^*) \geq f^*$. Then, taking expectation in $U_{0:k}$ in (25) and then taking summation, we obtain (8) as required. ■

ACKNOWLEDGMENT

The authors would like to thank the engineering staff at Toyota Higashi-Fuji Technical Centre, Susono, Japan, for their assistance in running the experiments included in this paper.

REFERENCES

- [1] C.-J. Lin, "Projected gradient methods for nonnegative matrix factorization," *Neural computation*, vol. 19, no. 10, pp. 2756–2779, 2007.
- [2] Z. Zhu, Q. Li, G. Tang, and M. B. Wakin, "Global optimality in low-rank matrix optimization," *IEEE Transactions on Signal Processing*, vol. 66, no. 13, pp. 3614–3628, 2018.
- [3] T. Zhao, Z. Wang, and H. Liu, "A nonconvex optimization framework for low rank matrix estimation," in *Advances in Neural Information Processing Systems*, 2015, pp. 559–567.
- [4] M. Anderson, X.-L. Li, P. Rodriguez, and T. Adali, "An effective decoupling method for matrix optimization and its application to the ICA problem," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2012, pp. 1885–1888.
- [5] P. Mertikopoulos, E. V. Belmega, R. Negrel, and L. Sanguinetti, "Distributed stochastic optimization via matrix exponential learning," *IEEE Transactions on Signal Processing*, vol. 65, no. 9, pp. 2277–2290, 2017.
- [6] S. Y. Shafi, M. Arcak, and L. El Ghaoui, "Graph weight allocation to meet Laplacian spectral constraints," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1872–1877, 2011.
- [7] A. Bertrand and M. Moonen, "Topology-aware distributed adaptation of Laplacian weights for in-network averaging," in *21st European Signal Processing Conference (EUSIPCO 2013)*. IEEE, 2013, pp. 1–5.
- [8] G. A. Bunin, F. F. Tirado, G. François, and D. Bonvin, "Run-to-run MPC tuning via gradient descent," in *Computer Aided Chemical Engineering*. Elsevier, 2012, vol. 30, pp. 927–931.
- [9] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe, "Automatic LQR tuning based on gaussian process global optimization," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 270–277.
- [10] M. Chaikovskii and I. Yadykin, "Optimal tuning of PID controllers for MIMO bilinear plants," *Automation and Remote Control*, vol. 70, no. 1, pp. 118–132, 2009.
- [11] J. L. Garriga and M. Soroush, "Model predictive control tuning methods: A review," *Industrial & Engineering Chemistry Research*, vol. 49, no. 8, pp. 3505–3515, 2010.
- [12] R. Shridhar and D. J. Cooper, "A tuning strategy for unconstrained multivariable model predictive control," *Industrial & engineering chemistry research*, vol. 37, no. 10, pp. 4003–4016, 1998.
- [13] G. Shah and S. Engell, "Tuning MPC for desired closed-loop performance for mimo systems," in *Proceedings of the 2011 American Control Conference*. IEEE, 2011, pp. 4404–4409.
- [14] P. Bagheri and A. Khaki-Sedigh, "An analytical tuning approach to multivariable model predictive controllers," *Journal of Process Control*, vol. 24, no. 12, pp. 41–54, 2014.
- [15] J. E. W. Santos, J. O. Trierweiler, and M. Farenzena, "Robust tuning for classical MPC through the multi-scenarios approach," *Industrial & Engineering Chemistry Research*, vol. 58, no. 8, pp. 3146–3158, 2019.
- [16] A. S. Yamashita, P. M. Alexandre, A. C. Zanin, and D. Odloak, "Reference trajectory tuning of model predictive control," *Control Engineering Practice*, vol. 50, pp. 1–11, 2016.
- [17] G. A. N. Júnior, M. A. Martins, and R. Kalid, "A PSO-based optimal tuning strategy for constrained multivariable predictive controllers with model uncertainty," *ISA transactions*, vol. 53, no. 2, pp. 560–567, 2014.
- [18] R. Suzuki, F. Kawai, H. Ito, C. Nakazawa, Y. Fukuyama, and E. Aiyoshi, "Automatic tuning of model predictive control using particle swarm optimization," in *2007 IEEE Swarm Intelligence Symposium*. IEEE, 2007, pp. 221–226.
- [19] J. Van der Lee, W. Svrcek, and B. Young, "A tuning algorithm for model predictive controllers based on genetic algorithms and fuzzy decision making," *ISA transactions*, vol. 47, no. 1, pp. 53–59, 2008.
- [20] A. S. Ira, C. Manzie, I. Shames, R. Chin, D. Nešić, H. Nakada, and T. Sano, "Tuning of multivariable model predictive controllers through expert bandit feedback," *International Journal of Control*, pp. 1–9, 2020.
- [21] G. S. Sankar, R. C. Shekhar, C. Manzie, T. Sano, and H. Nakada, "Fast calibration of a robust model predictive controller for diesel engine airpath," *IEEE Transactions on Control Systems Technology*, 2019.
- [22] Q. Lu, R. Kumar, and V. M. Zavala, "MPC controller tuning using bayesian optimization techniques," *arXiv preprint arXiv:2009.14175*, 2020.
- [23] F. Sorourifar, G. Makrygiros, A. Mesbah, and J. A. Paulson, "A data-driven automatic tuning method for mpc under uncertainty using constrained bayesian optimization," *IFAC-PapersOnLine*, vol. 54, no. 3, pp. 243–250, 2021.
- [24] M. Menner, P. Worsnop, and M. N. Zeilinger, "Constrained inverse optimal control with application to a human manipulation task," *IEEE Transactions on Control Systems Technology*, 2019.
- [25] A. Ramadan, J. Choi, C. J. Radcliffe, J. M. Popovich, and N. P. Reeves, "Inferring control intent during seated balance using inverse model predictive control," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 224–230, 2019.
- [26] P. J. Forrester, *Log-gases and random matrices (LMS-34)*. Princeton University Press, 2010.
- [27] G. W. Anderson, A. Guionnet, and O. Zeitouni, *An introduction to random matrices*. Cambridge university press, 2010, vol. 118.
- [28] C. Tracy and H. Widom, "On orthogonal and symplectic matrix ensembles," *Communications in Mathematical Physics*, vol. 177, no. 3, pp. 727–754, 1996.
- [29] Y. Nesterov and V. Spokoiny, "Random gradient-free minimization of convex functions," *Foundations of Computational Mathematics*, vol. 17, no. 2, pp. 527–566, 2017.
- [30] M. Athans and F. Schweppe, "Gradient matrices and matrix calculations," M.I.T. Lincoln Lab., Tech. Rep., 1965.

- [31] J. Dattorro, *Convex optimization & Euclidean distance geometry*. Meboo Publishing USA, 2011.
- [32] T. Hayakawa and Y. Kikuchi, "The moments of a function of traces of a matrix with a multivariate symmetric normal distribution," *South African Statistical Journal*, vol. 13, no. 1, pp. 71–82, 1979.
- [33] S. Liu, X. Li, P.-Y. Chen, J. Haupt, and L. Amini, "Zeroth-order stochastic projected gradient descent for nonconvex optimization," in *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 2018, pp. 1179–1183.
- [34] S. Ghadimi, G. Lan, and H. Zhang, "Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization," *Mathematical Programming*, vol. 155, no. 1-2, pp. 267–305, 2016.
- [35] S. Liu, B. Kailkhura, P.-Y. Chen, P. Ting, S. Chang, and L. Amini, "Zeroth-order stochastic variance reduction for nonconvex optimization," *Advances in Neural Information Processing Systems*, vol. 31, pp. 3727–3737, 2018.
- [36] J. Wahlström and L. Eriksson, "Modelling diesel engines with a variable-geometry turbocharger and exhaust gas recirculation by optimization of model parameters for capturing non-linear system dynamics," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 225, no. 7, pp. 960–986, 2011.
- [37] R. C. Shekhar, G. S. Sankar, C. Manzie, and H. Nakada, "Efficient calibration of real-time model-based controllers for diesel engines—Part I: Approach and drive cycle results," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 843–848.
- [38] J. B. Rawlings and D. Q. Mayne, *Model predictive control: Theory and design*. Nob Hill Pub., 2009.
- [39] L. Wang, *Model predictive control system design and implementation using MATLAB®*. Springer Science & Business Media, 2009.
- [40] G. S. Sankar, R. C. Shekhar, C. Manzie, T. Sano, and H. Nakada, "Fast calibration of a robust model predictive controller for diesel engine airpath," *IEEE Transactions on Control Systems Technology*, 2019.
- [41] N. J. Higham, "Computing a nearest symmetric positive semidefinite matrix," *Linear algebra and its applications*, vol. 103, pp. 103–118, 1988.
- [42] D. R. Jones, C. D. Perttunen, and B. E. Stuckman, "Lipschitzian optimization without the lipschitz constant," *Journal of optimization Theory and Applications*, vol. 79, no. 1, pp. 157–181, 1993.
- [43] G. A. Bunin and G. François, "Lipschitz constants in experimental optimization," *arXiv preprint arXiv:1603.07847*, 2016.
- [44] M. O. Ahmed, S. Vaswani, and M. Schmidt, "Combining bayesian optimization and lipschitz optimization," *Machine Learning*, vol. 109, no. 1, pp. 79–102, 2020.
- [45] D. Mayne, J. Rawling, C. Rao, and P. Sokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, pp. 789–814, 2000.



Alejandro I. Maass is a research fellow in the Department of Electrical and Electronic Engineering at the University of Melbourne, Australia. He received the "Ingeniero Civil Electrónico" degree and the M.Sc. degree in electronic engineering from the Universidad Técnica Federico Santa María (UTFSM), Chile, in 2013. In 2014, he was a research assistant and sessional lecturer in the Department of Electronic Engineering at UTFSM. He subsequently received the Ph.D. degree in engineering from the University

of Melbourne, Australia, in 2019. His main research interests include, but are not limited to, networked control systems, non-linear and hybrid systems, and numerical optimisation.



Chris Manzie is currently a full Professor and Head of Department of Electrical and Electronic Engineering at the University of Melbourne, and also the Director of the Melbourne Information, Decision and Autonomous Systems (MIDAS) Laboratory. Over the period 2003–2016, he was an academic in the Department of Mechanical Engineering, with responsibilities including Assistant Dean with the portfolio of Research Training (2011–2017), and Mechatronics Program Director (2009–2016). Professor Manzie was also

a Visiting Scholar with the University of California, San Diego in 2007 and a Visiteur Scientifique at IFP Energies Nouvelles, Rueil Malmaison in 2012. His research interests are in model-based and model-free control and optimization with applications in a range of areas, including systems related to autonomous systems, energy, transportation, and mechatronics.



Iman Shames is currently a Professor of Mechatronics at the School of Engineering at the Australian National University. Previously, he had been an Associate Professor at the Department of Electrical and Electronic Engineering, the University of Melbourne from 2014 to 2020 and a Senior Lecturer and a McKenzie fellow at the same department from 2012 to 2014, and before that he was an ACCESS Postdoctoral Researcher at the ACCESS Linnaeus Centre, the KTH Royal Institute of Technology, Stockholm, Sweden. He received his B.Sc. degree in Electrical Engineering from Shiraz University in 2006, and the Ph.D. degree in engineering and computer science from the Australian National University, Canberra, Australia in 2011. His current research interests include, but are not limited to, optimization theory and its application in control and estimation, mathematical systems theory, and security and privacy in cyber-physical systems.

Sweden. He received his B.Sc. degree in Electrical Engineering from Shiraz University in 2006, and the Ph.D. degree in engineering and computer science from the Australian National University, Canberra, Australia in 2011. His current research interests include, but are not limited to, optimization theory and its application in control and estimation, mathematical systems theory, and security and privacy in cyber-physical systems.



Hayato Nakada received B.Eng. degree in informatics and mathematical science and M.Inf. (Master of Informatics) and Ph.D. degrees in applied mathematics and physics from Kyoto University, Japan in 2000, 2002 and 2005, respectively, specializing in control theory. Since 2005, he has been with the Advanced Powertrain Management System Development Division of Toyota Motor Corporation, Japan. He is employed as a Project Manager, working on development of advanced controllers and models for a range

of powertrain projects.