Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Semi-Active Suspension Control Design via Bayesian Optimization

Gianluca Savaia *, **Simone Formentin** *, **Sergio M. Savaresi** *

* *Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Milano, MI 20133, Italy.*
Email: simone.formentin@polimi.it

Abstract: The fine tuning of semi-active suspension control systems for road vehicles is usually a costly and burdensome task, needing control expertise and many hours of professional driving. In this paper, we propose a data-driven tuning method enabling the automatic calibration of the parameters of the suspension controller using a small number of experiments and exploiting Bayesian Optimization tools. The effectiveness of the proposed approach is validated on a commercial multi-body simulator. As a side contribution, the approach is shown to be robust with respect to variations of the testing conditions.

Keywords: Semi-active suspension, Vehicle dynamics, Bayesian optimization, Calibration

1. INTRODUCTION

End-of-Line (EoL) calibration of semi-active suspension systems for road vehicles is known to be a costly and heavy task, usually needing a close collaboration among professional drivers, control system experts and test engineers, as well as many hours of driving. For this reason, in Milliken et al. [1995], the tuning of the damper was defined as a *black art*.

A popular option to make EoL calibration an easier and more efficient task is to rely upon *Hardware-in-the-Loop* (HiL) simulators, where the driver seats in the car cockpit and drives within a virtual environment, while the variations of the road profile are emulated by a robotic arm, see Schuette and Waeltermann [2005]. Another alternative to real-world tests is the so-called *4-poster test system*, where the car is placed on four moving platforms that emulate the road excitation, Chindamo et al. [2017]. Despite their major advantages in terms of costs and repeatability of the tests, such procedures are clearly limited by the lab environment and still require domain experts to perform a suitable calibration.

For the above reasons, in this paper we propose a data-driven calibration method, in which the problem is formulated as a *sequential decision making* task, see Barto et al. [1989], allowing the test engineer to perform fully automatic tuning of the suspension system. The input of the algorithm will only be the data measured by on-board sensors, while the control parameters will be returned as outputs with no need of human intervention, by minimizing a pre-defined performance-oriented objective function (*e.g.*, addressing comfort or road handling) via Bayesian Optimization, see Snoek et al. [2012], Klein et al. [2016]. We will argue that, using the latter tool, only a limited amount of driving time is required, thus making real experimental tests cost-competitive with respect to HiL and other indoor alternatives. This would not be true for other global optimization approaches, *e.g.*, Particle Swarm or Genetic Algorithms, since they require a high number of experiments, as discussed in Alkhatib et al. [2004]. For more details on Bayesian Optimization, see the surveys in Shahriari et al. [2015], Frazier [2018].

To summarize, the novelty of this work consists of a fully automatic calibration protocol of a fixed structure suspension controller for a road vehicle. The protocol defines the working conditions which the test must follow, the definition of an objective function corresponding to the major key performance index and a smart sampling strategy based on Bayesian Optimization, suggesting the calibration parameters which shall be tested to rapidly converge to the optimum.

The remainder of this paper is as follows. In Section 2, the calibration problem is formally stated for a parametric semi-active suspension controller and a specific layout is selected to illustrate the performance of the proposed algorithm. The automatic tuning algorithm based on Bayesian Optimization is illustrated in detail in Section 3. Simulation results, including a sensitivity analysis with respect to small variations of the testing conditions, are reported in Section 4. The paper is ended by some concluding remarks.

2. PROBLEM STATEMENT

Suspensions represent a key technology in road vehicles for both comfort and performance, see Savaresi et al. [2010]. A suspension device typically consists of a spring and a damper connected between the wheel and corner of the vehicle. In semi-active suspensions, the damping coefficient can be adjusted in real-time. This allows to combine the superior high-frequency filtering of a low-damping configuration and the good vehicle stability of a stiff damper.

A semi-active damper is a strictly passive system (*i.e.*, no energy is fed into the system) characterized by a controllability region defined in the I and III quadrants of the Force - Speed plane, as in Figure 1. This region defines

Copyright lies with the authors

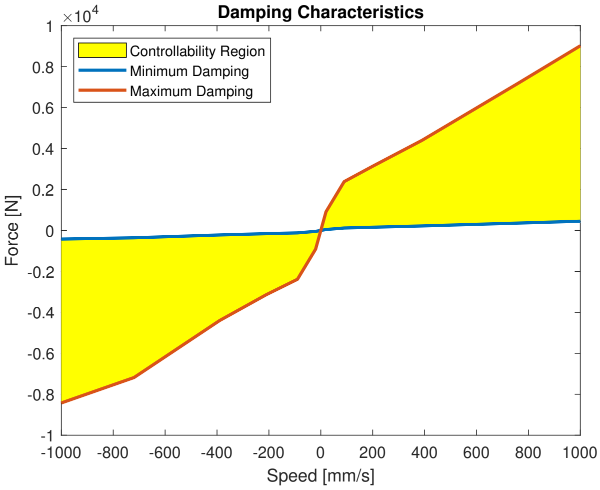

Fig. 1. Typical controllability region of a semi-active damper.

the boundaries of the minimum and maximum damping force F_d, which the device can exert as a function of the elongation rate $\Delta\dot{z}$. The damping force is therefore given by

$$F_d = C_{ref}\Phi_{max}(\Delta\dot{z}) + (1 - C_{ref})\Phi_{min}(\Delta\dot{z}), \quad (1)$$

where Φ is a non-decreasing function which shapes the damping characteristics (intrinsic to the physical hardware) and C_{ref} is the control signal which denotes the damping coefficient defined in the range $[0, 1]$ (interpolating between minimum and maximum). Ideally, Φ should be linear in $\Delta\dot{z}$, however for an actual physical device the damping characteristics looks more similar to the function shown in Figure 1.

The objective of the controller consists in choosing the damping coefficient achieving the optimal comfort at any given time. The most popular control algorithm in the semi-active suspension control literature is the so-called *skyhook*, see Karnopp [1995], which is here presented in its variant described by:

$$C_{ref} = C + K\dot{z}_c\Delta\dot{z}, \quad (2)$$

where $\dot{z}_c$ represents the vertical speed of the chassis.

The performance of the control algorithm thus depends on the control parameters (C, K), where the former specifies a nominal damping ratio, whilst the latter the sensitivity to the road excitation. Notice that if $K = 0$, the damper behaves as a passive suspension. As the control law is formulated in (2), the parameters (C, K) are adimensional.

The aim of this paper is to present a calibration procedure which allows to find the optimal values of the pair (C, K) from on-board measurements only, with limited experimental effort. In general, the proposed methodology could be applied to any parametric control law other than (2).

2.1 The Considered Key Performance Index

The main purpose of a suspension system is filtering vibrations due to the road excitation. The quality of the ride depends on the bandwidth where these vibrations occur, and there exists a major trade-off in the calibration of a suspension system with respect to its filtering capabilities.

In Figure 2, the spectrum of the vertical accelerations, relative to an experiment on a test road, is displayed for three configurations: passive low-damping, passive high-damping and semi-active damping.

The low-damping configuration achieves the best high-frequency filtering effect, but is poorly stable around the body resonance (around $2\,Hz$). On the contrary, the high-damping shows a good damping effect, but to the detriment of high-frequency vibrations. The skyhook logic, when accurately calibrated, can achieve the best compromise between the two extreme passive configurations.

The manner in which vibrations affect comfort is dependent on the vibration frequency spectrum. In ISO 2631-1:1997 [E], an evaluation of human exposure to whole-body vibrations is presented. The standard defines a complex key performance index consisting of tri-axial (longitudinal, lateral and vertical) and rotational (roll, pitch and yaw) accelerations; each component is weighted to take into account the range of frequencies which can be perceived by the human body. Eventually, the aforementioned terms are combined in a weighted summation whose result is the final *ISO Index*, defined as:

$$J = \sqrt{\sum_i (k_i J_i)^2}, \quad (3)$$

where

$$J_i = \sqrt{\frac{1}{T}\int_0^T A_i^w(t)^2\,dt}$$
$$A_i^w(t) = W_i(j\omega)A_i(t).$$

In the above expressions, A_i is the i^{th} acceleration (longitudinal, lateral, vertical and rotational), W_i is the frequency weight associated to that term, J_i the root mean squared of A_i^w and J the ultimate condensed ISO Index. The reader is referred to ISO 2631-1:1997 [E] for the exact expressions of the frequency weights W_i and gains k_i.

The quality of a given calibration of the suspension controller will be evaluated from now on after an experiment over a test road as specified by (3).

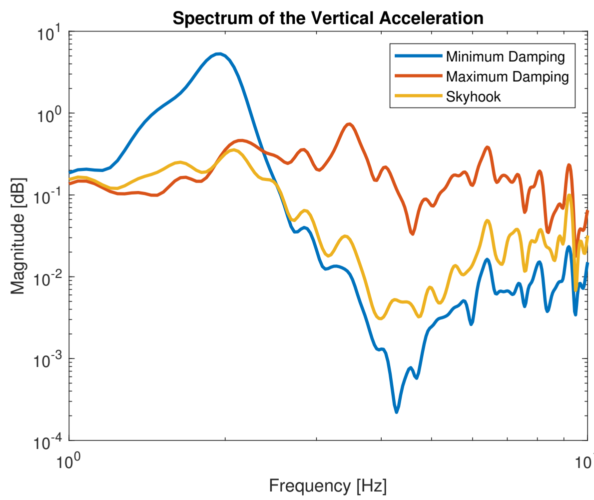

Fig. 2. Spectrum of the vertical acceleration sensed on the vehicle body in three different configurations.

3. THE PROPOSED DATA-DRIVEN APPROACH

This section introduces the proposed Bayesian Automatic Calibration (BAC) protocol for a parametric suspension control system like the skyhook controller of (2).

3.1 Experimental Protocol

The experiments shall be performed on a single test road representative of the dynamics of interest. This is a critical point, since the result of the optimization strongly depends on the road excitation; in other words, the optimal configuration for a driving cycle on a highway and off-road are expected to differ significantly. Hence, the choice of the road profile has an important impact on the outcome of the optimization and shall be carefully evaluated in the design phase.

In particular, each experiment must be performed over the same road segment at constant velocity in order to ensure the same frequency excitation.

3.2 Performance Evaluation

The evaluation of each experiment is assessed at the end of the experiment itself, based on the ISO Index sensed on the driver seat via an inertial measurement unit (IMU). It is important to remark that any other index based on sensors data could be employed to target a different objective.

The objective of the calibration consists in finding the optimal configuration for the control parameters, which minimizes the ISO Index (thus maximizing the riding comfort) in the least possible number of iterations.

3.3 Methodology

An effective technique to solve these problems would be to formulate an optimization problem including the model of the vehicle and the information about the road profile. However, models are always only an idealization of the real system. For this reason, we will follow a different, purely *data-driven* rationale, consisting of two main ingredients: a *surrogate function* to estimate the objective function from data, and an *acquisition function* to sample the next observation, see Bemporad [2019].

The *surrogate function* shall be ideally model-free or non-parametric, unless one desires to constrain the objective function to a certain class (*e.g.*, quadratic). In Bayesian Optimization (BO), the objective function is assumed to be drawn from a Gaussian Process (GP), where each observation is a random variable normally distributed and jointly Gaussian with one another. In particular, each pair of observations is bonded by a covariance function encoding the belief that closer points shall have higher correlation than those far apart; typical choices are the Gaussian and Màtern kernels.

Hence, given a set of observations, it is possible to fit a GP whose posterior mean is the expected estimation of the objective function, and its variance represents a confidence interval for the estimation. This means that, for any given set of control parameters, it is possible to compute the estimated ISO Index and a goodness indicator of this estimate.

Given this posterior model, there are several options for the *acquisition function* which aims at sampling the next observation where there is a higher probability to find a minimum, according to the posterior mean and variance. The most popular algorithms are *Expected Improvement* and *Upper Confidence Bound* which can be found declined in different variations to tackle the so-called exploration-exploitation trade-off.

3.4 Algorithm

A pseudocode summarizing this procedure is shown in Algorithm 1. The initialization is performed by picking at random (or any *best guess* policy) a set of control parameters (C, K), which are evaluated by performing an experiment and then computing the ISO Index defined in (3). The objective function is then estimated by fitting a Gaussian Process (GP). Eventually, the next set of control parameters are chosen by maximizing the expected improvement (EI), which identifies the region in the parameters space, where it is most likely to find the minimum according to such an acquisition function; the parameters are then evaluated by performing another experiment and the ISO Index is computed and updated accordingly.

This class of optimization problems does not rely upon a stopping condition since they are based on the assumption that the evaluation of an experiment is expensive and, thus, there shall be a budget cap to the maximum number of experiments to be performed. Hence, the procedure in Algorithm 1 is repeated until the allowed maximum number of iterations is reached.

Algorithm 1 Bayesian Automatic Calibration (BAC)

$(C^{obs}, K^{obs}) \leftarrow random$
Perform Experiment and Evaluate J^{obs}

while $i < MaxIterations$ **do**
  $\mu, \sigma \leftarrow fitGP(C^{obs}, K^{obs}, J^{obs})$

  $(C^{next}, K^{next}) \leftarrow argmax\ EI(\mu, \sigma)$
  Perform Experiment and Evaluate J^{next}

  Update $(C^{obs}, K^{obs}, J^{obs})$

$(C^{opt}, K^{opt}) \leftarrow argmin(\mu)$

4. SIMULATION RESULTS

In this section, different data-driven calibration strategies are compared: i) Grid Search, ii) Random Search and iii) Bayesian Automatic Calibration (BAC). These strategies are the best candidates to solve the global optimization problem in the scenario described here, where the evaluation of an experiment is an expensive task.

Grid Search represents the current calibration paradigm where the engineer tests each possible combination of the control parameters; it is the least efficient strategy which requires many experiments if the parameters space is fitted on a thin grid. Random Search is a standard benchmark

Fig. 3. The luxury sedan vehicle employed in the simulations of Section 4.

in literature, where the parameters are sampled at random from a uniform distribution in the parameters space; it is more efficient than Grid Search, but the optimum is not guaranteed to be found. BAC is the candidate framework proposed in this paper introduced in Section 3.

4.1 Simulation Setup

The experimental simulation of the proposed framework is performed on a full-fledged commercial vehicle simulator which can model complex multi-body dynamics.

The target vehicle is the luxury sedan illustrated in Figure 3, where the main parameters of interest are:

- sprung mass: $1370\,Kg$
- unsprung mass: $160\,Kg$
- spring stiffness: $153\,N/mm$

The vehicle is modeled as a MIMO system whose inputs are the damping forces at the four corners and the outputs are the accelerations the vehicle is subject to (used for the ISO Index evaluation), the elongation rates $\Delta\dot{z}$ and the vertical corner velocities $\dot{z}_c$ (used for feedback in control).

The control system is implemented in a Matlab/Simulink environment following Equations (1)-(2).

The candidate road profile is representative of a rough pavement with low-frequency valley, exciting a broad range of frequencies, traveled at the constant speed of $80\,km/h$.

4.2 Main Results

The utilization of a simulation environment made it possible to find the actual true minimum of the objective function via Gradient Descent. This is useful for benchmarking the different global optimization strategies.

The Grid Search strategy consisted in testing C in $[0, 1]$ with step 0.05, and K in $[0, 50]$ with step 1; therefore, the grid accounts for 1071 cells, corresponding to the equivalent number of experiments to be performed in order to find the optimal value. A projection of the ISO Index evaluated at each point of the grid is shown in Figure 4, where the green cross represents the true optimum computed via Gradient Descent.

In Random Search, C and K are drawn from uniform distributions, in $[0, 1]$ and $[0, 50]$, respectively. The maximum number of iterations has been set to 30; after performing 30 experiments, the one whose ISO Index is smaller is chosen as the best calibration.

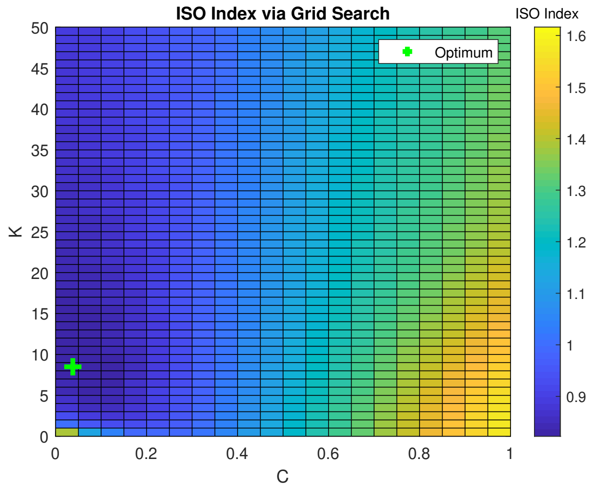

Fig. 4. ISO Index computed in a Grid Search fashion. The green cross denotes the true optimum (found off-line via gradient descent).

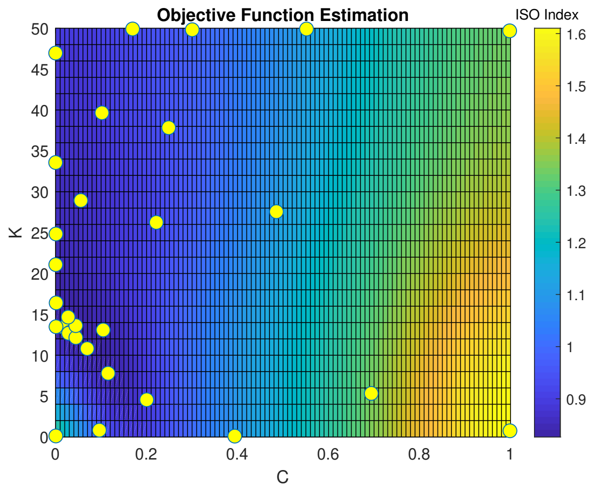

Fig. 5. ISO Index estimated as a Gaussian Process. The yellow dots represent the points of the parameters space explored via closed-loop experiments.

The proposed BAC strategy can identify the region where the optimum is located after only 30 experiments. The chronicle of the observations follows Algorithm 1, and it can be interpreted as follows: at first, pairs at the boundaries of the parameters space are evaluated to *learn* the approximate shape of the objective function; then, the observations are sampled in the neighborhood of the region where the optimum is most likely to be located. Figure 5 shows the ISO Index estimated after 30 iterations; the observations are represented by yellow dots and are arranged close to the estimated minimum.

The algorithm assigns an indicator for the goodness of the estimation, indicated by σ in Algorithm 1, shown in Figure 6: the points where the objective function has been observed are assigned a low uncertainty (ideally this should be zero for a deterministic process), whereas areas far apart from an observation are more uncertain. The main feature of BAC consists in an efficient sampling strategy, which does not consider the areas where the minimum is

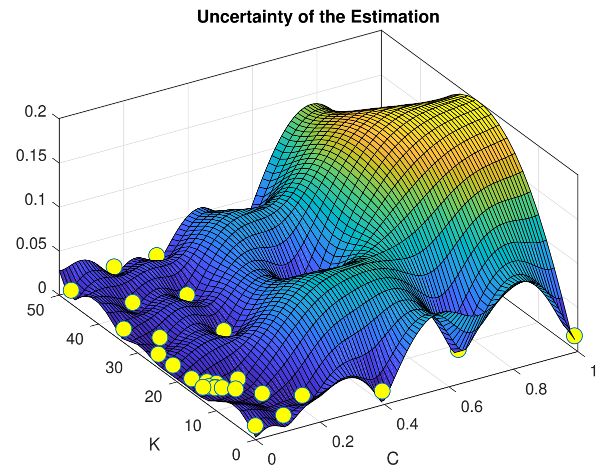

Fig. 6. Uncertainty of the ISO Index estimation. The yellow dots represent the points of the parameters space explored via closed-loop experiments.

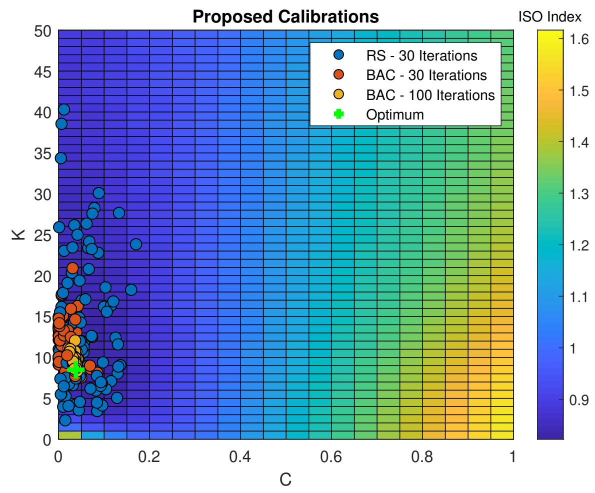

Fig. 7. Optimal calibration parameters given by the different methods highlighted in the legend. For each strategy, 100 runs are proposed (each point in the figure corresponds to the optimum of a single run).

not likely to be found, even though the estimation of the actual objective function is poor.

For the comparison to be fair, since Random Search and BAC are stochastic methodologies, 100 Monte Carlo simulations have been considered. The best candidate calibration (for each simulation) is reported in Figure 7. BAC is present in two configurations: one where each simulation consists of 30 iterations, and another consisting of 100. Random Search has the highest variance, whereas BAC achieves a more stable solution whose deviation from the optimum is narrower as the number of iterations grows.

The results are summarized in Table 1 where, among the 100 runs, the median and worst-case calibration are considered, this latter to avoid a *lucky strike* due to randomness. The reason why the results for BAC after 30 iterations ($K = 20.91$) is far from the optimal value ($K = 8.49$) is that it lies in a region where the gradient is *quasi* zero and, therefore, more iterations would be needed to converge towards the true optimum.

If the number of iterations was translated to actual time, assuming 120 s for each experiment and 8 working-hours per day, *the Grid Search approach would require 4 full days of work, whereas BAC would cut that down to 3 hours.*

4.3 Sensitivity to Road Perturbations

The results presented in the previous section were obtained by simulating a perfectly repeatable experiment, where two iterations with the same parameters lead to the same value for the cost function J.

In reality, this is seldom the case since it is impossible to perform such an experiment with very high accuracy, and a slightly different trajectory or velocity may bias the objective function. It is therefore of utmost importance for the proposed methodology to be robust with respect to (small) perturbations of the experiment.

Therefore, in the remainder of this section, the Bayesian framework will be tested in a scenario where the nominal road profile $\bar{r}(t)$ is perturbed as in

$$r(t) = \bar{r}(t) + \xi(t),$$

where ξ is a frequency weighted white noise

$$\xi(t) = W(j\omega)\hat{\xi}(t), \qquad \hat{\xi}(t) \sim \mathcal{N}(0, \sigma).$$

We will not consider here the case where the frequency composition of the road profile is changed, as it would be equivalent to traveling on a totally different road; further details on this matter can be found in ISO 8608:2016 [E] and Agostinacchio et al. [2014]. Figure 8 shows a detail of three realizations of the perturbed road profile which shall resemble three different experiments over the same road segment in a real world scenario; the perturbation on each iteration is independent from the previous.

The performance of the proposed algorithm in this perturbed setup is shown in Figure 9. Although the estimation of the actual objective function is worse than the previous (nominal) case in Figure 5, the sampling strategy is still able to find control parameters within a neighborhood of the true optimum.

In general, a higher number of experiments are needed to reach the optimum, since the GP must estimate the variance of an observation which is no longer deterministic. The other methods discussed in Section 4, Grid Search and Random Search, are not robust against these perturbation *per se*, although it may be possible to estimate a confidence interval to the detriment of a dramatic increase in the number of evaluations.

5. CONCLUSIONS

This paper presented a novel approach for the automatic calibration of a fixed structure parametric suspension control system. It is shown that the proposed BAC protocol, based on Bayesian Optimization, requires (one to two orders of magnitude) less iterations than Grid Search to reach the region of convergence where the optimum lies.

Future research will be devoted to an experimental validation of the proposed approach.

	True	GS	Median			Worst Case		
			RS	BAC		RS	BAC	
Iterations	-	1071	30	30	100	30	30	100
C	0.03	0.05	0.01	0.07	0.03	0.01	0.03	0.03
K	8.49	8.00	17.89	8.99	9.74	2.31	20.91	12.07
ISO Index	0.82	0.82	0.84	0.83	0.82	0.89	0.84	0.83

Table 1. Optimal suspension control parameters found by each strategy, after the specified number of iterations.

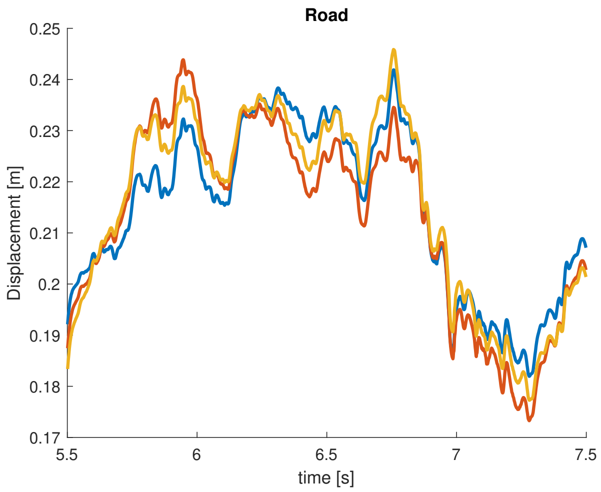

Fig. 8. Three random realizations of the perturbed road profile.

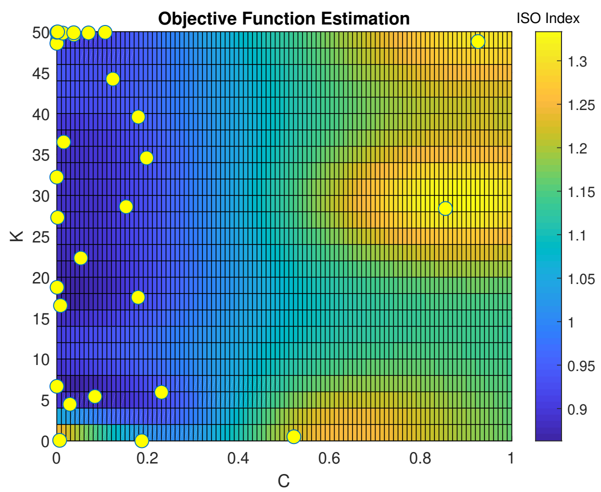

Fig. 9. ISO Index estimated in the perturbed scenario. The yellow dots represent the points of the parameters space explored via closed-loop experiments.

REFERENCES

Agostinacchio, M., Ciampa, D., and Olita, S. (2014). The vibrations induced by surface irregularities in road pavements–a matlab® approach. *European Transport Research Review*, 6(3), 267.

Alkhatib, R., Jazar, G.N., and Golnaraghi, M. (2004). Optimal design of passive linear suspension using genetic algorithm. *Journal of Sound and vibration*, 275(3-5), 665–691.

Barto, A.G., Sutton, R.S., and Watkins, C.J. (1989). Learning and sequential decision making. In *Learning and computational neuroscience*. Citeseer.

Bemporad, A. (2019). Global optimization via inverse distance weighting. *arXiv preprint arXiv:1906.06498*.

Chindamo, D., Gadola, M., and Marchesin, F.P. (2017). Reproduction of real-world road profiles on a four-poster rig for indoor vehicle chassis and suspension durability testing. *Advances in Mechanical Engineering*, 9(8), 1687814017726004.

Frazier, P.I. (2018). A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.

ISO 2631-1:1997(E) (1997). Mechanical vibration and shock – Evaluation of human exposure to whole-body vibration. Standard, International Organization for Standardization, Geneva, CH.

ISO 8608:2016(E) (2016). Mechanical vibration – Road surface profiles – Reporting of measured data. Standard, International Organization for Standardization, Geneva, CH.

Karnopp, D. (1995). Active and semi-active vibration isolation. In *Current Advances in Mechanical Design and Production VI*, 409–423. Elsevier.

Klein, A., Falkner, S., Bartels, S., Hennig, P., and Hutter, F. (2016). Fast bayesian optimization of machine learning hyperparameters on large datasets. *arXiv preprint arXiv:1605.07079*.

Milliken, W.F., Milliken, D.L., et al. (1995). *Race car vehicle dynamics*, volume 400. Society of Automotive Engineers Warrendale.

Savaresi, S.M., Poussot-Vassal, C., Spelta, C., Sename, O., and Dugard, L. (2010). *Semi-active suspension control design for vehicles*. Elsevier.

Schuette, H. and Waeltermann, P. (2005). Hardware-in-the-loop testing of vehicle dynamics controllers–a technical survey. *SAE transactions*, 593–609.

Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., and De Freitas, N. (2015). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1), 148–175.

Snoek, J., Larochelle, H., and Adams, R.P. (2012). Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, 2951–2959.