



HAL
open science

Robust Consumption Planning from Uncertain Power Demand Predictions

Mathieu Randon, Benjamin Quost, Dirk von Wissel, Nassim Boudaoud

► **To cite this version:**

Mathieu Randon, Benjamin Quost, Dirk von Wissel, Nassim Boudaoud. Robust Consumption Planning from Uncertain Power Demand Predictions. 35th IEEE Intelligent Vehicles Symposium (IV 2023), Jun 2023, Anchorage, United States. pp.1-7, 10.1109/IV55152.2023.10186569 . hal-04371329

HAL Id: hal-04371329

<https://hal.science/hal-04371329>

Submitted on 3 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust Consumption Planning from Uncertain Power Demand Predictions

Mathieu Randon
Heudiasyc lab., UTC
Renault SAS
 mathieu.randon@renault.com

Benjamin Quost
Heudiasyc, UMR CNRS 7253
UTC, Compiègne, France
 benjamin.quost@utc.fr

Dirk von Wissel
Renault SAS
DEA MWA powertrain control
 dirk.von-wissel@renault.com

Nassim Boudaoud
Roberval, FRE CNRS 2012
UTC, Compiègne, France
 nassim.boudaoud@utc.fr

Abstract—A plug-in hybrid electric vehicle (PHEV) satisfies the driver’s power demand with two types of energy potentials: fuel and electrical energy provided by a battery. Classically, the battery consumption is planned over a trip to minimize the expected fuel consumption. A cautious driver will save battery potential to cross restricted areas (with desired low or even zero fuel consumption) without the fuel engine. This paper proposes an approach to minimize energy consumption while controlling the risk of a PHEV falling short of battery potential when crossing a restricted area. We use a nonlinear Gaussian process, trained on real vehicle data, for predicting the vehicle consumption. We take into account prediction uncertainty by ensuring that the driver’s highest power demand will be satisfied with a high probability. The interest of the approach is demonstrated by a simulated trip around Paris.

I. INTRODUCTION AND PROBLEM DESCRIPTION

A. Context

Since the European Union measures on reducing air pollution [1], 247 European cities have created low emission zones [2] so as to restrict access to vehicles meeting specific emission criteria [3]. Confirming positive results [4], [2], most cities are now planning a gradual increase of restrictions, resulting by 2030 in so-called “zero-emission vehicle zones” (ZEV zone) [5]. In order to boost the deployment of ZEV zones, the European Commission proposed a new standard [6] motivating plug-in hybrid electric vehicle (PHEV) to be equipped with *geofencing technologies* so as to avoid the internal combustion engine (ICE) being used in such areas. Other ZEV zones projects are planned outside of the EU, such as in cities of Foshan and Luoyang (China), or Kevadia (India) [5].

A PHEV, satisfies the driver’s power demand with two types of energy potentials: fuel (stored in a tank) and electrical energy (provided by a battery). Therefore, PHEVs will only be allowed in certain cities, provided that the ICE remains switched off within the ZEV zone. Due to PHEVs having a reduced range in full-electric mode, it is crucial to manage the battery potential (ΔE_b) so as to guarantee that the vehicle will be able to reach a charging station or to drive out of the ZEV zone.

B. Purpose of this work

Trip description: When a driver enters a destination into the vehicle’s navigation system, the future trip is planned. The powertrain controller computes a vector of descriptive variables \mathbf{x}_i (or “driving context”) for each future road section¹. These latter are computed in advance, as portions of the road with fixed length—this length can be made larger provided it does not significantly modify the characteristics of the section, which are average speed, average road slope, and expected auxiliary consumption².

1) *Consumption plan:* On any future road section, fuel and battery energy consumptions obviously compensate each other to satisfy the driver’s power demands. For a PHEV with limited battery energy potential, the optimal planning of battery energy consumption over the whole trip significantly reduces the global fuel energy consumption. When crossing a ZEV zone, the fuel energy consumption must be set to 0: the PHEV thus relies on the battery only.

To manage such additional consumption constraints, this paper builds a predictive energy management strategy (PEMS) [7], [8], that aims at optimally planning both fuel and battery energy consumptions ΔE_{fi} and ΔE_{bi} on a driving context \mathbf{x}_i . Let $\Delta E_i = [\Delta E_{fi}, \Delta E_{bi}]^T$ be the joint energy consumption vector. The corresponding plan of remaining energy potential are E_i , E_{fi} and E_{bi} (for instance, $E_{bi} = E_{bini} + \sum_{k=0}^i \Delta E_{bk} \cdot dl_k$), where dl_k is the length of road section k . The PEMS solves the following problem:

$$\max_{\Delta E} E_{fn}, \quad \text{such that for all } i = 1, \dots, n, \quad (1a)$$

$$\Delta E_i = W_i, \quad (1b)$$

$$E_i = E_{wi}, \quad (1c)$$

$$\Delta E_i \in [\underline{\Delta E}_i, \overline{\Delta E}_i], \quad (1d)$$

$$E_i \in [E_i, \bar{E}_i]. \quad (1e)$$

It computes a consumption plan ΔE for both energy sources so as to maximize the remaining fuel at arrival E_{fn} (1a)

¹Indices refer to road sections, e.g. \mathbf{x}_i is the descriptive vector of road section $i \in \{0, \dots, n\}$; as a convention, we refer to a set of values by dropping the index, e.g. the horizon of events is $\mathbf{x} = \{\mathbf{x}_0, \dots, \mathbf{x}_n\}$.

²The data used in the example, and the road pre-processing (including the road section characteristics) are actually provided by Renault S.A.S., and cannot be disclosed due to a privacy agreement.

while meeting the driver’s local (1b) and global (1c) power demands, respectively $\mathbf{w}_i = [\mathbf{w}_{fi}, \mathbf{w}_{bi}]^\top$ and \mathbf{E}_{w_i} such that

$$\mathbf{E}_{w_i} = \mathbf{E}_{ini} + \sum_{k=0}^i \mathbf{w}_k \cdot d\mathbf{l}_k. \quad (2)$$

Here, $\mathbf{E}_{ini} = [\mathbf{E}_{fini}, \mathbf{E}_{bini}]^\top$; and the road section lengths $d\mathbf{l}_k$ converts energy consumptions (in kWh/km and L/100km) into energy potentials (in kWh and L).

Additional constraints may restrict both energy consumptions ΔE_i within powertrain limits $[\underline{\Delta E}_i, \overline{\Delta E}_i]$ for some road sections (1d) (typically, in a ZEV zone). Similarly, both energy potentials E_i are restricted within energy storage limits $[\underline{E}_i, \overline{E}_i]$ (1e) (the battery must have at least 10% to avoid damaging cells).

2) *Robust estimation*: The PEMS generally assumes that future powertrain local and global power demands \mathbf{w} and \mathbf{E}_w are perfectly known, which is seldom the case. Indeed, various kinds of uncertainty may arise. *Aleatory uncertainty* is due to random events (e.g., traffic jams) corrupting initial estimates. *Epistemic uncertainty* results from the dynamic model at hand being ill-known due to scarce observations, which reduces the power of statistical inference (e.g., the actual behavior of the vehicle at a very high speed is much less frequently observed). *Model uncertainty*, due to approximations (e.g., neglecting battery aging), results in higher discrepancies between actual plan of both energy consumptions and both power demands.

To account for these uncertainties, we model local power demands \mathbf{w}_f and \mathbf{w}_b using two independent *Gaussian Process Regression models* (GPR) [9]. A GPR results in a Gaussian pdf on an estimated power demand, the uncertainty being quantified by its variance. Here, one GPR estimates the fuel local power demand \mathbf{w}_{fi} given the driving context \mathbf{x}_i and the planned battery energy consumption ΔE_{bi} ; and the other estimates the battery local power demand \mathbf{w}_{bi} from the planned fuel energy consumption ΔE_{fi} and \mathbf{x}_i (for more details on these models, we refer the reader to our previous work [10]). GPR estimates (e.g., $\widehat{\mathbf{w}}_{fi}$) are composed of an expected value (\mathbf{w}_{fi}^*) and a noise term ($\widetilde{\mathbf{w}}_{fi}$) accounting for the uncertainties:

$$\widehat{\mathbf{w}}_{fi}(\Delta E_{bi}, \mathbf{x}_i) = \mathbf{w}_{fi}^*(\Delta E_{bi}, \mathbf{x}_i) + \widetilde{\mathbf{w}}_{fi}(\Delta E_{bi}, \mathbf{x}_i), \quad (3a)$$

$$\widehat{\mathbf{w}}_{bi}(\Delta E_{fi}, \mathbf{x}_i) = \mathbf{w}_{bi}^*(\Delta E_{fi}, \mathbf{x}_i) + \widetilde{\mathbf{w}}_{bi}(\Delta E_{fi}, \mathbf{x}_i). \quad (3b)$$

Global power demand estimates $\widehat{\mathbf{E}}_w$ (and thus $\widehat{\mathbf{E}}_{w_f}$ and $\widehat{\mathbf{E}}_{w_b}$) can be computed using Eq. (2). These estimates are Gaussian, being a linear combination of Gaussian demands. We will refer to their expectation and noise as \mathbf{E}_w^* and $\widetilde{\mathbf{E}}_w$. Power demand estimates being uncertain, the PEMS should be *cautious* so as to avoid violating the energy consumption constraints (e.g., by falling short of battery storage in a ZEV zone, where the ICE shouldn’t be used). Obviously, sparing battery energy potentials for ZEV zones exclusively would

make a PHEV of little interest. Therefore, we propose to minimize the vehicle’s fuel energy consumption, while guaranteeing—at least with a high probability—that both fuel and battery power demand will be satisfied.

C. Outline

This paper addresses the PEMS optimization problem defined by (1), where local and global power demands \mathbf{w} and \mathbf{E}_w are replaced with their probabilistic estimates $\widehat{\mathbf{w}}$ and $\widehat{\mathbf{E}}_w$, obtained from the fuel and battery energy consumption plans ΔE . Energy consumptions constraints such as (1d) allow to adapt the PEMS to take ZEV zones into account. After reviewing related works in Section II, Section III presents a new strategy which guarantees with a controlled risk that the vehicle’s full electric range will be sufficient to cross the ZEV zone identified when planning the trip. Section IV discusses the results obtained on two simulated scenarios, where the models used are nevertheless learnt from real data measured with a PHEV. Finally, Section V concludes the paper and draws some perspectives of work.

II. RELATED WORK

To our knowledge, few works consider adapting the PEMS so as to guarantee a sufficient electric range to cross a ZEV zone [11], [12]. None of them take into account the fact that the estimated power demands are uncertain.

Optimization under uncertainty can take several forms [13]. Our problem, formulated within a probabilistic framework, is a *stochastic program* [14], [15], and more specifically a *chance-constrained program*. Chance-constrained programming has been introduced by [16] to determine a production plan robust to adversarial demands. It makes use of cautious statistical decision functions [17] which transform the optimization of uncertain variables into the optimization of an accurate worst-case regret [18]. The resulting conservative decision is sub-optimal, with the additional cost being called the price for robustness [19]. When uncertainty is described by a multivariate normal distribution, this problem was shown to be convex [20]. Closer to our problem, [21] reviewed recent advances in safe decision-making for robotics, and compared frameworks used in reinforcement learning and control theory. Our work connects to optimal trajectory planning methods which guarantee a so-called safety level II for state and command probabilistic constraints.

Dynamic programming (DP) is the reference approach for optimizing offline the PEMS of a vehicle [22]. Derived from Bellman’s principle of optimality [23], DP solves the optimization in a recursive manner. Starting from the final state over a discretized horizon of time, it selects the optimal action at each previous possible state. For online optimization, the A^* algorithm reduces the space of actions with an heuristic [24]. When actions cause random changes

of the system's state, DP becomes a Markov Decision Process (MDP) [25], [26]. Constrained MDPs [27] and robust MDPs [28] reduce the action domain so that the optimized set of actions, or policy, complies with probabilistic safety constraints. However, [29], [30] showed that expressing such a chance-constrained domain is NP-hard. Moreover, when cumulative resource constraints are considered as for PEMS, [31] showed that the optimization domain must be reduced in several additional iterations. As a consequence, reinforcement learning (RL) has been developed to replace the optimization by learning a policy which predicts near-optimal actions depending on the states of the system from the cost of observed trajectories. However, RL does not necessarily guarantee that constraints are satisfied during the execution phase [32]. Despite the recent advances in the cautious exploration of the state-actions domain [33], [34] with constraints approximated around an initial policy [35], or with probabilistic stability guarantees [36], safe RL remains an open research problem, and control theory generally achieves better formal safety guarantees.

In control theory, the PEMS is a model-predictive control (MPC) strategy [37] which plans an optimal command, with future events being predicted by a dynamical model. However, MPC is sensitive to uncertain model predictions. In [38], MPC strategies providing formal probabilistic safety guarantees are reviewed. The first approach, referred to as robust MPC, defines a probabilistic reachable set [39] of the state-action space before performing the optimization. The reachable set can be computed offline as in [40], or online using iterated GPR predictions as in [41], [42]. It can also be approximated via moment matching [43], with a smoothness Lipschitz constant [44], or with deep quantile regression [45]. The second approach, safety filtering, corrects the outputs of any optimal controller, such as MPC or RL, to ensure that safety constraints are satisfied. For this purpose, control barrier functions [46], [47] or model predictive safety filters [48] can be used.

Our contribution is inspired by the iterative Gaussian Process-based MPC [49], [50], which iteratively optimizes a command, the uncertainty being estimated by a GPR given the previous command; the robust MPC [41], which approximates inequality constraints with logarithmic barrier terms; and the collision avoidance algorithm [51], which uses a conditional value at risk metric. Section III will introduce a new robust MPC algorithm for the non-linear PEMS problem.

III. METHOD : A NEW ROBUST PEMS ALGORITHM

A. Cautious optimization

Solving the optimization problem defined above requires to turn the uncertain interval constraints (1b) and (1c) into deterministic constraints: the optimisation uses a single power demand value (consistent with the estimated distribution) for each future road section. Classically, the

expectation of the distribution is chosen. However, a robust approach rather considers “pessimistic” fuel and battery energy demands, by picking a quantile (either left-tail or right-tail) of the estimated distribution: such worst-case replacements aim at ensuring that the probability of the constraint being violated is small. Choosing this “adversarial quantile” may be difficult, to such an extent that some authors [30] proposed to make it a sub-problem of the robust optimization process. To alleviate the complexity, and inspired by [51], which uses a conditional value at risk metric, we proceed as follows. First, the fuel consumption plan is constrained to meet the expected fuel demand, estimated based on the battery energy consumption plan. Conversely, the battery consumption plan is constrained to meet the expected battery power demand, estimated from the fuel consumption plan. Safety margins are applied to both upper and lower bounds of the interval constraints for both energy consumptions and both energy potentials. These margins are computed so that local and global power demands, which might deviate from the expected ones, remain within a safety interval over energy consumptions and energy potentials with a desired risk.

1) *Demand satisfaction constraints*: The energy consumptions ΔE_i and energy potentials E_i should be chosen so as to satisfy the driver's power demands, i.e. they should be “close” to the expected estimation of local and global power demands $w_i^*(\Delta E_i, \mathbf{x}_i)$ and $E_{w_i}^*(\Delta E, \mathbf{x})$. For each road section $i = 0, \dots, n$, we thus replace Eq. (1b)-(1c) with

$$\Delta E_i - w_i^*(\Delta E_i, \mathbf{x}_i) \in [-\varepsilon_{\Delta E, i}, \varepsilon_{\Delta E, i}], \quad (4a)$$

$$E_i(\Delta E) - E_{w_i}^*(\Delta E, \mathbf{x}) \in [-\varepsilon_{E, i}, \varepsilon_{E, i}]; \quad (4b)$$

the positive thresholds $\varepsilon_{\Delta E, i}$ and $\varepsilon_{E, i}$ specify how firmly the energy consumption and energy potential plans are driven to satisfy the expected local and global power demands.

2) *Cautious consumption restriction constraints*: The PEMS should determine both energy consumptions ΔE so as to keep the probability of the PHEV running out of battery supply E_b in a ZEV zone close to zero. This can be naturally done by specifying an additional constraint similar to (1e).

For this purpose, we start from (1c), in which we replace E_{w_i} with \hat{E}_{w_i} ; and we introduce safety margins \bar{s}_i and \underline{s}_i , i.e. (positive) slack variables measuring the gap between the estimated demands $\hat{E}_{w_i}(\Delta E, \mathbf{x})$ and the bounds $\bar{E}_i, \underline{E}_i$:

$$\bar{s}_i = \bar{E}_i - \hat{E}_{w_i}(\Delta E, \mathbf{x}), \quad \underline{s}_i = \hat{E}_{w_i}(\Delta E, \mathbf{x}) - \underline{E}_i. \quad (5)$$

Instead of replacing the random variables $\hat{E}_{w_i}(\Delta E, \mathbf{x})$ in Eq. (5) with their expectations, we make *cautious (pessimistic) replacements* by using the quantiles $q_{i, 1-\alpha}$ and $q_{i, \alpha}$:

$$\bar{s}_i = \bar{E}_i - q_{i, 1-\alpha}(\Delta E, \mathbf{x}), \quad \underline{s}_i = q_{i, \alpha}(\Delta E, \mathbf{x}) - \underline{E}_i;$$

these replacements guarantee that each constraint in Eq. (5) is violated with probability α —which we obviously set to a small value. Both quantiles obviously depend on ΔE_i :

$$\begin{aligned} \mathbf{q}_{i,1-\alpha}(\Delta E, \mathbf{x}) &= E_{W_i^*}(\Delta E, \mathbf{x}) + u_{1-\alpha} \cdot \sigma_i(\Delta E, \mathbf{x}), \\ \mathbf{q}_{i,\alpha}(\Delta E, \mathbf{x}) &= E_{W_i^*}(\Delta E, \mathbf{x}) - u_{1-\alpha} \cdot \sigma_i(\Delta E, \mathbf{x}), \end{aligned}$$

with $u_{1-\alpha}$ the $1 - \alpha$ -level quantile of the standard Gaussian distribution, and $\sigma_i(\Delta E, \mathbf{x})$ the standard deviation of the distribution of $\tilde{w}_i(\Delta E, \mathbf{x})$. Finally, using the positivity constraints $\bar{s}_i \geq 0$ and $\underline{s}_i \geq 0$, and Eq. (4b), we get

$$E_i(\Delta E) \leq \bar{E}_i + \varepsilon_{E,i} - u_{1-\alpha} \cdot \sigma_i(\Delta E, \mathbf{x}), \quad (6a)$$

$$E_i(\Delta E) \geq \underline{E}_i - \varepsilon_{E,i} + u_{1-\alpha} \cdot \sigma_i(\Delta E, \mathbf{x}). \quad (6b)$$

This set of constraints can be seen as a generalization of (1e), where $1 - \alpha$ determines the level of cautiousness of the PEMS strategy, whereas $\varepsilon_{E,i}$ adjusts the strategy so as to find a feasible solution in presence of uncertainty (which is antagonistic). For this reason, we replace (1e) with (6a)-(6b). In summary, our cautious PEMS strategy solves:

$$\min_{\Delta E} -E_{fn}, \quad \text{s. t. } \forall i = 1, \dots, n, \quad (7a)$$

$$\Delta E_i - W_i^*(\Delta E, \mathbf{x}) \in [-\varepsilon_{\Delta E,i}, \varepsilon_{\Delta E,i}], \quad (7b)$$

$$E_i(\Delta E) - E_{W_i^*}(\Delta E, \mathbf{x}) \in [-\varepsilon_{E,i}, \varepsilon_{E,i}]. \quad (7c)$$

$$\Delta E_i \in [\underline{\Delta E}_i, \bar{\Delta E}_i], \quad (7d)$$

$$E_i(\Delta E) + u_{1-\alpha} \cdot \sigma_i(\Delta E, \mathbf{x}) \leq \bar{E}_i + \varepsilon_{E,i}, \quad (7e)$$

$$E_i(\Delta E) - u_{1-\alpha} \cdot \sigma_i(\Delta E, \mathbf{x}) \geq \underline{E}_i - \varepsilon_{E,i}. \quad (7f)$$

B. Practical implementation

The cautious PEMS problem defined above by (7) has nonlinear inequality constraints. It is not necessarily convex, pending on the monotonicity (and smoothness) of the GPR prediction function. To cope with nonconvexity, we first transform the constrained PEMS problem into an unconstrained problem using Lagrangian relaxation, and use a gradient descent algorithm with momentum.

In a standard form, the PEMS problem minimizes the cost function $f_0(\Delta E) = -E_{fn}$ (7a), and is subject to five constraint functions $g_1(\Delta E), \dots, g_5(\Delta E)$ associated to each of the inequality constraints (7b) to (7f), which become positive when a solution does not comply with the constraints. The PEMS problem is thus relaxed into an unconstrained problem, with noncompliant constraints being penalized using rectified linear units (ReLU) integrated into the objective function using positive multipliers $\lambda_0, \dots, \lambda_5$:

$$J(\Delta E) = \lambda_0 \cdot f_0(\Delta E) + \sum_{p=1}^5 \lambda_p \cdot \text{ReLU}(g_p(\Delta E)).$$

The logarithmic transformation of this function is minimized via the Adam gradient descent algorithm with adaptive momentum from TensorFlow [52].

In the experiments, the multipliers are set to $\lambda_0 = 1$, $\lambda_1 = 5$, $\lambda_2 = 5$, $\lambda_3 = 10$, $\lambda_4 = 50$, and $\lambda_5 = 100$.


The learning rate of the Adam algorithm is set to 0.05, with exponential decay rates of 0.9 for the first moment estimates, and of 0.999 for the second moment estimates. The algorithm iterates over 1450 steps for each of the PEMS optimizations reported in the next section. We also set $\varepsilon_{\Delta E_f,i} = 0.5\text{L}/100\text{km}$, $\varepsilon_{\Delta E_b,i} = 0.1\text{kWh}/\text{km}$, $\varepsilon_{E_f,i} = 10^{-4}\text{L}$, and $\varepsilon_{E_b,i} = 10^{-5}\text{kWh}$.

IV. RESULTS

A. Accuracy of the energy demand models

In the experiments, both GPR models for the local power demands (3) are trained from data collected with a Renault PHEV in actual driving conditions on open road, for a total of 3877 sections. We selected for training a representative sample of the data using uncertainty sampling, as follows. We started with a small subset of training samples, and then iteratively repeated (1) training a GPR model and (2) processing the remaining samples with it so as to identify the samples with highest predictive uncertainty: these latter were then incorporated to the training set. Using this procedure, we reduced the training data to a subset of 8% (296 sections) of the initial dataset³. The remaining 92% were used to assess model performance.

Figure 1 displays the residuals for both GPR models. The model for w_f^* achieves a root mean squared error (RMSE) of 1.42 L/100km (see Figure 1a); the model for w_b^* , a RMSE of 0.046 kWh/km (see Figure 1b).



(a) residuals of w_f^* [L/100km] (b) residuals of w_b^* [kWh/km]

Fig. 1: Fuel and battery model accuracy

These prediction accuracies on each road section are acceptable. Note that the relative uncertainty on the global power demand at the end of the trip decreases with the distance traveled, uncertainties compensating each other over the trip. The remaining uncertainty on the predictions \hat{w}_b and \hat{w}_b , predicted according to each context by the GPR, will be propagated into the cautious PEMS.

³We stress that this training set size is sufficient to achieve precise predictions, and thus to have the chance-constrained programming method (at the core of this paper) exhibit good performances.

B. Simulated trip

Figure 2 displays a simulated trip around the current low emission zone of the “Métropole du Grand Paris”, scheduled to become a ZEV zone by 2030. It starts at the Université de Technologie de Compiègne (A), and ends at the historical headquarters of Renault in Boulogne-Billancourt (B). The future ZEV zone is represented by the blue shaded area. This simulated trip provides us with the

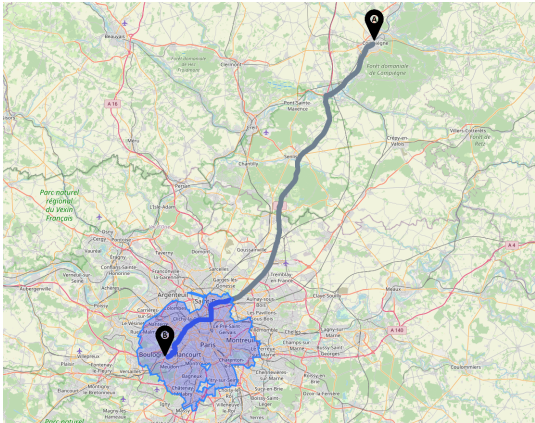


Fig. 2: Simulated trip around a low emission zone of the Paris area, scheduled to become a ZEV zone by 2030.

vectors of driving contexts \mathbf{x} fed to both GPR models in the PEMS optimization. We assume that the PHEV starts with $E_{bini} = 50\% \Delta E_b$ and $E_{fini} = 75\%$ fuel, and reaches a charging station at the end of the trip.

C. Results

1) *PEMS modes*: We apply the optimization strategy described in Section III-B in three different cases. The first one (nominal) considers the blue-shaded zone in Fig. 2 as a low emission zone and not a ZEV zone. The fuel energy consumption can thus be optimized on the whole trip in an optimal way, i.e. by choosing $\alpha = 0.5 \Leftrightarrow u_{1-\alpha} = 0$. The second one (optimal) considers the blue-shaded zone as a ZEV zone: we thus forbid the use of ICE by reducing the upper bound of the fuel consumption constraint (7d) to $\bar{E}_i = 0$; however, we solve the PEMS optimization problem without safety margins, i.e. $u_{1-\alpha} = 0$. The third one (optimal-cautious) implements our cautious PEMS: we consider the zone as a ZEV zone, and solve the PEMS using a safety margin of $u_{1-\alpha} = 1.96$, corresponding to a 2.5% risk that the PHEV will lack battery potential in the ZEV zone. The PEMS results are displayed in Figure 3.

2) *Nominal case*: In the nominal case (Fig. 3a), the battery energy consumption is distributed evenly over the trip so as to minimize the global fuel energy consumption, leading to a remaining fuel energy potential at arrival of 69.5%. Fuel savings are notably achieved by recharging

the battery using the ICE in low power demand sections at a lower speed downhill, e.g. around kilometers 30 and 75.

3) *Optimal case*: In the optimal case (Fig. 3b), the entire battery consumption is also evenly distributed over the trip. The PEMS saves enough battery power to drive inside the ZEV zone in full-electric mode. Although the ZEV zone constraint prevents the battery from being recharged around kilometer 75, the same remaining fuel energy potential at arrival (69.5%) is reached. However, the credibility interval indicates a significant risk ($\alpha = 50\%$) to run out of battery before reaching a charging station.

4) *Optimal-cautious case*: In the optimal-cautious case (Fig. 3c), the additional safety margin in the PEMS results in a battery energy potential of 19.8% at arrival, which is sufficient to guarantee (with a 2.5% risk) that the vehicle has enough electric range to reach its destination. This comes at the price of increasing the remaining fuel energy potential at arrival by 0.9%, which seems reasonable when it comes to providing further guarantees against driving around the ZEV zone or stopping the vehicle.

V. CONCLUSION AND PERSPECTIVES

This article addresses the problem of navigating areas with zero-emission zones, in which the use of internal combustion engines are prohibited. It provides a solution to guarantee with a high probability that PHEVs can cross these areas with their engine remaining switched off.

Assuming that the future trip profile is given by the vehicle’s navigation system, we propose a cautious PEMS optimization strategy, which determines battery and fuel energy consumption plans that minimize the global fuel consumption while saving a sufficient electric range for ZEV zones. It takes into account the uncertainty over power demands estimated using two separate GPR models. The resulting chance-constrained optimization problem is solved using a simple gradient descent algorithm with momentum.

Using models trained on actual PHEV data, we evaluate our strategy on a simulated trip enclosing a future ZEV zone of the Paris area. Results show that it succeeds in saving battery energy potential to navigate a ZEV zone, with a low increase in the global fuel consumption plan.

Future work will focus on learning a joint GPR model over both fuel and battery power demands, so as to obtain linear relationships between both given driving conditions. This would result in sets of linear constraints, and therefore to a convex optimization problem. The strategy will also be compared with the dynamic programming strategy previously used, in particular to evaluate the gain in terms of computational complexity. We eventually plan to conduct additional experiments, in particular to validate our approach in a real-life setting involving poor traffic conditions.

VI. ACKNOWLEDGMENTS

We acknowledge the ANRT for funding, and Renault SAS for providing the data used to train the models. We

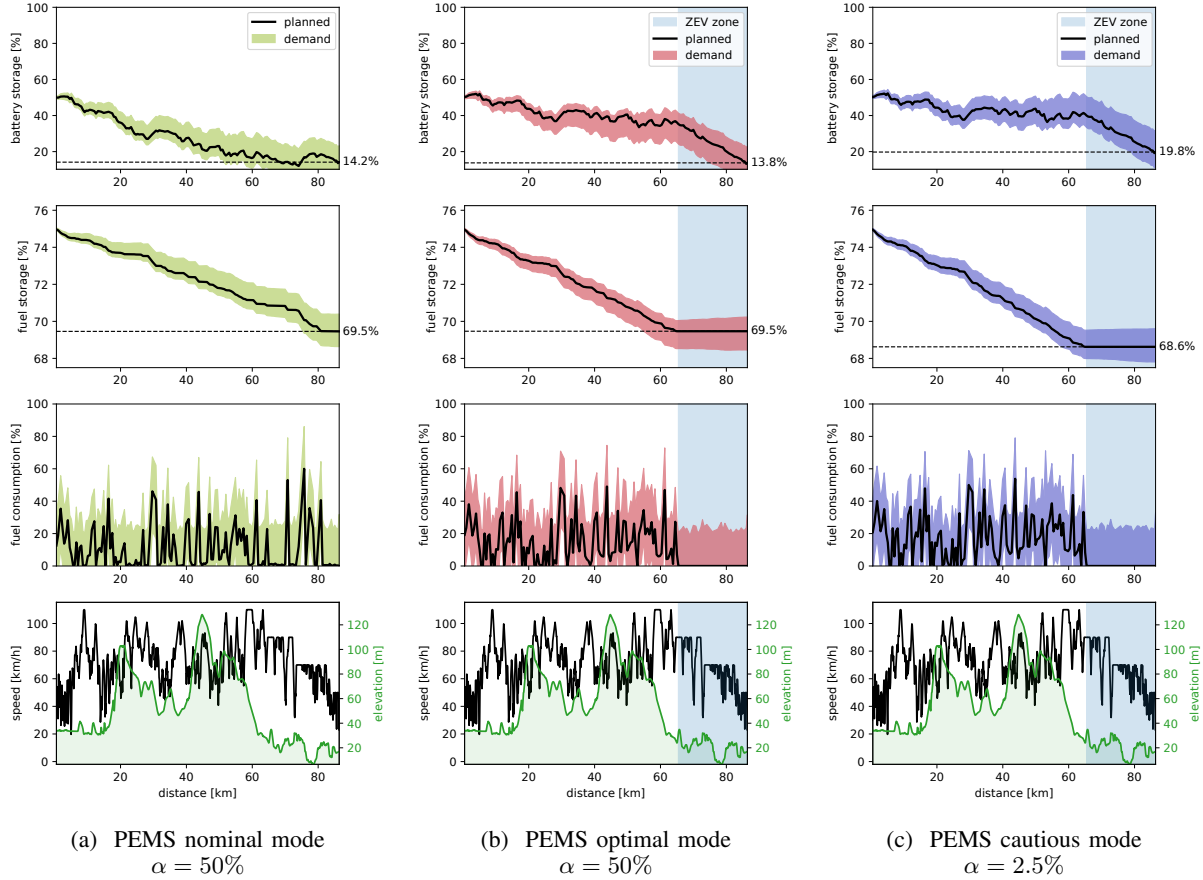


Fig. 3: Consumption trajectories provided by the PEMS in the three cases (nominal, optimal, optimal-cautious). The ZEV zones are represented with the blue areas on the x -axis. The planned consumption is represented by the black solid line, the 95% credibility intervals over the power demand by the colored (green, red and blue) areas. For a better comparison, the numerical expected value of the global energy potential at arrival is displayed in the same color in each figure.

also express our thanks to Romain Guillaume and Sylvain Lagrue for their sound advices on robust optimization.

REFERENCES

- [1] *Directive 2008/50/EC*, the European Parliament and the Council Std., May 2008.
- [2] M. Pouponneau, B. Forestier, and F. Cape, “Zones à faibles émissions (Low Emission Zones) à travers l’Europe,” Agence de l’Environnement et de la Maîtrise de l’Énergie (ADEME), Tech. Rep., 2020.
- [3] M. Browne, J. Allen, and S. Anderson, “Low emission zones: the likely effects on the freight transport sector,” *International Journal of Logistics Research and Applications*, vol. 8, no. 4, pp. 269–281, 2005.
- [4] D. Ku, M. Bencekri, J. Kim, S. Leec, and S. Leed, “Review of European Low Emission Zone Policy,” *Chemical Engineering Transactions*, vol. 78, pp. 241–246, Apr. 2020.
- [5] H. Cui, P. Gode, and S. Wappelhorst, “A global overview of zero-emission zones in cities and their development progress,” the International Council on Clean Transportation (ICCT), Tech. Rep., 2021.
- [6] *COM(2022) 586 final Proposal for a Regulation of the European Parliament and of the Council*, European Commission Std., Rev. 2022/0365 (COD), Nov. 2022.
- [7] A.-D. Ourabah, X. Jaffrezic, A. Gayed, B. Quost, and T. Denoeux, “Method for calculating a setpoint for managing the fuel and electricity consumption of a hybrid motor vehicle,” US Patent US20180281620A1, Oct., 2018.
- [8] H. A. Borhan, A. Vahidi, A. M. Phillips, M. L. Kuang, and I. V. Kolmanovsky, “Predictive energy management of a power-split hybrid electric vehicle,” in *2009 American Control Conference*, 2009, pp. 3970–3976.
- [9] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. the MIT Press, 2006.
- [10] M. Randon, B. Quost, N. Boudaoud, and D. v. Wissel, “Vehicle consumption estimation via calibrated Gaussian Process regression,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*, 2022, pp. 71–76.
- [11] D. Biswas, S. Ghosh, S. Sengupta, and S. Mukhopadhyay, “A Predictive Supervisory Controller for an HEV Operating in a Zero Emission Zone,” in *2019 IEEE Transportation Electrification Conference and Expo (ITEC)*, 2019, pp. 1–6.
- [12] A. Capancioni, “Development of predictive energy management strategies for hybrid electric vehicles supported by connectivity,” Ph.D. dissertation, Università di Bologna, 2022.
- [13] N. V. Sahinidis, “Optimization under uncertainty: state-of-the-art and opportunities,” *Computers & Chemical Engineering*, vol. 28, no. 6, pp. 971–983, 2004.
- [14] J. R. Birge and F. Louveaux, *Introduction to stochastic programming*. Springer Science & Business Media, 2011.

- [15] P. Kall and S. W. Wallace, *Stochastic programming*. John Wiley & Sons, 1994.
- [16] A. Charnes, W. W. Cooper, and G. H. Symonds, "Cost Horizons and Certainty Equivalents: An Approach to Stochastic Programming of Heating Oil," *Management Science*, vol. 4, no. 3, pp. 235–263, 1958.
- [17] A. Wald, "Statistical decision functions." 1950.
- [18] L. J. Savage, "The Theory of Statistical Decision," *Journal of the American Statistical Association*, vol. 46, no. 253, pp. 55–67, 1951.
- [19] D. Bertsimas and M. Sim, "The Price of Robustness," *Oper. Res.*, vol. 52, pp. 35–53, 2004.
- [20] A. Prékopa, "Logarithmic concave measures with applications to stochastic programming," *Acta Scientifica Mathematica*, vol. 32, pp. 301–316, 1971.
- [21] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.
- [22] A. Sciarretta and L. Guzzella, "Control of hybrid electric vehicles," *IEEE Control Systems*, vol. 27, pp. 60–70, 2007.
- [23] R. Bellman, "The theory of dynamic programming," *Bulletin of the American Mathematical Society*, vol. 60, no. 6, pp. 503–515, 1954.
- [24] A.-D. Ourabah, X. Jaffrezic, A. Gayed, B. Quost, and T. Denœux, "Method for calculating a setpoint for managing the fuel and electricity consumption of a hybrid motor vehicle," Patent WO2017001740A1, 2015.
- [25] R. A. Howard, *Dynamic programming and Markov Processes*. The Technology Press of The Massachusetts Institute of Technology and John Wiley & Sons, Inc., New York : London, 1960.
- [26] I. Chadès, G. Chapron, M.-J. Cros, F. Garcia, and R. Sabbadin, "MDPtoolbox: a multi-platform toolbox to solve stochastic dynamic programming problems," *Ecography*, vol. 37, no. 9, pp. 916–920, 2014.
- [27] E. Altman, *Constrained Markov decision processes*. CRC Press, 1999, vol. 7.
- [28] W. Wiesemann, D. Kuhn, and B. Rustem, "Robust Markov decision processes," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 153–183, 2013.
- [29] T. M. Moldovan and P. Abbeel, "Safe Exploration in Markov Decision Processes," *CoRR*, vol. abs/1205.4810, 2012.
- [30] R. Guillaume, C. Thierry, and P. Zieliński, "Robust material requirement planning with cumulative demand under uncertainty," *International Journal of Production Research*, vol. 55, no. 22, pp. 6824–6845, 2017.
- [31] F. Blahoudek, T. Brázdil, P. Novotný, M. Ornik, P. Thangeda, and U. Topcu, "Qualitative Controller Synthesis for Consumption Markov Decision Processes," in *Computer Aided Verification*, S. K. Lahiri and C. Wang, Eds. Springer International Publishing, 2020, pp. 421–447.
- [32] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe reinforcement learning via shielding," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018, pp. 2669–2678.
- [33] M. Turchetta, F. Berkenkamp, and A. Krause, "Safe Exploration in Finite Markov Decision Processes with Gaussian Processes," *CoRR*, vol. abs/1606.04753, 2016.
- [34] A. Wachi and Y. Sui, "Safe Reinforcement Learning in Constrained Markov Decision Processes," in *Proceedings of the 37th International Conference on Machine Learning*, vol. 119, 2020, pp. 9797–9806.
- [35] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained Policy Optimization," in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, 2017, pp. 22–31.
- [36] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," *Advances in neural information processing systems*, vol. 30, 2017.
- [37] B. Kouvaritakis and M. Cannon, "Model predictive control," *Switzerland: Springer International Publishing*, vol. 38, 2016.
- [38] A. Mesbah, K. P. Wabersich, A. P. Schöllig, M. N. Zeilinger, S. Lucia, T. A. Badgwell, and J. Paulson, "Fusion of machine learning and MPC under uncertainty: What advances are on the horizon?" in *American Control Conference (ACC)*, 2022, pp. 342 – 357.
- [39] L. Hewing and M. N. Zeilinger, "Stochastic Model Predictive Control for Linear Systems Using Probabilistic Reachable Sets," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 5182–5188.
- [40] R. Soloperto, M. A. Müller, S. Trimpe, and F. Allgöwer, "Learning-Based Robust Model Predictive Control with State-Dependent Uncertainty," *IFAC-PapersOnLine*, vol. 51, no. 20, pp. 442–447, 2018.
- [41] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, "Robust Constrained Learning-based NMPC enabling reliable mobile robot path tracking," *International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.
- [42] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious Model Predictive Control Using Gaussian Process Regression," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2736–2743, 2020.
- [43] S. Kamthe and M. Deisenroth, "Data-Efficient Reinforcement Learning with Probabilistic Model Predictive Control," in *Proceedings of Machine Learning Research (PMLR)*, vol. 84, 2018, pp. 1701–1710.
- [44] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-based model predictive control for safe exploration," in *2018 IEEE conference on decision and control (CDC)*, 2018, pp. 6059–6066.
- [45] D. D. Fan, A. akbar Agha-mohammadi, and E. A. Theodorou, "Deep Learning Tubes for Tube MPC," *CoRR*, vol. abs/2002.01587, 2020.
- [46] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.
- [47] S. V. Koevring, Y. Lyu, W. Luo, and J. Dolan, "Provable Probabilistic Safety and Feasibility-Assured Control for Autonomous Vehicles using Exponential Control Barrier Functions," *2022 IEEE Intelligent Vehicles Symposium (IV)*, pp. 952–957, 2022.
- [48] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," *Automatica*, vol. 129, p. 109597, 2021.
- [49] G. Männel, J. Graßhoff, P. Rostalski, and H. S. Abbas, "Iterative Gaussian Process Model Predictive Control with Application to Physiological Control Systems," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 2203–2210.
- [50] W. Hashimoto, K. Hashimoto, Y. Onoue, and S. Takai, "Learning-Based Iterative Optimal Control for Unknown Systems Using Gaussian Process Regression," in *2022 European Control Conference (ECC)*, 2022, pp. 1554–1559.
- [51] F. S. Barbosa and J. Tumova, "Risk-Aware Navigation on Smooth Approximations of Euclidean Distance Fields Among Dynamic Obstacles," KTH Royal Institute of Technology and Digital Futures, Stockholm, Sweden.
- [52] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2014.