

Backpropagation-Based Analytical Derivatives of EKF Covariance for Active Sensing

Jonas Benhamou, Silvère Bonnabel and Camille Chapdelaine

Abstract—To enhance accuracy of robot state estimation, active sensing (or perception-aware) methods seek trajectories that maximize the information gathered by the sensors. To this aim, one possibility is to seek trajectories that minimize the estimation error covariance matrix output by an extended Kalman filter (EKF), w.r.t. its control inputs over a given horizon. However, this is computationally demanding. In this article, we derive novel backpropagation analytical formulas for the derivatives of the covariance matrices of an EKF w.r.t. all its inputs. We then leverage the obtained analytical gradients as an enabling technology to derive perception-aware optimal motion plans. Simulations validate the approach, showcasing improvements in execution time, notably over PyTorch’s automatic differentiation. Experimental results on a real vehicle also support the method.

Keywords: perception-aware, extended Kalman filter, trajectory optimization, backpropagation.

I. INTRODUCTION

In robotics, perception-aware (PA) approaches, [1]–[4], or active sensing approaches, seek trajectories that maximize information gathered from sensors so as to perform robotic tasks safely. Notably, in the context of ground vehicles, when localization is based on ranging or bearing measurements relative to beacons, the efficiency of active sensing has been shown by [5], [6]. In [7], trajectories are generated to perform optimal online calibration between GPS and inertial measurement unit (IMU), see also [8]. In [3], for visual-inertial navigation systems, the authors have optimized the duration in which landmarks remain within the field of view. In the context of simultaneous localization and mapping (SLAM), those methods pertain to active SLAM, see [9].

One way to attack active sensing is through the use of Partially Observable Markov Decision Processes (POMDPs) [10], see [11], which offer a proper mathematical framework, but whose complexity is often prohibitory [12]. Sampling-based planners [13], [14] may be subject to the same issues. A more tractable option, that we presently adopt, is to work with the widespread extended Kalman filter (EKF). It estimates in real time the state x_n from various sensor measurements, and (approximately) conveys the associated extent of uncertainty through the state error covariance matrix P_n . The magnitude of P_n may then be used as an objective to minimize, as advocated in, e.g., [15].

J. Benhamou and S. Bonnabel are with Mines Paris, PSL Research University, Centre for Robotics, 60 bd Saint-Michel, 75006 Paris, France. silvere.bonnabel@minesparis.psl.eu.

J. Benhamou and C. Chapdelaine are with SAFRAN TECH, Groupe Safran, Rue des Jeunes Bois - Chateaufort, 78772 Magny Les Hameaux CEDEX, France. camille.chapdelaine,jonas.benhamou@safranrgroup.com.

To derive trajectories that minimize P_n , a first step is to compute the gradients of P_n with respect to the control inputs. To this aim, several approaches are possible. One can use a brute force approach (in the vein of [16]) or numerical differentiation, that may be ill-conditioned and intractable, owing to the complexity of the EKF’s equations. To get around those problems, [15] advocates using backpropagation, through automatic differentiation, and argues that deriving analytical expressions would be difficult.

We fill a gap in this paper, by providing novel closed-form analytical expressions over a horizon for the derivatives of any smooth function of the covariance matrices of an EKF, w.r.t. all control inputs, leveraging the powerful backpropagation method. Those equations partly extend our recent results for the linear Kalman filter [17] to the nonlinear case, using an EKF. Besides being analytical, these equations lead to further speedups even over automatic differentiation, as we will show in this paper. We will then use those gradients to solve perception-aware optimal motion planning.

Note that some active sensing methods revolve instead around the Observability Gramian (OG), used to elicit observability as in [5], or as a surrogate for P_n , as advocated in [5]–[7], [18], [19]. Although the covariance matrix and the OG may be related, see [5], or see the probabilistic interpretation of the OG in [20], we prefer to focus on the covariance matrix, as advocated in [5]–[7], [19], that more realistically captures the noise characteristics, see [21].

The main contributions of this paper are as follows:

- Deriving novel analytical backpropagation equations for the gradient of the covariance of an EKF with respect to all inputs of the filter, including control variables, thus partly extending [17] to the relevant nonlinear context;
- Hence providing a computationally efficient and general method whose cost to compute the gradients w.r.t. *all* control inputs at once, over an N -step horizon, is similar to that of running one EKF over this horizon, i.e., of complexity $O(Nd^3)$ with d the dimension of the state.
- Applying the technique to derive a computationally efficient perception-aware method;
- Validate and compare the technique through simulations, and real experiments on a car-size ground vehicle.

Section II introduces the considered problems. Section III establishes backpropagation equations for computing this gradient. In Section IV, our path planning formulation is introduced. Finally, we demonstrate the benefits of our approach in simulations (Section V) and real experiments (Section VI).

II. CONSIDERED PROBLEMS

For partially observed linear dynamical systems affected by white Gaussian noise, the Kalman filter (KF) computes the statistics of the state given past observations, namely $p(x_n|y_0, \dots, y_n)$, in real time. The Kalman filter (KF) relies on parameters such as the process noise with covariance Q_n and the observation covariance R_n . There exists various approaches to compute the derivatives of the KF's outputs w.r.t. those parameters. The early sensitivity equations [22], see also [23], allow for computing the derivative of the likelihood $\mathcal{L} := \log p(y_0, \dots, y_n)$ w.r.t. the noise parameters. A much faster approach is to use backpropagation, either using numerical auto-differentiation, as advocated in [15], or closed-form formulas as very recently derived in [17].

In this paper, we target closed-form backpropagation formulas for computing the gradients of an EKF's final covariance w.r.t. its inputs, including the control inputs. This will provide a nontrivial (partial) extension of the results of [17] to nonlinear systems. We start with a few primers.

A. The Extended Kalman filter (EKF)

Let us consider a nonlinear discrete-time system:

$$\begin{cases} x_n = f(x_{n-1}, u_n, w_n), \\ y_n = h(x_n) + v_n. \end{cases} \quad (1)$$

where $x_n \in \mathbb{R}^p$ is the system's state, $u_n \in \mathbb{R}^q$ is the control input, $w_n \in \mathbb{R}^r$ represents the process noise which follows a Gaussian distribution with zero mean and covariance matrix Q_n . The measured output is denoted by $y_n \in \mathbb{R}^l$, corrupted by a Gaussian measurement noise v_n with zero mean and covariance matrix R_n . Owing to unknown noises corrupting the equations, and to the state being only partially observed through y_n , one needs to resort to a state estimator.

The Extended Kalman Filter (EKF) provides a joint estimation of the state, denoted as \hat{x}_n , and its covariance matrix, denoted as P_n , consisting of two steps.

At the propagation step, the estimated state is evolved through the noise-free model, that is,

$$\hat{x}_{n|n-1} = f(\hat{x}_{n-1|n-1}, u_n, 0), \quad (2)$$

and the covariance of the state error is evolved as

$$P_{n|n-1} = F_n P_{n-1|n-1} F_n^T + G_n Q_n G_n^T. \quad (3)$$

At the update step, the estimated state and the covariance matrix are updated in the light of observation y_n as

$$S_n = H_n P_{n|n-1} H_n^T + R_n, \quad (4)$$

$$K_n = P_{n|n-1} H_n^T S_n^{-1}, \quad (5)$$

$$\hat{x}_{n|n} = \hat{x}_{n|n-1} + K_n (y_n - h(\hat{x}_{n|n-1})), \quad (6)$$

$$P_{n|n} = (I - K_n H_n) P_{n|n-1}. \quad (7)$$

The Riccati update step (7) proves equivalent to the following update in so-called information form:

$$P_{n|n}^{-1} = P_{n|n-1}^{-1} + H_n^T R_n^{-1} H_n. \quad (8)$$

In these equations, matrices F_n , G_n and H_n are all jacobians that depend on state estimation \hat{x}_n and control inputs u_n :

$$\begin{aligned} F_n &= \frac{\partial f}{\partial x}(\hat{x}_{n-1|n-1}, u_n, 0), & G_n &= \frac{\partial f}{\partial w}(\hat{x}_{n-1|n-1}, u_n, 0), \\ H_n &= \frac{\partial h}{\partial x}(\hat{x}_{n|n-1}). \end{aligned} \quad (9)$$

Through the Jacobians, the control inputs affect the covariance matrices, which is in contrast with the linear case where all trajectories are equivalent. It thus makes sense to compute the sensitivity of covariance matrices w.r.t. control inputs.

B. Backpropagation based gradient computation

The final covariance computed in the EKF $P_{N|N}$ over a fixed window, say $n = 0, \dots, N$, is the result of an iterative algorithm (2)-(9), and can thus be viewed as a composition of many functions from $n = 0$ up to N , as layers in a neural network. As a result, it lends itself to backpropagation, a way of computing the chain rule backwards. Backpropagation ("backprop") is very efficient when there are numerous inputs and the output is a scalar function (as opposed to the more intuitive method of applying the chain rule forward). In [24], in the context of linear Kalman filtering, it is used to compute the gradient of the negative logarithm of the marginal likelihood (NLL) w.r.t all observations y_1, \dots, y_N . In [15], perception-aware trajectory generation is performed using an optimisation-based method where gradients are computed using an automatic differentiation algorithm.

C. Optimal motion planning problem

Let \mathcal{L} be a scalar function that reflects the magnitude of its argument. The considered minimum uncertainty (or active sensing) motion planning problem writes

$$\left\{ \begin{array}{l} \min_{u_1, \dots, u_N, x_1, \dots, x_N} \mathcal{L}(P_{N|N}) \\ \text{subject to} \\ \forall n \leq N \quad x_n = f(x_{n-1}, u_n, 0), \quad x_0 = x_I, \\ \forall n \leq N \quad P_{n|n-1} = F_n P_{n-1|n-1} F_n^T + G_n Q_n G_n^T, \\ \forall n \leq N \quad P_{n|n} = (I - K_n H_n) P_{n|n-1}, \\ \forall n \leq N \quad u_{\min} \leq u_n \leq u_{\max} \\ \text{with } F_n = \frac{\partial f}{\partial x_n}(x_{n-1}, u_n, 0), \\ G_n = \frac{\partial f}{\partial w_n}(x_{n-1}, u_n, 0), \quad H_n = \frac{\partial h}{\partial x_n}(x_n). \end{array} \right. \quad (10)$$

where N is the time horizon. The most common choices for \mathcal{L} to reflect the final uncertainty are the trace $\mathcal{L} = \text{Tr}(P_{N|N})$, used in [7], [15], [21], or the maximum eigenvalue, which has to be regularized using Schatten's norm, see [5]. As $\text{Tr}(P_{N|N})$ sums the diagonal terms, which are variances expressed in possibly different units and choice of scales, we renormalize the matrix, as suggested in [7], leading to the objective $\mathcal{L} = \text{Tr}(P_0^{-1} P_{N|N})$.

The difficulty of the (optimal control) problem above stems from the complicated relationship between the controls u_1, \dots, u_N and the final covariance matrix. The rationale in

the present paper is to derive an efficient method to compute all gradients $\frac{\partial \mathcal{L}}{\partial u_n}$, and then to perform gradient descent.

Although we focus on $P_{N|N}$, our method seamlessly applies to a loss that also depends on matrices $P_{n|n}$, $n \leq N$.

III. NOVEL SENSITIVITY EQUATIONS FOR THE EKF

We consider a fixed window, say, $n = 0 \dots N$, and we seek to differentiate a function $\mathcal{L}(P_{N|N})$ of the final uncertainty $P_{N|N}$, w.r.t. all the previous EKF inputs, that are, the noise parameters R_n, Q_n , the control inputs u_n , for $n = 1 \dots N$, and the initial values $\hat{x}_0, P_{0|0}$. They all affect $\mathcal{L}(P_{N|N})$ in a complicated manner, through the EKF equations (2) to (9). For instance a modification of u_1 , namely $u_1 + \delta u_1$ affects initial Jacobians R_0, G_0, H_0 in (9), that in turn affect all subsequent quantities output by the EKF through eq. (2) to (8), finally affecting $\mathcal{L}(P_{N|N})$.

To analytically compute the derivatives, there are two routes. The historical one is to forward propagate a perturbation, say δu_n , $n \leq N$, through the equations. It is known as the sensitivity equations, and has been done—at least—for the linear Kalman filter in [22], essentially for adaptive filtering. The other route is to compute derivatives backwards, using the backprop method. It is far less straightforward, and has been proposed only very recently, leading to drastic computation speedups, see [17], in the context of linear systems. In this paper we heavily rely on our prior work [17], but go a step further, by accommodating nonlinear equations, that is, going from the Kalman Filter to the EKF, with the additional difficulty that the Jacobians depend on the estimates. Our end goal being active sensing, though, we restrict ourselves to losses of the form $\mathcal{L}(P_{N|N})$. We show this additional dependency lends itself to the backprop framework too. We also extend the calculations to get the derivatives w.r.t. the control inputs.

A. Matrix derivatives

We now explain the methodology, the main steps, and provide the final equations. As deriving closed-form formulas through backward computation is somewhat tedious, a full-blown mathematical proof can be found in the online arxiv preprint of the present paper [25].

Our method heavily relies on two ingredients. First, the notion of derivative of a scalar function w.r.t. a matrix, and the associated formulas based on the chain rule. Then, dependency diagrams which encapsulate how the functions are composed.

Consider $\mathcal{L}(M)$, a scalar function of a matrix defined as $M = f(X, Y)$, where X and Y are also matrices. $\frac{\partial \mathcal{L}}{\partial M}$ denotes the matrix defined by $(\frac{\partial \mathcal{L}}{\partial M})_{ij} = \frac{\partial \mathcal{L}}{\partial M_{ij}}$. We write $\frac{\partial \mathcal{L}}{\partial X}$ to denote $\frac{\partial \mathcal{L} \circ f}{\partial X}$, otherwise notation would become impractical.

The chain rule provides rules of calculus for matrix

derivatives. We have the following formulas [26]:

$$\begin{aligned} M = XYX^T &\Rightarrow \frac{\partial \mathcal{L}}{\partial X} = \frac{\partial \mathcal{L}}{\partial M} XY^T + \frac{\partial \mathcal{L}}{\partial M}^T XY, \\ M = YXY^T &\Rightarrow \frac{\partial \mathcal{L}}{\partial X} = Y^T \frac{\partial \mathcal{L}}{\partial M} Y, \\ M = X^{-1} &\Rightarrow \frac{\partial \mathcal{L}}{\partial X} = -M^T \frac{\partial \mathcal{L}}{\partial M} M^T, \end{aligned} \quad (11)$$

In the particular case where X and M are vectors, we have:

$$M = f(X) \Rightarrow \frac{\partial \mathcal{L}}{\partial X} = J^T \frac{\partial \mathcal{L}}{\partial M}, \quad (12)$$

with J the Jacobian matrix of f w.r.t. X .

When the matrix X depends on a scalar variable s , we have the following formula:

$$M = X(s) \Rightarrow \frac{d\mathcal{L}}{ds} = \text{Tr} \left(\frac{dX^T}{ds} \frac{\partial \mathcal{L}}{\partial M} \right). \quad (13)$$

B. Backprop equations for the involved matrices

The graph in Figure 1 shows the relationships involved in the calculation of the state's error covariance, which is our variable of interest.




Fig. 1. Dependencies of EKF's variables in Riccati equations (3) and (8). Each variable (node) is a function of its predecessors.

The backprop method consists in running an EKF until time N , to fix the values of all variables, and then compute gradients backwards to get the derivatives. Let us assume we have computed an expression for $\frac{\partial \mathcal{L}}{\partial P_{N|N}}$ (initialization step) at final covariance $P_{N|N}$. It gives how a small variation in $P_{N|N}$ affects the objective, ignoring all the other variables. Starting with $n = N$, we go backwards as follows. As $P_{n|n}$ is a function of $P_{n|n-1}$ which is a function of $P_{n-1|n-1}$ in turn, through (3) and (8), we can use formulas (11) to assess how a small perturbation in $P_{n|n-1}$ and $P_{n-1|n-1}$ affects the loss in turn (as they affect subsequent quantities, whose variation on \mathcal{L} has been computed). This yields

$$\frac{\partial \mathcal{L}}{\partial P_{n|n-1}} = (I - K_n H_n)^T \frac{\partial \mathcal{L}}{\partial P_{n|n}} (I - K_n H_n), \quad (14)$$

$$\frac{\partial \mathcal{L}}{\partial P_{n-1|n-1}} = F_n^T \frac{\partial \mathcal{L}}{\partial P_{n|n-1}} F_n. \quad (15)$$

In the same way, we apply formula (11) to (3) and (8) to obtain the following relationships:

$$\frac{\partial \mathcal{L}}{\partial F_n} = \left(\frac{\partial \mathcal{L}}{\partial P_{n|n-1}} + \frac{\partial \mathcal{L}}{\partial P_{n|n-1}}^T \right) F_n P_{n-1|n-1}, \quad (16)$$

$$\frac{\partial \mathcal{L}}{\partial G_n} = \left(\frac{\partial \mathcal{L}}{\partial P_{n|n-1}} + \frac{\partial \mathcal{L}}{\partial P_{n|n-1}}^T \right) G_n Q_n, \quad (17)$$

$$\frac{\partial \mathcal{L}}{\partial H_n^T} = -2P_{n|n} \left(\frac{\partial \mathcal{L}}{\partial P_{n|n}} + \frac{\partial \mathcal{L}}{\partial P_{n|n}}^T \right) P_{n|n} H_n^T R_n^{-1}, \quad (18)$$

$$(19)$$

$$\frac{\partial \mathcal{L}}{\partial R_n} = R_n^{-1} H_n P_{n|n} \frac{\partial \mathcal{L}}{\partial P_{n|n}} P_{n|n} H_n^T R_n^{-1}. \quad (20)$$

Knowing $\frac{\partial \mathcal{L}}{\partial P_{n|n}}$ and $\frac{\partial \mathcal{L}}{\partial P_{n|n-1}}$, these equations allow in passing to calculate the partial derivative of the loss with respect to the intermediate variables F_n , G_n , H_n and R_n at each step.

C. Backprop equations for the vector variables

We can now compute the derivatives w.r.t. the state estimates \hat{x} and the control inputs u , which are vectors. However, a remark is in order. Our end goal is to derive optimal controls u_1, \dots, u_N that minimize loss $\mathcal{L}(P_{N|N})$. As this is performed ahead of time, the (noisy) observations y_n in (6) are not available. The most reasonable choice is then to plan using the *a priori* value of the y_n , i.e., $y_n = h(\hat{x}_{n|n-1})$. We may thus alleviate notation writing $\hat{x}_{n|n} = \hat{x}_{n|n-1} := \hat{x}_n$ and $\hat{x}_{n-1|n-1} := \hat{x}_{n-1}$.

The graph in Figure 1 only focuses on the covariance variables. If we step back, we see the Jacobians depend on the linearization point \hat{x}_{n-1} and the control inputs u_n , see (9). A bigger picture encapsulating all the dependencies in the EKF is represented in Figure 2.




Fig. 2. Dependency diagram of all the variables involved in an EKF.

First we compute $\frac{\partial \mathcal{L}}{\partial u_n}$. There is a general rule, derived from the chain rule, which is that $\frac{\partial \mathcal{L}}{\partial u_n}$ is the sum of all the derivatives w.r.t. the direct successors of u_n in the graph, see e.g., [17], provided they have been already computed in a previous step of the backward calculation. Additionally using (12) and (13), and computing w.r.t. to each scalar component u_n^k of vector u_n , this yields

$$\frac{\partial \mathcal{L}}{\partial u_n^k} = Tr \left(\frac{\partial \mathcal{L}^T}{\partial F_n} \frac{\partial F_n}{\partial u_n^k} \right) + Tr \left(\frac{\partial \mathcal{L}^T}{\partial G_n} \frac{\partial G_n}{\partial u_n^k} \right) + (J_n^u e_k)^T \frac{\partial \mathcal{L}}{\partial \hat{x}_n}. \quad (21)$$

where $J_n^u := \frac{\partial f}{\partial u} \Big|_{\hat{x}_{n-1}, u_n, 0}$ with \hat{x}_{n-1}, u_n computed when running the EKF forward, and e_k the k -th vector of the canonical basis (details are given the online preprint [25]).

Similarly, we compute the derivative w.r.t. the k -th scalar component of the system's state x_{n-1}^k , by adding terms corresponding to each successor in the graph:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \hat{x}_{n-1}^k} &= Tr \left(\frac{\partial \mathcal{L}^T}{\partial H_{n-1}} \frac{\partial H_{n-1}}{\partial \hat{x}_{n-1}^k} \right) + Tr \left(\frac{\partial \mathcal{L}^T}{\partial R_{n-1}} \frac{\partial R_{n-1}}{\partial \hat{x}_{n-1}^k} \right) \\ &+ (J_n^x e_k)^T \frac{\partial \mathcal{L}}{\partial \hat{x}_n} + Tr \left(\frac{\partial \mathcal{L}^T}{\partial F_n} \frac{\partial F_n}{\partial \hat{x}_{n-1}^k} \right) + Tr \left(\frac{\partial \mathcal{L}^T}{\partial G_n} \frac{\partial G_n}{\partial \hat{x}_{n-1}^k} \right). \end{aligned} \quad (22)$$

where $J_n^x := \frac{\partial f}{\partial x} \Big|_{\hat{x}_{n-1}, u_n, 0}$, which is equal to F_n .

D. Backprop initialization

To initialize the backward process, one needs to compute $\frac{\partial \mathcal{L}(P_{N|N})}{\partial P_{N|N}}$. In the case of the normalized trace, we have [26]

$$\frac{\partial \mathcal{L}(P_{N|N})}{\partial P_{N|N}} = \frac{\partial Tr(P_0^{-1} P_{N|N})}{\partial P_{N|N}} = P_0^{-1}. \quad (23)$$

In the case where one targets the maximum eigenvalue of $P_{N|N}$ as a minimization objective, the loss must be chosen consequently. To handle the non differentiability of this cost function, we resort to its regularized version using Schatten's norm, see [5] (details are given in our online preprint [25]).

Note also that at $n = N$, (22) needs to be adapted, as \hat{x}_N only has two successors in the graph, thus:

$$\frac{\partial \mathcal{L}}{\partial \hat{x}_N^k} = Tr \left(\frac{\partial \mathcal{L}^T}{\partial H_N} \frac{\partial H_N}{\partial \hat{x}_N^k} \right) + Tr \left(\frac{\partial \mathcal{L}^T}{\partial R_N} \frac{\partial R_N}{\partial \hat{x}_N^k} \right) \quad (24)$$

E. Final equations

The gradients may be obtained as follows. We first run the EKF forward, to get all the EKF variables given a sequence of inputs. Then, we may compute the derivative of the loss w.r.t. $P_{N|N}$ at the obtained final covariance matrix. Letting $n = N$, (18)-(20), provide the derivatives w.r.t. H_N and R_N , and (14) w.r.t. $P_{N|N-1}$. In turn (24) provides the derivative w.r.t. \hat{x}_N , so that (21) yields the gradient w.r.t. control input u_N . Continuing the process backward and using the equations above, we get the derivatives w.r.t. to all control inputs.

The process allows for drastic computation speedups, as the backward equations yield derivatives w.r.t. *all* control inputs in one pass only. By contrast, forward propagating perturbations would demand running an entire EKF-like process from scratch for $k = n$ to N to derive the derivative w.r.t. u_n . Akin to dynamic programming, backpropagation allows for reusing previous computations at each step.

IV. APPLICATION TO PERCEPTION-AWARE PLANNING

The computation of the gradients w.r.t. all the EKF's variables is a contribution in itself that may prove useful beyond active sensing. However, our present goal is to leverage it to address the perception-aware optimal path planning problem (10). This approach is known as partial collocation, as in [15], i.e., we explicitly include states as optimization variables and implicitly compute the covariance. A commonly employed method for solving this problem involves a first-order optimization algorithm, the main steps of which are outlined in Algorithm 1. From a computational perspective, the gradient computation stands out as the most computationally demanding step, hence the interest for the method developed above.

While (10) seeks trajectories that minimize the accumulated uncertainty on the robot state over a given horizon, we note that it is easy to add constraints or another term in the loss to perform a specific task. For example, adding the constraint $x_N = x_F$ allows for reaching a specific state while being perception-aware. Note it is also easy to make \mathcal{L} depend on previous $P_{n|n}$, $n < N$, as in [7].

A. Algorithm

With known gradients, a *first order* nonlinear optimization algorithm such as Sequential Quadratic Programming (SQP) may be brought to bear. This leads to Algorithm 1.

Algorithm 1 Path planning algorithm

Require: x_0, P_0, N, \mathcal{L}
 $u \leftarrow (u_0, \dots, u_N)$
while $\mathcal{L}(P_{N|N})$ not converge **do**
 $g \leftarrow \text{gradient_computation}(u)$
 $\alpha \leftarrow \text{line_search}(u, g)$
 $u \leftarrow \text{update}(u, g, \alpha)$ ▷ Using SQP for example
end while
Return: (u_0, \dots, u_N)

The process involves computing the gradient g of the cost function w.r.t. the control variables using forward and backward passes through the EKF. Additionally, we compute the gradient of the constraints with respect to decision variables. Subsequently, line search is conducted to determine an appropriate step size α for efficient convergence. However, line search requires evaluating the cost function, which corresponds to running a full EKF in our case. The decision variables are then updated using SQP. Further details are given in the experimental sections.

B. Discussion

In prior work, it proved difficult to use a loss \mathcal{L} depending on the state's covariance $P_{N|N}$. Indeed, the gradient computation of the loss w.r.t each control variable is expensive when using forward difference [7], [21], as explained in Section III-E. This has motivated [15] to use backpropagation, through automatic differentiation (AD). [15] also argues deriving analytical formulas would be difficult.

The interest of our work, that provides analytical formulas, is twofold in this regard. First, it is often preferable to have closed-form formulas when possible, to rule some numerical errors out and to keep better control over the calculation process (possibly opening up for some guarantees about the execution). Then, it leads to computation speedups, as will be shown experimentally in the sequel.

V. SIMULATION RESULTS

We now apply our results to the problem of wheeled robot localization. We consider a car-like robot modelled through the unicycle equations, and equipped with a GPS returning position measurements. There are two difficulties associated with this estimation problem. First, the heading is not directly measured. Then, and more importantly, we assume that the position of the GPS antenna in the robot's frame (we call lever arm) is unknown, or inaccurately known, or may slightly vary over time. The resulting problem pertains to simultaneous self-calibration and navigation, in the vein of [7] but in a simpler context. In straight lines, for instance, the lever arm is not observable, so perception-aware trajectories should lead to more accurate robot state estimation.

In this section we present the model, and we assess and compare our method through simulations. In the next section, we apply it to a real world off-road vehicle.

A. Bicycle model

The system state consists of the orientation (heading) $\theta_n \in \mathbb{R}$ in 2D, the position of the vehicle $p_n \in \mathbb{R}^2$, and the lever arm $l_n \in \mathbb{R}^2$. The control inputs are the steering angle v_n and the forward velocity μ_n . The kinematic equations based on a roll-without-slip assumption are as follows:

$$\begin{cases} \theta_n = \theta_{n-1} + \frac{dt}{L} (\mu_n + w_n^\mu) \tan(v_n + w_n^v), \\ p_n = p_{n-1} + dt \Omega(\theta_{n-1}) (\mu_n + w_n^\mu) e_1, \\ l_n = l_{n-1} \end{cases} \quad (25)$$

where $\Omega(\theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}$, L is the distance between the front and the rear wheels, $e_1 = (1, 0)^T$ indicates that the velocity is aligned with the robot's heading, and dt is the sampling time. A Gaussian white noise $w_n = (w_n^\mu, w_n^v)^T$ with covariance matrix Q corrupts the forward velocity μ_n and steering angle v_n to account for actuators' imperfections, and the mismatch with idealized kinematic model, i.e., slip.

Letting $l_n \in \mathbb{R}^2$ be the position of the GPS antenna in the vehicle's frame w.r.t. to its center (the midpoint of the rear axle), the position measured by the GPS is

$$y_n = p_n + \Omega(\theta_n) l_n + \varepsilon_n \quad (26)$$

where ε_n is a 2D white noise with covariance matrix R_n .

In the simulations, we let $L = 4\text{m}$, $dt = 1\text{s}$. In terms of noise parameters, we let $Q = \text{diag}(0.1, \pi/180)^2$, and R_n be the identity matrix, i.e., a standard deviation of 1 m for the GPS position measurements. To account for actuator physical limits, we assume $|v_n| \leq 30\pi/180$ rad and $0 \leq \mu_n \leq 5 \text{ m}\cdot\text{s}^{-1}$. To account for acceleration limits, we add the constraints $|\Delta v_n| \leq 15\pi/180 \text{ rad}\cdot\text{s}^{-1}$ and $|\Delta \mu_n| \leq 1 \text{ m}\cdot\text{s}^{-2}$.

B. Simulation results




Fig. 3. On the left, an example of an initial random guess in blue. The other two trajectories are solutions to the perception-aware problem where the loss is the trace (in orange) and the Schatten norm (in green). The right plot shows the expected trace of the covariance evolution for each trajectory.

We start by sampling admissible control inputs over a horizon $N = 150\text{s}$. Then, the corresponding trajectory is

obtained by integration. We then use the sequential least-squares programming (SLSQP) algorithm from Scipy [27] to optimize the sequence of controls to apply. The calculation of the gradient of the loss with respect to the control inputs is performed using the equations detailed in Section III. An example of the initial random trajectory and the solution of the perception-aware problem can be found in Figure 3. The optimal trajectory for the Schatten norm and for the trace are quite similar in this case, and both oscillate around the initial random trajectory. These oscillations are manoeuvres that increase the observability of the lever arm. They lead to an actual reduction of the expected theoretical error covariance.

To demonstrate that the final covariance minimization translates into an actual reduction of the average state estimation error, we simulate 200 trials of each optimal trajectory by adding process noise and observation noise. During the simulation, the state is estimated with an EKF based on (25), (26). The evolution of the absolute estimation error of the lever arm is displayed on Figure 4. For PA trajectories, the error decreases and converges more rapidly, illustrating the benefit of perception-aware optimal trajectory generation.

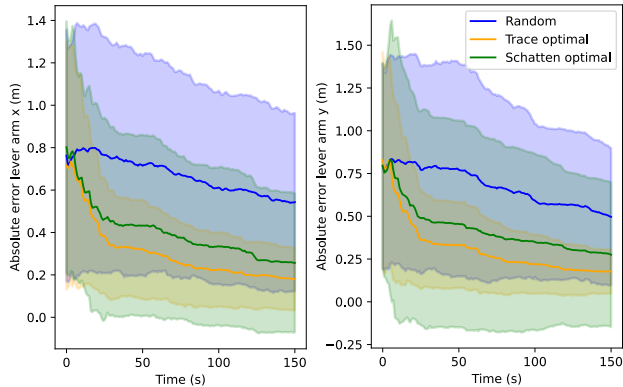


Fig. 4. Absolute estimation error of the lever arm during the trajectory. On the left the error for the lever arm in x and on the right in y . One- σ envelope illustrates the dispersion of errors over trials.

C. Computation time

We compared the computation times for the gradient of the loss w.r.t. control inputs with three different methods. The first method uses (forward) finite differences to compute the gradient, as in [28]. The second method is an automatic differentiation method, which adopts the backpropagation paradigm, but through a numerical tool, akin to [15]. Namely, we used the state-of-the-art PyTorch automatic differentiation (AD) [29]. Finally, the last method uses our backpropagation analytical formulas of Section III. The code was executed on computer with an Intel Core i5-1145G7 at 2.60 GHz.

TABLE I

AVERAGE AND STANDARD DEVIATION OF GRADIENT CALCULATION TIME (OVER 100 CALCULATIONS) USING DIFFERENT METHODS.

Method	Execution time
Finite differences	$26.92 \pm 8.45s$
PyTorch Autograd	$0.55 \pm 0.19s$
Ours	$0.19 \pm 0.07s$

Table I shows that gradient calculation using backpropagation (Autograd and ours) leads to large speedups. Indeed, when using a finite difference based method, optimization is computationally expensive, as mentioned in [7], where optimizing the trace of the sum of all covariances is reported to take 13 hours for a 3D inertial navigation system. Even when compared to state-of-the-art PyTorch Autograd, our closed-form formulas are much faster, and more stable in terms of variability of computation time.

The computation time for the complete resolution of the optimization problem (10) using our gradient calculation method and Algorithm 1 is 351s. Although this perfectly suits off-line trajectory generation, it means the robot should stop for a few minutes to plan in a real-time context.

D. Discussion

A few remarks are in order. First, we see that, by replacing state-of-the-art autograd differentiation with our formulas, one may (roughly speaking) double the planning horizon for an identical computation budget. Besides, analytical formulas better suit onboard implementation. It is interesting to note that they are totally akin to the EKF equations, which must be implemented on the robot anyway. Finally, analytical formulas—when available—may be preferable to numerical methods, as one keeps a better control over what is being implemented, possibly opening up for some guarantees (the behavior of autograd may be harder to anticipate, and leads to higher computation time variability). Moreover, we anticipate that coding them in C++ may lead to further speedups.

Note that the overall computation time of the path planning algorithm could be significantly reduced. First, we may obviously optimize over a shorter horizon. Then, the optimizer currently uses a line search algorithm to determine the descent step size. In our case, evaluating the objective function is computationally expensive because it requires running the EKF. Therefore, finding a method that reduces the number of calls to the objective function should prove efficient. Finally, another option to reduce computation time is to decrease the number of decision variables, by for instance parameterizing trajectories using B-splines, as in [5].

We may also comment on the obtained trajectories. As the optimization problem is highly nonlinear, non-convex, and constrained, one should expect an optimization method to fall into a close-by local minimum. In Figure 3, the local nature of the optimum proves visible, as the obtained trajectory oscillates around the initial trajectory. Methods to step out of local minima go beyond the scope of this paper. However, it is worth noting that albeit a (close-by) local minimum, the obtained trajectory succeeds in much reducing state uncertainty. It reduces the final average error on the lever arm, which is the most difficult variable to estimate, by a factor 3, see Figure 4.

In the context of real-time online planning, this suggests a sensible way to use the formulas of the present paper would be to compute a real-time trajectory that optimizes a control objective, and then to refine it in real time, by

performing a few gradient descent steps. This shall increase the information gathered by the sensors.

VI. REAL-WORLD EXPERIMENTS

Real experiments were conducted jointly with the company Safran, a large group that commercializes (among others) navigation systems. With the help from its engineers, we used an experimental off-road car owned by the company, which is approximately 4m long and 2.1m wide¹.

A. Experimental setting

The vehicle is equipped with a standard GPS, odometers, and a RTK (Real Time Kinematic) GPS, which is not used by the EKF, but serves as ground-truth for position owing to its high accuracy. The lever arm between the GPS and the RTK is denoted by $l_{GPS/RTK}$, and has been calibrated (it is only used for comparison to the ground-truth).

To further test our method, we conducted localization experiments on both ordinary and PA trajectories and compared their performance. We used the model described in Section V-A, and devised an EKF that fuses odometer data (dynamical model) with GPS position measurements. The vehicle state is represented by a 5-dimensional vector: two dimensions for position, one for the orientation, and two for the lever arm between the GPS and the center of the vehicle frame (midpoint of the rear axle).

Initially, we generated ordinary trajectories at two different speeds (5 km/h, 10 km/h). Subsequently, we employed our PA path planning Algorithm 1 initialized with those trajectories, to address optimization problem (10), using the final covariance trace as the loss. The constraints used were those relative to the actuators. To demonstrate the ability of optimization-based planning methods to handle further constraints, we also constrained the distance between the initial trajectory and the PA trajectory to be less than 1.5m. Then, we used local tangent plane coordinates computed near the experiment locations to convert 2D planning to world frame coordinates. Then, we followed each trajectory and computed the localization error committed by the EKF.

To measure localization accuracy, we utilized the RTK-GPS system as a ground-truth, since its uncertainties are of the order of a few centimeters. The position error of the vehicle was computed as follows:

$$e_n = \left\| p_n^{RTK} - (\hat{p}_n + \Omega(\hat{\theta}_n)(\hat{l}_n + l_{GPS/RTK})) \right\| \quad (27)$$

B. Results

The off-road experiment was performed in a field. It consists of 4 runs: 2 reference runs and 2 perception-aware runs (at 5 km/h and 10 km/h). Trajectories returned by RTK-GPS over the first run are displayed in Figure 5. We note the oscillations of the PA trajectory around the reference trajectory, enhancing state observability.

Figure 6 shows the trajectories for the second run, along with the evolution of the trace of $P_{n|n}$ output by the EKF

¹Because of confidentiality requirements, the company has not wished to publish a picture of its experimental vehicle.




Fig. 5. Off-road trajectories of the RTK-GPS in a local tangent plane coordinates oriented East-North-Up of the scenario at 5 km/h.

(on the first run, the latter is wholly similar and was not included owing to space limitations, to improve legibility of Fig. 5). We see the improvement in terms of trace through the optimized trajectories. The visible wiggling of the covariance corresponds to the correction steps of the EKF. Indeed, the dynamical model is run at 100Hz, corresponding to the odometer frequency, while corrections are made when GPS data are available, at approximately 1Hz.




Fig. 6. On the left, trajectories of the RTK-GPS in local tangent plane coordinates at 10 km/h. On the right, the evolution of the covariance trace along each trajectory in the scenario at a speed of 10 km/h.

TABLE II
ESTIMATION ERROR ACCORDING TO (27) FOR ALL TRAJECTORIES

Speed (km/h)	Type	Mean error (m)
5	Ref	3.32
	PA	2.88
10	Ref	4.73
	PA	4.34

Although the trace of the covariance is smaller along the PA trajectory, it doesn't necessarily guarantee that the error with respect to ground truth is reduced. The uncertainty calculated by the filter assumes all random variables are Gaussian, and the EKF is based on approximations. For each trial, we calculated the average error obtained for each type of trajectory. The results, presented in Table II, point out that PA trajectories enhance localization accuracy. Indeed, in the scenario at 5 km/h, the error decreases by 13.24% between the reference and the PA trajectory, and by 8.25% in the scenario at 10 km/h.

C. Analysis

When compared with results from simulations, the improvements are smaller. Several factors may explain this difference. Firstly, in simulations, we calculated an average result using the Monte Carlo method, whereas here we have a single trial. Additionally, in simulations, model and noises are perfectly known, whereas in reality, they are not. The roll-without-slip assumption is especially challenged when using an off-road vehicle. This also makes the trajectory tracking quite imperfect. Finally, due to the limited area, we explicitly constrained the optimized trajectory to stay around the reference one, producing a trajectory that optimally “refines” the reference trajectory, limiting the accuracy increase.

VII. CONCLUSION

Our first contribution has been to introduce novel back-propagation analytical equations to compute the gradient of any loss based on the covariance of an EKF, w.r.t all its inputs. Beyond the theoretical contribution, they lead to actual numerical speedups. Our second contribution has been to leverage those formulas to address PA path planning, and to test the method in simulation and on real-world experiments over large off-road trajectories over a span of more than 50 meters. In future work, we would like to apply the method to more challenging problems, such as inertial navigation [7], and to improve its scalability, notably by improving the line search. We also would like to combine the method with other objectives such as reaching a desired goal, collision avoidance, or reducing energy consumption.

REFERENCES

- [1] L. Bartolomei, L. Teixeira, and M. Chli, “Perception-aware Path Planning for UAVs using Semantic Segmentation,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 5808–5815.
- [2] G. Costante, C. Forster, J. Delmerico, P. Valigi, and D. Scaramuzza, “Perception-aware Path Planning,” <https://arxiv.org/abs/1605.04151>, Feb. 2017.
- [3] V. Murali, I. Spasojevic, W. Guerra, and S. Karaman, “Perception-aware trajectory generation for aggressive quadrotor flight using differential flatness,” in *2019 American Control Conference (ACC)*, Jul. 2019, pp. 3936–3943.
- [4] J. Xing, G. Cioffi, J. Hidalgo-Carrió, and D. Scaramuzza, “Autonomous Power Line Inspection with Drones via Perception-Aware MPC,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2023, pp. 1086–1093.
- [5] P. Salaris, M. Cagnetti, R. Spica, and P. R. Giordano, “Online Optimal Perception-Aware Trajectory Generation,” *IEEE Transactions on Robotics*, vol. 35, no. 6, pp. 1307–1322, Dec. 2019.
- [6] O. Napolitano, D. Fontanelli, L. Pallottino, and P. Salaris, “Gramian-based optimal active sensing control under intermittent measurements,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, pp. 9680–9686.
- [7] J. Preiss, K. Hausman, G. Sukhatme, and S. Weiss, “Simultaneous self-calibration and navigation using trajectory optimization,” *The International Journal of Robotics Research*, vol. 37, Aug. 2018.
- [8] S. Xu *et al.*, “Observability-Aware Active Extrinsic Calibration of Multiple Sensors,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 2091–2097.
- [9] J. S. Willners *et al.*, “Adaptive Heading for Perception-Aware Trajectory Following,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 3161–3167.
- [10] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, 1998.
- [11] S. Candido and S. Hutchinson, “Minimum uncertainty robot path planning using a POMDP approach,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2010, pp. 1408–1413.
- [12] C. H. Papadimitriou and J. N. Tsitsiklis, “The Complexity of Markov Decision Processes,” *Mathematics of Operations Research*, vol. 12, no. 3, pp. 441–450, 1987.
- [13] R. He, S. Prentice, and N. Roy, “Planning in information space for a quadrotor helicopter in a GPS-denied environment,” in *2008 IEEE International Conference on Robotics and Automation (ICRA)*, May 2008, pp. 1814–1820.
- [14] S. Prentice and N. Roy, “The Belief Roadmap: Efficient Planning in Belief Space by Factoring the Covariance,” *The International Journal of Robotics Research*, vol. 28, no. 11-12, pp. 1448–1465, Nov. 2009.
- [15] S. Patil *et al.*, “Scaling up gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation,” in *Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics*, H. L. Akin, N. M. Amato, V. Isler, and A. F. van der Stappen, Eds. Cham: Springer International Publishing, 2015, pp. 515–533.
- [16] P. Abbeel, A. Coates, M. Montemerlo, A. Ng, and S. Thrun, “Discriminative Training of Kalman Filters,” in *Robotics: Science and Systems*, Jun. 2005, pp. 289–296.
- [17] C. Parellier, A. Barrau, and S. Bonnabel, “Speeding-Up Backpropagation of Gradients Through the Kalman Filter via Closed-Form Expressions,” *IEEE Transactions on Automatic Control*, vol. 68, no. 12, pp. 8171–8177, Dec. 2023.
- [18] M. Cagnetti, P. Salaris, and P. Robuffo Giordano, “Optimal Active Sensing with Process and Measurement Noise,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 2118–2125.
- [19] C. Böhm, P. Brault, Q. Delamare, P. R. Giordano, and S. Weiss, “COP: Control & Observability-aware Planning,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 3364–3370.
- [20] J. Benhamou, S. Bonnabel, and C. Chapdelaine, “Optimal Active Sensing Control for Two-Frame Systems,” in *2023 62nd IEEE Conference on Decision and Control (CDC)*, Dec. 2023, pp. 5967–5974.
- [21] M. Rafieisakhaei, S. Chakravorty, and P. R. Kumar, “On the Use of the Observability Gramian for Partially Observed Robotic Path Planning Problems,” in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, Dec. 2017, pp. 1523–1528.
- [22] N. Gupta and R. Mehra, “Computational aspects of maximum likelihood estimation and reduction in sensitivity function calculations,” *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 774–783, Dec. 1974.
- [23] J. V. Tsyganova and M. V. Kulikova, “SVD-Based Kalman Filter Derivative Computation,” *IEEE Transactions on Automatic Control*, vol. 62, no. 9, 2017.
- [24] C. Parellier, A. Barrau, and S. Bonnabel, “A Computationally Efficient Global Indicator to Detect Spurious Measurement Drifts in Kalman Filtering,” in *2023 62nd IEEE Conference on Decision and Control (CDC)*, Dec. 2023, pp. 8641–8646.
- [25] J. Benhamou, S. Bonnabel, and C. Chapdelaine, “Backpropagation-Based Analytical Derivatives of EKF Covariance for Active Sensing,” <https://arxiv.org/abs/2402.17569>, Feb. 2024.
- [26] K. B. Petersen and M. S. Pedersen, “[<http://matrixcookbook.com>],”
- [27] P. Virtanen *et al.*, “SciPy 1.0: Fundamental algorithms for scientific computing in Python,” *Nature Methods*, vol. 17, no. 3, pp. 261–272, Mar. 2020.
- [28] V. Indelman, L. Carlone, and F. Dellaert, “Towards Planning in Generalized Belief Space,” in *Robotics Research: The 16th International Symposium ISRR*, ser. Springer Tracts in Advanced Robotics, M. Inaba and P. Corke, Eds. Cham: Springer International Publishing, 2016, pp. 593–609.
- [29] A. Paszke *et al.*, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” Dec. 2019.