



HAL
open science

Risk-Sensitive Extended Kalman Filter

Armand Jordana, Avadesh Meduri, Etienne Arlaud, Justin Carpentier,
Ludovic Righetti

► **To cite this version:**

Armand Jordana, Avadesh Meduri, Etienne Arlaud, Justin Carpentier, Ludovic Righetti. Risk-Sensitive Extended Kalman Filter. 2023. hal-04107207

HAL Id: hal-04107207

<https://inria.hal.science/hal-04107207v1>

Preprint submitted on 26 May 2023

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.

Risk-Sensitive Extended Kalman Filter

Armand Jordana, Avadesh Meduri, Etienne Arlaud, Justin Carpentier and Ludovic Righetti

Abstract—In robotics, designing robust algorithms in the face of estimation uncertainty is a challenging task. Indeed, controllers often do not consider the estimation uncertainty and only rely on the most likely estimated state. Consequently, sudden changes in the environment or the robot’s dynamics can lead to catastrophic behaviors. In this work, we present a risk-sensitive Extended Kalman Filter that allows doing output-feedback Model Predictive Control (MPC) safely. This filter adapts its estimation to the control objective. By taking a pessimistic estimate concerning the value function resulting from the MPC controller, the filter provides increased robustness to the controller in phases of uncertainty as compared to a standard Extended Kalman Filter (EKF). Moreover, the filter has the same complexity as an EKF, so that it can be used for real-time model-predictive control. The paper evaluates the risk-sensitive behavior of the proposed filter when used in a nonlinear model-predictive control loop on a planar drone and industrial manipulator in simulation, as well as on an external force estimation task on a real quadruped robot. These experiments demonstrate the abilities of the approach to improve performance in the face of uncertainties significantly.

I. INTRODUCTION

Adapting the decisions robots make based on their perception of the world is key to deploying robots outside factories. More precisely, controllers should adapt to the degree of certainty or confidence of the robot’s belief of the world. For instance, it is important that a quadruped chooses conservative footholds and slows body movements when its confidence in the location of the ground decreases. Robust output feedback Model Predictive Control (MPC) studies methods that can adapt robot decisions based on the confidence of the perception module. However, the general nonlinear problem is very difficult, and practical algorithms remain often limited to linear systems [1].

The common practice in robotics is to decouple estimation and control (i.e., assume that the certainty equivalence principle holds) [2]–[6]. This approach is often chosen due to the availability of separate and tractable control and estimation algorithms that can be deployed on the robot. The estimation module is often a variation of a Gaussian filter, such as an Extended Kalman Filter (EKF) [7], which computes both

the mean and uncertainty of the state estimates from sensor information. In control, an increasingly popular approach is MPC, which consists in solving an optimal control problem numerically, e.g., using Differential Dynamic Programming (DDP) [3], [8]–[11], at each time step or at a fixed frequency. The controller can then adapt its behavior online based on the current robot and environment states. During deployment, the estimation module is used to compute the mean of the state estimate, which is then passed on to the controller to compute the optimal behavior [2]–[6]. Unfortunately, relying on the most likely outcome can lead to catastrophic behavior. For instance, on a load-carrying task with a quadruped where the load’s mass is unknown, the notion of mean might not be appropriate as this could lead the quadruped to apply insufficient force on the ground and then fall.

Some approaches try to address this issue by adding robustness or safety bounds in either the estimation or control block while keeping them independent. For instance, Robust Extended Kalman filtering [12] adds robustness to inaccuracies of the EKF or the model. However, the control objective is disregarded, and therefore the controller cannot be robust to estimation uncertainties. Robust MPC has been studied and applied to robots. [13] used tube-based MPC to control a biped. [14] used linear stochastic MPC to account for uncertainties in bipedal walking. However, this line of work assumes the state to be known. In contrast to such approaches, we aim to link estimation and control by adding into the estimation module a notion of control performance to improve robustness to the estimation uncertainty.

In this work, we leverage our previous theoretical results on dynamic game control with imperfect state observation [15] to introduce the Risk-Sensitive Extended Kalman Filter (RS-EKF), a novel filter that enables online risk-sensitive output feedback MPC. The RS-EKF computes state estimates robust to measurement uncertainty while taking into account the value function provided by the controller, i.e., the estimator tailors risk reduction to the control objectives. This, in turn enables automatic modification of robot decisions to be cautious in times of high environmental perturbation. Furthermore, RS-EKF has a similar computational cost to an EKF, allowing real-time deployment. To demonstrate the ability of the filter, we use it together with a DDP-based online non-linear controller to perform risk-sensitive output-feedback MPC on various simulated robots, such as a quadrotor subjected to arbitrary changes in its mass, and a KuKa robot facing unforeseen environmental disturbances. Finally, we test the filter on a real quadruped robot Solo12 [16] to perform an external force estimation and balancing task. These experiments demonstrate that the robots are more

This work was in part supported by the National Science Foundation grants 1932187, 1925079, 2026479, 2222815 and the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), the European Union through the AGIMUS project (GA no.101070165) and the Louis Vuitton ENS Chair on Artificial Intelligence.

Armand Jordana, Avadesh Meduri and Ludovic Righetti are with the Tandon School of Engineering, New York University, Brooklyn, NY. (e-mail: aj2988@nyu.edu, am9789@nyu.edu, lr114@nyu.edu).

Justin Carpentier and Etienne Arlaud are with Inria, Département d’informatique de l’ENS, École normale supérieure, CNRS, PSL Research University, Paris, France (e-mail: justin.carpentier@inria.fr).

robust to perturbations with the RS-EKF algorithm than a classical EKF. To the best of our knowledge, this is the first time that a non-linear risk-sensitive output-feedback MPC controller has been deployed on a robot.

II. BACKGROUND

A. Dynamic game output feedback MPC

To design a controller sensitive to the risk related to estimation uncertainty, Whittle [17] introduced a zero-sum game that aims at solving jointly the estimation and control problem. Given a history of measurements $y_{1:t}$, a history of control inputs, $u_{0:t-1}$ and a prior on the initial state \hat{x}_0 , we aim to find a control sequence $u_{t:T-1}$ that minimizes a given cost ℓ while an opposing player aims to find the disturbances $(w_{0:T}, \gamma_{1:t})$ that maximize this cost ℓ minus a weighted norm of the disturbances. Such a problem is formally written as:

$$\min_{u_{t:T-1}} \max_{w_{0:T}} \max_{\gamma_{1:t}} \ell_T(x_T) + \sum_{i=0}^{T-1} \ell_i(x_i, u_i) \quad (1)$$

$$- \frac{1}{2\mu} \left(\omega_0^T P^{-1} \omega_0 + \sum_{j=1}^t \gamma_j^T R_j^{-1} \gamma_j + \sum_{j=1}^T w_j^T Q_j^{-1} w_j \right)$$

$$\text{subject to } x_0 = \hat{x}_0 + w_0, \quad (2a)$$

$$x_{j+1} = f_j(x_j, u_j) + w_{j+1}, \quad 0 \leq j < T, \quad (2b)$$

$$y_j = h_j(x_j) + \gamma_j, \quad 1 \leq j \leq t. \quad (2c)$$

where $\mu > 0$. x_j is the state, ω_j the process disturbance, y_j the measurement, γ_j the measurement disturbance, T the time horizon, t the current time. The transition model f_j , the measurement model h_j , and the cost ℓ_j are assumed to be \mathcal{C}^2 . The measurement uncertainty R_j , the process uncertainty Q_j and the initial state uncertainty P are positive-definite matrices. This weighted sum of the disturbances can be seen as the estimation of maximum a posteriori probability (MAP) under Gaussian assumption. Hence, R_j , Q_j and P can be thought of as the covariance matrix of Gaussian noise.

Interestingly, this problem encompasses both formulations of control and estimation. If $t = 0$, in the limit where μ tends to zero, we find the generic OCP formulation [18] which directly minimizes the cost function assuming standard deterministic dynamics. And, if $t = T$ and if we consider all costs ℓ_i to be null, then, (1) is equivalent to maximizing the estimation maximum a posteriori probability (MAP). Here, the parameter μ is referred to as the risk-sensitive parameter and regulates how adversarial the problem is.

Whittle [17] provided a solution to this min-max problem in the case of linear dynamics and quadratic costs. This solution can be iteratively obtained with two recursions, one on past disturbances and one on future disturbances. Recently, [15] showed how this solution could be used to implement an efficient Newton's method that iteratively searches for a saddle point of the more general Problem (1). Exploiting the sparsity of a problem, the proposed Newton step has a linear complexity in the time horizon (note that naive optimization has at least a cubic complexity [19]). The

solution can then be interpreted as a risk-sensitive Kalman smoother coupled to minimax DDP [20].

In this work, we use these insights to derive a computationally efficient risk-sensitive extended Kalman filter that can be used for output feedback MPC. This is done by simplifying Problem (1) to match assumptions common to EKF and DDP while keeping the adversarial min-max formulation.

B. Extended Kalman Filter

We now discuss the structure of the EKF necessary to derive our filter. The EKF is usually derived by computing the probability a posteriori of the state given measurements, using the linearized dynamics and a Gaussian noise assumption [21]. However, the EKF can also be derived from an optimization point of view [22]. More precisely, the EKF can be seen as a Gauss-Newton step around a well-chosen point on the log-likelihood of the maximum a posteriori probability (MAP), i.e. $\log(p(x_t, x_{t-1} | y_t))$ [22]. Assuming disturbances follow Gaussian distributions, $\gamma_t \sim \mathcal{N}(0, R_t)$, $\omega_t \sim \mathcal{N}(0, Q_t)$, the MAP can be written as:

$$\max_{x_t, x_{t-1}} - (y_t - h_t(x_t))^T R_t^{-1} (y_t - h_t(x_t)) - (x_t - f_{t-1}(x_{t-1}, u_{t-1}))^T Q_t^{-1} (x_t - f_{t-1}(x_{t-1}, u_{t-1})) - (x_{t-1} - \hat{x}_{t-1})^T P_{t-1}^{-1} (x_{t-1} - \hat{x}_{t-1}) \quad (3)$$

where \hat{x}_{t-1} is the prior knowledge on the past state and P_{t-1} its associated covariance matrix. As shown in [22], a Gauss-Newton step around \hat{x}_{t-1} and $\bar{x}_t = f_{t-1}(\hat{x}_{t-1}, u_{t-1})$ on (3) leads to well-known recursion [21]:

$$\bar{P}_t = Q_t + F_{t-1} P_{t-1} F_{t-1}^T \quad (4)$$

$$K_t = \bar{P}_t H_t^T (R_t + H_t \bar{P}_t H_t^T)^{-1} \quad (5)$$

$$P_t = (I - K_t H_t) \bar{P}_t \quad (6)$$

$$\hat{\mu}_t = K_t (y_t - h_t(\bar{x}_t)) \quad (7)$$

$$\hat{x}_t = \bar{x}_t + \hat{\mu}_t \quad (8)$$

where $F_{t-1} = \partial_x f_{t-1}(\hat{x}_{t-1}, u_{t-1})$ and $H_t = \partial_x h_t(\bar{x}_t)$. Here, \hat{x}_t is the most likely estimate and P_t is the covariance uncertainty. In practice, at the next time step, \hat{x}_t is used as the prior knowledge on the state.

We notice the similar structure of the costs of Problem (1) and Eq. (3), except that the EKF only uses one measurement and does not include the control cost ℓ_j . Hence, Eq. (1) can be seen as a maximization of the estimation log-likelihood up to some cost terms. We will leverage this similarity to derive a risk-sensitive version of the EKF. More precisely, we will add cost-dependent terms in the maximization (3) in order for the filter to adapt to the control objective.

C. Nonlinear MPC

Once the state is estimated, it can then be used in the control module. In this work, we focus on controllers defined as the solution to a minimization problem. At each time step, a stagewise cost defined over a horizon H is minimized over future control inputs $u_t, u_{t+1}, \dots, u_{H-1}$:

$$\mathcal{L}_t(u_t, \dots, u_{H-1}) = \ell_{t+H}(x_{T+H}) + \sum_{j=t}^{t+H-1} \ell_j(x_j, u_j) \quad (9)$$

where the state sequence is implicitly defined by the dynamics $x_{j+1} = f_j(x_j, u_j)$ and where x_t is assumed to be the state estimated by the filter. At each time step, only the term u_t is applied to the real system. In the next time step, the state estimate is updated given the new measurement and the optimal control problem is solved again. There are various ways to solve efficiently this minimization problem. In robotics, a popular algorithm is DDP [23] which reassembles the Newton method but with linear complexity in the time horizon. Additionally, DDP provides a quadratic approximation of the value function which we exploit in our derivation of the risk-sensitive filter.

III. RISK SENSITIVE FILTER

We now introduce RS-EKF, which builds on the dynamic game defined in Equation (1). First, we modify the game to account for typical assumptions made for MPC while keeping the adversarial part that provides the risk-sensitive behavior. Then, we show how the solution can be computed with a Gauss-Newton step similar to the EKF, leading to an algorithm of similar complexity. Our formulation leads to a modified update in the filter of which the standard EKF can be seen as a limit case.

First, as for the EKF, we consider a history of measurements of length one. Furthermore, we disregard future uncertainties and assume deterministic dynamic equations for the future as is done in classical MPC formulations. Indeed, we expect that the high-frequency re-planning will compensate for model discrepancies. In the end, the problem is supposed to be adversarial only with respect to the uncertainties related to the estimation. This can be written as:

$$\min_{u_{t:t+H-1}} \max_{w_t} \max_{w_{t-1}} \max_{\gamma_t} \mathcal{L}_t(u_t, \dots, u_{H-1}) \quad (10)$$

$$- \frac{1}{2\mu} (\gamma_t^T R_t^{-1} \gamma_t + w_t^T Q_t^{-1} w_t + w_{t-1}^T P_{t-1}^{-1} w_{t-1})$$

$$\text{s.t. } x_{t-1} = \hat{x}_{t-1} + w_{t-1}, \quad (11a)$$

$$x_t = f_{t-1}(x_{t-1}, u_{t-1}) + w_t, \quad (11b)$$

$$y_t = h_t(x_t) + \gamma_t. \quad (11c)$$

$$x_{j+1} = f_j(x_j, u_j), \quad t < j < T. \quad (11d)$$

As presented in [15], one of the key features of the dynamic game is that some of the constraints can be removed with an appropriate change of variable. Indeed, we can use the equality constraints of Equations (11a), (11b) and (11c) to replace the disturbance maximization into a maximization over x_{t-1}, x_t :

$$\min_{u_{t:t+H-1}} \max_{x_{t-1}, x_t} \mathcal{L}_t(u_t, \dots, u_{H-1}) \quad (12)$$

$$- \frac{1}{2\mu} (y_t - h_t(x_t))^T R_t^{-1} (y_t - h_t(x_t))$$

$$- \frac{1}{2\mu} (x_t - f_{t-1}(x_{t-1}, u_{t-1}))^T Q_t^{-1} (x_t - f_{t-1}(x_{t-1}, u_{t-1}))$$

$$- \frac{1}{2\mu} (x_{t-1} - \hat{x}_{t-1})^T P_{t-1}^{-1} (x_{t-1} - \hat{x}_{t-1})$$

$$\text{subject to } x_{j+1} = f_j(x_j, u_j), \quad t < j < T,$$

By definition of the MAP [21], this can be written:

$$\begin{aligned} \min_{u_{t:t+H-1}} \max_{x_{t-1}, x_t} \mathcal{L}_t(u_t, \dots, u_{H-1}) - \frac{1}{\mu} \log(p(x_t, x_{t-1} | y_t)) \\ \text{subject to } x_{j+1} = f_j(x_j, u_j), \quad t < j < T. \end{aligned} \quad (13)$$

Problem (13) is intractable in the general case. However, by taking the concave-convex assumption, the minimization and maximization can be interchanged according to the min-max theorem. Consequently, the problem is equivalent to:

$$\max_{x_{t-1}, x_t} \log(p(x_{t-1}, x_t | y_t)) + \mu V_t(x_t), \quad (14)$$

where V_t is the value function of the OCP:

$$V_t(x_t) = \min_{u_{t:t+H-1}} \mathcal{L}_t(u_t, \dots, u_{H-1}) \quad (15)$$

Note that in the simplification from Eq. (1) to Eq. (10), it is not necessary to disregard future uncertainties as the value function could be the one resulting from minimax DDP [20]. If $\mu = 0$, we will obtain the unbiased estimate of Kalman filtering and the estimate will be independent of the control objective. Otherwise, if $\mu > 0$, the term $\mu V(x_t)$ will bias the estimate towards regions with higher value function, which in turn will force the controller to be more conservative.

We now take a Gauss-Newton step on the objective of Eq. (14) around the prior: \hat{x}_{t-1} and $\bar{x}_t = f_{t-1}(\hat{x}_{t-1}, u_{t-1})$. $V_t(x_t)$ is independent of x_{t-1} therefore, as shown in the appendix, the maximization over x_{t-1} can be simplified to:

$$\begin{aligned} \max_{x_t} - \frac{1}{2} (x_t - \hat{x}_t)^T P_t^{-1} (x_t - \hat{x}_t) \\ + \mu \frac{1}{2} (x_t - \bar{x}_t)^T V_t^{xx} (x_t - \bar{x}_t) + \mu (x_t - \bar{x}_t)^T v_t^x \end{aligned} \quad (16)$$

where \hat{x}_t and P_t are defined as in Eq. (8) and (6). where V_t^{xx} (respectively v_t^x) is the hessian (respectively the gradient) of the value function. Those are typically provided by optimal control algorithms such as DDP. In the end, the solution on the maximization over x_t is:

$$\hat{x}_t^{RS} = \bar{x}_t + (I - \mu P_t V_t^{xx})^{-1} (\hat{\mu}_t + \mu P_t v_t^x) \quad (17)$$

Interestingly, if $\mu = 0$, we recover the EKF. This was to be expected as, when μ tends to zero, the solution of problem (1) is exactly the solution of the neutral case where estimation and control are solved independently [17]. Otherwise, the estimate is shifted towards regions with higher cost values. Importantly, the magnitude of the shift depends on P_t the covariance matrix of the estimation. Note that μ cannot be arbitrarily large as $(I - \mu P_{t+1} V_{t+1}^{xx})$ needs to be positive definite. Larger values of μ would make the min-max problem defined in Eq. (1) ill-posed. More details on this limit value can be found in [17]. In the end, the estimate is shifted towards $P_t v_t^x$, i.e. towards a region with a larger cost function, and the magnitude of this shift is increased in the direction corresponding to large eigenvalues of $P_t V_t^{xx}$.

We obtained the solution to the maximization problem (13). Therefore, the cost function can now be minimized with respect to the control inputs by taking \hat{x}_t^{RS} as an initial condition of the optimal control problem, which can be solved for example with DDP.

Algorithm 1 summarizes the estimation procedure. It can then be used to do output-feedback MPC efficiently. At each time step, given a measurement, past control input, and a quadratic approximation of the value function, a risk-sensitive estimate can be computed. This estimate is then used to minimize the cost function (9) for MPC and the first control input is applied to the real system. Lastly, the quadratic approximation of the value function at $t + 1$ is saved as it will be used at the next estimation step.

Algorithm 1: Risk Sensitive EKF

Input: $\hat{x}_{t-1}, u_{t-1}, y_t, P_{t-1}, Q_t, R_t, V_t, v_t$

/ Predict */* */

- 1 $\bar{P}_t \leftarrow Q_t + F_{t-1}P_{t-1}F_{t-1}^T$
- 2 $\bar{x}_t \leftarrow f(\hat{x}_{t-1}, u_{t-1})$

/ Classical Update */* */

- 3 $K_t \leftarrow \bar{P}_t H_t^T (R_t + H_t \bar{P}_t H_t^T)^{-1}$
- 4 $P_t \leftarrow (I - K_t H_t) \bar{P}_t$
- 5 $\hat{\mu}_t \leftarrow K_t (y_t - h_t(\bar{x}_t))$

/ Value function bias */* */

- 6 $p_{x_t} \leftarrow (I - \mu P_t V_t^{xx})^{-1} (\hat{\mu}_t + \mu P_t v_t^x)$
- 7 $\hat{x}_t^{RS} \leftarrow \bar{x}_t + p_{x_t}$

Output: \hat{x}_t^{RS}, P_t

IV. EXPERIMENTS

This section presents simulation and real robot experiments to illustrate the benefits of the proposed algorithm and demonstrate its applicability to real problems. We study three test problems where we deploy the RS-EKF inside a MPC loop: a planar quadrotor with a load estimation task where we demonstrate qualitatively that the RS-EKF can bring conservatism appropriately in phases of high uncertainty, a push-recovery experiment on a 7-dof industrial manipulator on which we perform a quantitative study and lastly, an external force estimation task on a real quadruped robot to showcase the viability of the method on real systems. In each experiment, we use the DDP implementation provided by Crocodyl [24] to solve the optimal control problem given the filter estimate. All the code is available online¹.

A. Planar quadrotor

In this first scenario, we consider a planar quadrotor executing a load-carrying task. The goal is to move the quadrotor from position $(p_x, p_y) = (0, 0)$ to position $(1, 0)$ while carrying a load during the first half of the itinerary. The robot mass is 2 kg and the mass's load which is unknown a priori is 3 kg. The system dynamics is:

$$\begin{aligned}
 m\ddot{p}_x &= -(u_1 + u_2) \sin(\theta), \\
 m\ddot{p}_y &= (u_1 + u_2) \cos(\theta) - mg, \\
 m d\dot{\theta} &= r(u_1 - u_2),
 \end{aligned} \tag{18}$$

where m is the mass of the robot, d the distance between the rotors, θ the orientation of the quadrotor. u_1 and u_2 are the control inputs representing the force applied at each rotor.

In this experiment, we want to estimate online the mass parameter that changes in the middle of the flying phase. As it is standard in parameter identification [25], we augment the system's state with the unknown parameter and let it be estimated recursively by the filter (EKF or RS-EKF). The state of the system is thus: $x = (p_x \ p_y \ \theta \ \dot{p}_x \ \dot{p}_y \ \dot{\theta} \ m)^T$ and it is assumed that $\dot{m} = 0$ up to some random Gaussian noise. The dynamics are integrated with an Euler scheme and a time step of 0.05. We consider that $P_0 = 10^{-4}I_7$, $R = 10^{-4}I_3$, and Q is a 7×7 diagonal matrix where all terms are equal to 10^{-4} except the last one that we set to 2 to represent the uncertainty in the changes of the load. Lastly, we set $\mu = 4 \times 10^{-3}$. A stationary cost function of the following form is considered:

$$\begin{aligned}
 \ell(x, u) &= \alpha_1 (\|p_x - p_x^{des}\|^2 + \|p_y - p_y^{des}\|^2) + \alpha_2 \|\theta\|^2 \\
 &+ \alpha_3 (\|\dot{p}_x\|^2 + \|\dot{p}_y\|^2 + \|\dot{\theta}\|^2) + \alpha_4 \|u - \bar{u}\|^2
 \end{aligned} \tag{19}$$

where $\bar{u} = (\frac{mg}{2}, \frac{mg}{2})^T$ and where: $\alpha_1 = 100$, $\alpha_2 = 10$, $\alpha_3 = 0.01$ and $\alpha_4 = 0.1$. We consider a horizon of 20 nodes and re-plan at each new measurement, i.e. every 0.05s. Furthermore, we only measure: $y = (p_x \ p_y \ \theta)^T$ to illustrate the estimator capabilities. We simulate 4s with both output-feedback MPC controllers: the first one relying on the standard EKF estimate and the second one relying on the RS-EKF estimate.

Figure 1 shows the real mass variation and the estimates of both methods. It can be seen that the RS-EKF is more reactive when the load is added or dropped. The RS-EKF estimate spikes in the phases of uncertainty, which adds some robustness. Those spikes can be explained by the fact that the uncertainty increases on the components of the state that are important in the cost function. In other words, some of the eigenvalue of $P_t V_t^{xx}$ become larger in the phases of uncertainty, which augments the shift on the estimate as shown in Equation (17).




Fig. 1. Mass estimation for both EKF and RS-EKF.

Figure 2 depicts the trajectory in space for both methods. It can be seen that the controller relying on the risk-sensitive estimate is more reactive. To evaluate both controllers, we compute the average Mean Square Error (MSE) relative to the reference trajectory. The MSE of RS-EKF is 0.0011, and the one of EKF is 0.0024. Hence, in that situation, the risk sensitivity brings a 54% improvement in tracking. Furthermore, the average cost along the trajectory is equal to 0.0569 for the RS-EKF-based controller, while it is equal to 0.0880 for the EKF-based controller, yielding a 35%

¹<https://github.com/machines-in-motion/risk-sensitive-EKF>

improvement. This illustrates how a filter informed of the cost objective can improve the controller’s performance.




Fig. 2. Quadrotor trajectory for both the EKF-MPC and the RS-EKF-MPC.

As it can be seen in both Figures 1 and 2, RS-EKF introduces a steady state error. This is because the estimation uncertainty never goes to zero due to the nonzero diagonal term of the matrices Q and R . Intuitively, it makes sense that the controller that plans for the worst is slightly sub-optimal in the phase with no environment perturbation uncertainty. The risk-sensitive filter demonstrates its benefits when there are perturbations in the environment as this corresponds to phases where the most likely estimate might be far from reality. In the end, this example illustrates the advantages of having a risk-sensitive controller which is conservative only in phases with large environmental perturbations.

B. Kuka robot

In this example, we consider the 7-DoF torque-controlled KUKA LWR iiwa R820 14. The 14-dimensional state is composed of the joint positions and velocities. The control input is a 7-dimensional vector of the torque applied on each joint. The continuous dynamics and its analytical derivatives are provided by Pinocchio [26]. The goal is to track a reference trajectory with the end effector. To do so, we consider the following cost:

$$\begin{aligned} \ell_k(x_k, u_k) &= 10^{-2} \|x_k - \bar{x}\|_2^2 + 10^{-4} \|u_k - \bar{u}(x_k)\|_2^2 \\ &\quad + 10^2 \|p_k^{\text{target}} - \bar{p}(x_k)\|_2^2 \\ \ell_T(x_T) &= 10^2 \|p_T^{\text{target}} - \bar{p}(x_k)\|_2^2 + 10^{-2} \|x_k - \bar{x}\|_2^2, \end{aligned} \quad (20)$$

\bar{x} , the initial state, is used for the state regularization and is the concatenation of the initial configuration $(0.1 \ 0.7 \ 0 \ -1 \ -0.5 \ 1.5 \ 0)^T$ and a 7-dimensional zero vector corresponding to the velocity. $\bar{u}(x_k)$ is the gravity compensation term given by the rigid body dynamics. $\bar{p}(x_k)$ is the end-effector position obtained through forward kinematics. p_k^{target} is defined such that the end effector follows trajectories forming a circle in the xy plane. We use a horizon of 20 collocation point with an integration step of $0.05s$ and re-plan at 500 Hz.

In this experiment, we aim to showcase that the risk-sensitive filter can bring conservatism on phases with large perturbations from the environment which we simulate with large forces applied on the end effector. For the measurement model, we assume that all the states are observed with high accuracy, therefore, we set $R = P_0 = 10^{-6} I_{14}$. However, to model the disturbances in the dynamics, we set $Q = 10^{-1} I_{14}$. Finally, we consider $\mu = 7.5 \times 10^4$.




Fig. 3. End effector trajectory on a tracking task for both the EKF-MPC and the RS-EKF-MPC. An unexpected force is applied between 1s and 2s.

Figure 3 depicts the end effector trajectory for both controllers and their respective estimates. From time 1s to 2s, an external force of norm 80 is applied on the end-effector on the x and z direction. Both controllers are pushed away from the reference trajectory. However, the risk-sensitive estimate overestimates the distance between the reference and the end-effector. Consequently, the robot is more aggressive in its response, and the end-effector remains closer to the reference. This illustrates how taking a pessimistic estimate with respect to the cost improves the running cost. Note that both estimates are originally state estimates but are mapped to end effector space through the forward kinematics to draw Figure 3. The fact that those estimates are pessimistic with respect to the task originally defined in end effector space illustrates well how the method can handle nonlinear dynamics and cost functions.




Fig. 4. Median MSE on the tracking over 10,000 experiments with random external disturbances. The envelope represents the 25th and 75th percentiles.

To validate the consistency of the filter, 10,000 experiments are performed with random external forces. The timing and direction of the forces are uniformly sampled while the duration is fixed to 1s and the norm is fixed to 80. Figure 4 shows the median end effector error trajectory. In average, we obtain a 32% improvement in the MSE. Furthermore, the mean cost is 22% lower with the risk-sensitive filter.




Fig. 5. Comparison of both methods after an external force of 20N is applied by pulling the robot vertically. The vertical line indicates the moment when the robot is dropped. The top sub-figure overlays the trajectory of both the EKF in blue and RS-EKF in solid. The RS-EKF-based controller is more reactive to the perturbation and returns to the reference sooner.

C. Load estimation on a quadruped robot

In this experiment, we deploy the RS-EKF on a real 12-degree-of-freedom, torque-controlled quadruped robot - Solo12 [16]. We demonstrate the superior performance of the RS-EKF in estimating wrenches applied on Solo12 while it is standing. We generate the standing behavior on Solo12 using a non-linear MPC scheme. At each control cycle, we minimize a cost function using a centroidal model to compute the optimal forces and trajectory that keep the robot's base at a desired height and orientation. Additionally, we use an augmented state to estimate external forces-torques applied to the robot [27]:

$$\dot{c} = \frac{1}{m} l \quad (21)$$

$$\dot{l} = mg + \sum_{i=1}^{M_c} F_i + F_{\text{ext}} \quad (22)$$

$$\dot{k} = \sum_{i=1}^{M_c} (p_i - c) \times F_i + \tau_{\text{ext}} \quad (23)$$

$$\dot{F}_{\text{ext}} = 0, \quad \dot{\tau}_{\text{ext}} = 0 \quad (24)$$

where m denotes the mass, M_c the number of end effectors in contact. Each p_i represents a contact location. Here, the state is $x = (c \ l \ k \ F_{\text{ext}} \ \tau_{\text{ext}})^T$ which includes the Center of Mass (c), linear momentum (l), Angular Momentum (k) and External Force-torques ($F_{\text{ext}}, \tau_{\text{ext}}$). The measurement is $y = (c \ l \ k)^T$ up to some noise. A motion capture system measures the base position, velocity, and orientation and an IMU gives the orientation velocity. Joint encodings are provided and their velocities are derived with finite differences. Then, given, q, \dot{q} , we can compute c, l, k with

Pinocchio [26], [28] and can then be used as a measurement by the centroidal-based filter.

The control input, $u = (F_1 \ \dots \ F_{M_c})^T$, is a $M_c \times 3$ dimensional vector, representing the force applied at each end effector. For this experiment, the robot is standing, therefore, $M_c = 4$. The cost function for the OCP is:

$$\begin{aligned} \ell_t(x, u) &= (x - x^*)^T H_x (x - x^*) + (u - u^*)^T H_u (u - u^*) \\ &\quad + 10^5 \sum_{i=1}^{M_c} \ell_{\text{barrier}}(u_{3i}) \\ \ell_T(x) &= (x - x^*)^T H_x (x - x^*) \end{aligned} \quad (25)$$

where $H_x = \text{BlockDiag}(10^2 I_3, 10 I_6)$ and H_u is a diagonal matrix where the diagonal terms are made of M_c times the following sequence $(10^{-4}, 10^{-4}, 10^{-6})$. Lastly,

$$\ell_{\text{barrier}}(u) = \begin{cases} u^2 & \text{if } u < 0 \\ 0 & \text{if } 0 < u < 10 \\ (u - 10)^2 & \text{if } u > 10 \end{cases} \quad (26)$$

Here, x^* is designed to keep a constant height above the ground and to keep the base horizontal, u^* is the gravity compensation. The reference desired angular momentum for the OCP is adapted to bring the base back to a horizontal position. We do this by computing $k^* = \frac{1}{T} \log_3(R_t R_{\text{des}}^T)$, where R_t is the current base rotation matrix, R_{des} the rotation matrix corresponding to the quaternion $q_{\text{des}} = [0, 0, 0, 1]$ and T the horizon length. The \log_3 is a mapping from $SE(3)$ to $\mathfrak{se}(3)$. The nominal angular momentum aims to bring the base back to the desired orientation over the time horizon [29]. The ℓ_{barrier} is a quadratic barrier function that creates a soft constraint on the maximum forces the robot can apply on the ground.

We solve this OCP at 100 Hz using Crocodyl [24] and then track the desired forces using a task space impedance dynamics QP [30] that we solve at 1 kHz using [31].

$$\begin{aligned} \min_{f, \tau, a} \quad & \frac{1}{2} \|f - F\|^2 \\ \text{subject to} \quad & Ma + g = J^T f + S^T \tau + S^T f_{\text{friction}} \\ & Ja = -\dot{J}\dot{q}, \end{aligned} \quad (27)$$

where J is the robot contact Jacobian, M is the mass matrix, S is the selection matrix that projects on the actuated joints, and g is the robot gravity vector at the current time step. The static friction in each joint was estimated independently and is approximately equal to 0.07. However, to keep a continuous model, we consider $f_{\text{friction}} = -0.07 \frac{2}{\pi} \arctan(2S\dot{q})$. The two constraints ensure dynamics satisfaction and model that the end effectors do not move. We update our state estimate (including the external force torque) using the RS-EKF at 200 Hz with $\mu = 6$. We also use the EKF in place of the RS-EKF to compare the performance of a risk-sensitive filter while keeping the update frequency the same. For both filter, we consider the following parameters: $P_0 = Q = \text{BlockDiag}(10^{-3}I_6, 10^{-4}I_3, 10^{-1}I_3, 10^{-2}I_3)$ $R = \text{BlockDiag}(10^{-4}I_3, 10^{-2}I_3, 10^{-4}I_3)$.

The first experiment ran on the robot is shown in Fig. 5. Here, the base of Solo12 is pulled up (in the z direction) until an external estimate of 20N is computed by both the filters (Shown in bottom sub-figure 5 at time 1.4s with a vertical line). The base is then released to let Solo12 recover and bring its base back to the nominal desired height. The top sub-figure in Fig 5 shows an overlaid pictorial comparison of Solo12 when the two filters (RS-EKF and EKF) are used on the robot. Each frame corresponds to the state of Solo12 at the same time. The opaque Solo12 image shows the state with RS-EKF, and the blue one shows the EKF. This experiment shows that the RS-EKF helps the OCP to react quicker and bring the base to the nominal location sooner. This happens because the RS-EKF underestimates the base height in z as compared to EKF, which makes the OCP generate higher ground reaction forces to bring the base up sooner. In the end, both filter estimate similar external forces. As it can be seen in Fig. 5, the external vertical force does not converge exactly to zero. We find experimentally that this estimated force is due to friction. Lastly, as it can be seen, the cost of the RS-EKF-based controller is lower after the robot is dropped. The average cost of the RS-EKF is 0.065 while the one of the EKF is 0.130, which corresponds to a 50% improvement. This demonstrates how a filter aware of the cost objective can improve the overall performance.

In order to get a more systematic comparison between the filters and get rid of the human error, we perform two additional experiments where the filters are initialized with exactly the same priors. First, we initialize both the filters with a wrong prior on the external vertical force of 20N, while in reality, no force is applied on the robot. This experiment creates an identical situation as the previous experiment while also ensuring the very same initial conditions for the

robot. The results are shown in Fig 6, where the RS-EKF still performs better. Also, the performance is similar to the first experiment. In that experiment, we obtain a 62.9% improvement in the average cost.




Fig. 6. Comparison of the RS-EKF and EKF when initialized with a wrong prior of 20N on the estimated vertical external force.

Finally, we replicate the previous experiment but now, initialize the filters with a wrong prior of -10 N on the external force, while, in fact, there is no force on the robot. The RS-EKF reacts sooner than EKF once again. It brings the base of Solo12 back to the desired location sooner than EKF, as can be seen in Fig 7. In that experiment, we obtain a 58.9% improvement in the average cost.




Fig. 7. Comparison of the RS-EKF and EKF when initialized with a wrong prior of -10 N on the estimated vertical external force.

V. CONCLUSION

To conclude, we have introduced a risk-sensitive variation of the EKF based on the zero-sum game introduced by Whittle [17]. The filter biases estimates towards high regions of the control cost which result in more robust controllers. Furthermore, the complexity of this filter is similar to the EKF. Lastly, we have demonstrated on several robotics problems, both in simulation and on real hardware, the benefits of this filter for output-feedback MPC. Importantly, we have shown that the proposed filter makes the controller more conservative in phases of high uncertainty leading to better overall control performance. Future work will address scenarios involving more dynamic tasks and more complex measurement uncertainty.

APPENDIX

By taking a quadratic approximation of the value function, the Gauss-Newton step can be written as:

$$\begin{aligned} \max_{x_{t-1}} \max_{x_t} \quad & \mu(x_t - \bar{x}_t)^T V_t^{xx}(x_t - \bar{x}_t) + 2\mu(x_t - \bar{x}_t)^T v_t^x \\ & - (\Delta y - H_t \Delta x_t)^T R_t^{-1} (\Delta y - H_t \Delta x_t) - \Delta x_{t-1}^T P_{t-1}^{-1} \Delta x_{t-1} \\ & - (\Delta x_t - F_{t-1} \Delta x_{t-1})^T Q_t^{-1} (\Delta x_t - F_{t-1} \Delta x_{t-1}) \end{aligned} \quad (28)$$

where $\Delta y = y_t - h(\hat{x}_t)$, $\Delta x_{t-1} = x_{t-1} - \hat{x}_{t-1}$, $\Delta x_t = x_t - \hat{x}_t$. It can then be found that $x_{t-1} = \tilde{Q}^{-1} \tilde{q}$, where:

$$\begin{aligned} \tilde{Q} &= P_{t-1}^{-1} + F_{t-1}^T Q_t^{-1} F_{t-1} \\ \tilde{q} &= -P_{t-1}^{-1} \hat{x}_{t-1} - F_{t-1}^T Q_t^{-1} (x_t - \hat{x}_t) - F_{t-1}^T Q_t^{-1} F_{t-1} \hat{x}_{t-1} \end{aligned} \quad (29)$$

by using the Woodbury lemma [21], It can be shown that:

$$\begin{aligned} \max_{x_t} \quad & \mu \frac{1}{2} (x_t - \bar{x}_t)^T V_t^{xx}(x_t - \bar{x}_t) + \mu(x_t - \bar{x}_t)^T v_t^x \\ & - \frac{1}{2} (\Delta y - H_t \Delta x_t)^T R_t^{-1} (\Delta y - H_t \Delta x_t) - \frac{1}{2} \Delta x_t^T \bar{P}_t^{-1} \Delta x_t \end{aligned}$$

where \bar{P}_t is defined as in (4). Finally, using $P_t = (H_t^T R_t^{-1} H_t + \bar{P}_t^{-1})^{-1} = (I - K_t H_t) \bar{P}_t$, we can show that:

$$\begin{aligned} \max_{x_t} \quad & \mu \frac{1}{2} (x_t - \bar{x}_t)^T V_t^{xx}(x_t - \bar{x}_t) + \mu(x_t - \bar{x}_t)^T v_t^x \\ & - \frac{1}{2} (x_t - \hat{x}_t - \hat{\mu}_t)^T P_t^{-1} (x_t - \hat{x}_t - \hat{\mu}_t) \end{aligned} \quad (30)$$

where $P_t, \hat{\mu}_t$ are defined as in (6), (7).

REFERENCES

- [1] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.
- [2] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, "Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot," *Autonomous robots*, vol. 40, pp. 429–455, 2016.
- [3] M. Neunert, M. Stäuble, M. Gifftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli, "Whole-body nonlinear model predictive control through contacts for quadrupeds," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1458–1465, 2018.
- [4] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter, "A unified mpc framework for whole-body dynamic locomotion and manipulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4688–4695, 2021.
- [5] D. Kim, J. Di Carlo, B. Katz, G. Bledt, and S. Kim, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv preprint arXiv:1909.06586*, 2019.
- [6] E. Dantec, M. Naveau, P. Fernbach, N. Villa, G. Saurel, O. Stasse, M. Taïx, and N. Mansard, "Whole-body model predictive control for biped locomotion on a torque-controlled humanoid robot," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*. IEEE, 2022, pp. 638–644.
- [7] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [8] T. Erez, K. Lowrey, Y. Tassa, V. Kumar, S. Kolev, and E. Todorov, "An integrated system for real-time model predictive control of humanoid robots," in *2013 13th IEEE-RAS International conference on humanoid robots (Humanoids)*. IEEE, 2013, pp. 292–299.
- [9] J. Koenemann, A. Del Prete, Y. Tassa, E. Todorov, O. Stasse, M. Bennewitz, and N. Mansard, "Whole-body model-predictive control applied to the hrp-2 humanoid," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 3346–3351.
- [10] R. Budhiraja, J. Carpentier, C. Mastalli, and N. Mansard, "Differential dynamic programming for multi-phase rigid contact dynamics," in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2018, pp. 1–9.
- [11] S. Kleff, A. Meduri, R. Budhiraja, N. Mansard, and L. Righetti, "High-frequency nonlinear model predictive control of a manipulator," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7330–7336.
- [12] G. A. Einicke and L. B. White, "Robust extended kalman filtering," *IEEE Transactions on Signal Processing*, vol. 47, no. 9, pp. 2596–2599, 1999.
- [13] N. A. Villa and P.-B. Wieber, "Model predictive control of biped walking with bounded uncertainties," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. IEEE, 2017, pp. 836–841.
- [14] A. Gazar, M. Khadiv, A. Del Prete, and L. Righetti, "Stochastic and robust mpc for bipedal locomotion: A comparative study on robustness and performance," in *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2020, pp. 61–68.
- [15] A. Jordana, B. Hammoud, J. Carpentier, and L. Righetti, "Stagewise newton method for dynamic game control with imperfect state observation," *IEEE Control Systems Letters*, vol. 6, pp. 3241–3246, 2022.
- [16] F. Grimminger, A. Meduri, M. Khadiv, J. Viereck, M. Wüthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols *et al.*, "An open torque-controlled modular robot architecture for legged locomotion research," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3650–3657, 2020.
- [17] P. Whittle, "Risk-sensitive linear/quadratic/gaussian control," *Advances in Applied Probability*, vol. 13, no. 4, pp. 764–777, 1981.
- [18] M. C. Campi and M. R. James, "Nonlinear discrete-time risk-sensitive optimal control," *International Journal of Robust and Nonlinear Control*, vol. 6, no. 1, pp. 1–19, 1996.
- [19] D. A. Copp and J. P. Hespanha, "Simultaneous nonlinear model predictive control and state estimation," *Automatica*, vol. 77, pp. 143–154, 2017.
- [20] J. Morimoto and C. Atkeson, "Minimax differential dynamic programming: An application to robust biped walking," *Advances in neural information processing systems*, vol. 15, 2002.
- [21] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.
- [22] B. M. Bell and F. W. Cathey, "The iterated kalman filter update as a gauss-newton method," *IEEE Transactions on Automatic Control*, vol. 38, no. 2, pp. 294–297, 1993.
- [23] D. Murray and S. Yakowitz, "Differential dynamic programming and newton's method for discrete optimal control problems," *Journal of Optimization Theory and Applications*, vol. 43, no. 3, pp. 395–414, 1984.
- [24] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2536–2542.
- [25] H. U. Voss, J. Timmer, and J. Kurths, "Nonlinear dynamical system identification from uncertain and indirect measurements," *International Journal of Bifurcation and Chaos*, vol. 14, no. 06, pp. 1905–1933, 2004.
- [26] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiraux, O. Stasse, and N. Mansard, "The pinocchio c++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *2019 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2019, pp. 614–619.
- [27] N. Rotella, A. Herzog, S. Schaal, and L. Righetti, "Humanoid momentum estimation using sensed contact wrenches," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 556–563.
- [28] J. Carpentier, F. Valenza, N. Mansard *et al.*, "Pinocchio: fast forward and inverse dynamics for poly-articulated systems," <https://stack-of-tasks.github.io/pinocchio>, 2015–2021.
- [29] A. Meduri, P. Shah, J. Viereck, M. Khadiv, I. Havoutis, and L. Righetti, "Biconmp: A nonlinear model predictive control framework for whole body motion planning," *IEEE Transactions on Robotics*, 2023.
- [30] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, and L. Righetti, "Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid," *Autonomous Robots*, vol. 40, pp. 473–491, 2016.
- [31] A. Bambade, S. El-Kazdadi, A. Taylor, and J. Carpentier, "Prox-qp: Yet another quadratic programming solver for robotics and beyond," in *RSS 2022-Robotics: Science and Systems*, 2022.