

Document Version

Final published version

Citation (APA)

Bertipaglia, A., Shyrokau, B., Alirezaei, M., & Happee, R. (2022). A Two-Stage Bayesian optimisation for Automatic Tuning of an Unscented Kalman Filter for Vehicle Sideslip Angle Estimation. In *Proceedings of the 2022 IEEE Intelligent Vehicles Symposium (IV)* (pp. 670-677). (IEEE Intelligent Vehicles Symposium, Proceedings; Vol. 2022-June). IEEE. <https://doi.org/10.1109/IV51971.2022.9826998>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

A Two-Stage Bayesian Optimisation for Automatic Tuning of an Unscented Kalman Filter for Vehicle Sideslip Angle Estimation*

Alberto Bertipaglia¹, Barys Shyrokau¹, Mohsen Alirezaei² and Riender Happee¹

Abstract—This paper presents a novel methodology to auto-tune an Unscented Kalman Filter (UKF). It involves using a Two-Stage Bayesian Optimisation (TSBO), based on a t-Student Process to optimise the process noise parameters of a UKF for vehicle sideslip angle estimation. Our method minimises performance metrics, given by the average sum of the states' and measurement' estimation error for various vehicle manoeuvres covering a wide range of vehicle behaviour. The predefined cost function is minimised through a TSBO which aims to find a location in the feasible region that maximises the probability of improving the current best solution. Results on an experimental dataset show the capability to tune the UKF in 79.9 % less time than using a genetic algorithm (GA) and the overall capacity to improve the estimation performance in an experimental test dataset of 9.9 % to the current state-of-the-art GA.

I. INTRODUCTION

An accurate and real-time estimation of the vehicle sideslip angle is essential to strengthen the performance of active vehicle control systems. The state-of-the-art estimation of sideslip angle mostly relies on model-based approaches using an Unscented Kalman Filter (UKF). Despite its efficiency and robustness, the process and observation noise parameters must be accurately tuned to achieve good performance [1], [2]. The noise parameters selection is particularly relevant because they need to capture the following aspects:

- Process and observation model mismatch with respect to the real vehicle dynamics.
- Discretisation error because the time-step in the Kalman filter influences the noise parameters to achieve the best state estimation [3].
- Different working conditions of the sensors installed in the vehicle.

If the tuning is obtained through numerical optimisation, the cost function is highly non-convex, non-smooth and nonlinear, calling for robust optimisation methods [1].

Several optimisation algorithms for UKF tuning [4] have been studied: GA, Sequential Quadratic Programming, Nelder-Mead, Artificial Bee Colony, Fruit Fly Optimisation (FFO) and Differential Evolution. The comparison demonstrates that each algorithm leads to different optimum locations, illustrating the importance of the

*The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484).

¹Alberto Bertipaglia, Barys Shyrokau and Riender Happee are with the Department of Cognitive Robotics, Delft University of Technology, 2628 CD Delft, The Netherlands A.Bertipaglia@tudelft.nl, B.Shyrokau@tudelft.nl and R.Happee@tudelft.nl

²Mohsen Alirezaei is with the Department of Mechanical Engineering, University of Eindhoven, 5612 AZ Eindhoven, The Netherlands m.alirezaei@tue.nl




Fig. 1. Framework utilizing a Two-Stage Bayesian Optimisation (TSBO) to tune the process noise parameters of a UKF.

optimization method. However, a comparison regarding the amount of data and the optimisation time has been omitted. Furthermore, only stochastic optimisation algorithms have been compared that can be inefficient if the UKF takes a long simulation time. For vehicle sideslip angle estimation, large-scale training sets representing vehicle behaviour [5] make the cost function evaluation very time-consuming. In different engineering applications, several solutions have been proposed, e.g. Reinforcement Learning (RL) [6] or Bayesian Optimisation (BO) [1]. However, the majority of these are tested only on simulation data or toy examples. A more detailed analysis of previous works is in Section II.

This paper proposes a new UKF tuning methodology using Two-Stage Bayesian Optimisation (TSBO) [7], based on a t-Student Process (tSP) for vehicle sideslip angle estimation. The proposed methodology to tune the process noise parameters reduces the optimisation time and improves the optimum localisation. Both process and observation noise are assumed Gaussian, zero mean and uncorrelated, as for the majority of the model-based vehicle state estimation filters [8]. The methodology framework is represented in Figure 1, a more detailed explanation will be presented in Section IV. The observation noise parameters are tuned by performing a statistical analysis of the vehicle sensor measurements [5]. The process noise parameters' tuning is based on optimising a cost function defined as the sum of the Normalized Root Mean Squared Errors (NRMSE) of the vehicle states and measurements for 8 manoeuvres. The 8 manoeuvres which compose the training set and the 23 forming the test set are selected

from a real-world experimental dataset composed of 216 manoeuvres recorded at the Automotive Testing Papenburg GmbH. The manoeuvres are chosen to cover a wide range of vehicle motions. The cost function is optimised using the modified TSBO [7]. It consists of a fast exploration and a pure exploitation stage to find the optimum, reducing the number of simulations required. Differently from the TSBO proposed in [7], the surrogate model is a tSP to enhance the robustness and the ability to capture heavy-tailed NRMSE. The methodology is experimentally validated using a UKF, but can also be applied to an Extended Kalman Filter (EKF).

II. PREVIOUS WORKS

Despite the extensive use of UKF for vehicle sideslip angle estimation, the tuning of the process and observation noise parameters is rarely investigated in the literature. The observation noise parameters are usually tuned through statistical analysis, obtained after sensor calibration tests [5]. The only exception is when the observation noise parameters are tuned together with the process noise parameters [9]. This rarely happens because the filter performance depends mainly on the ratio between the eigenvalues of the observation and process covariance matrices. Thus, it is more convenient to fix the observation noise parameters, which can be determined through sensor calibration, and optimise the process noise parameters. The methodologies for process noise parameters tuning can be split into three different categories: *manual tuning*, *metaheuristic optimisation* and *data-driven* techniques. An overview of optimisation algorithms for UKF tuning is shown in Table I.

Manual tuning is a trial-and-error approach driven by the user experience. The idea is to perform a grid search in the parameter space and reach the best tuning iteratively. Despite the oversimplicity of the concept, it is still very often used, especially for vehicle sideslip angle estimation [10], because the performance of Kalman filters is stable for a range of process noise parameters settings [5]. The main drawback is the inefficiency of the approach. Moreover, there is no proof to reach the optimal performance. This leads to an unreliable comparison when different filter architectures are analysed.

Thus, Kalman filter tuning with numerical optimisation techniques are widely used. Gradient-based optimisations are the most straightforward algorithms which can be implemented [11], but the non-convex and nonlinear cost functions let them easily be trapped into local minima. Thus, gradient-free optimisation algorithms are introduced. One of the simplest algorithms is the coordinate descent algorithm [2] which consists of a cycling increase or decrease of each parameter of the process covariance. A similar optimisation technique is the downhill simplex algorithm [12], which consists of evaluating the cost function at the simplex vertices of the predefined sample space. After this, the algorithm should converge to the optimum through a series of movements: reflection, expansion, and contraction. Interested reader can find more information on both algorithms in [2] and [12] respectively. These algorithms may easily fall in local optima

TABLE I
OVERVIEW OF THE OPTIMISATION ALGORITHM MOST USED TO TUNE
THE NOISE PARAMETERS OF NONLINEAR KALMAN FILTER.

Authors	Process noise tuning	Dataset
van Aalst, 2018 [10]	Manual Tuning	Experimental
Abbeel, 2005 [2]	Coordinate Ascent Algorithm	Experimental
Powell, 2002 [12]	Downhill Simplex Algorithm	Simulated
Mazzilli, 2021 [5], Oshman, 2000 [13]	Multi-Objective GA, GA	Experimental
Acosta, 2019 [4]	FFO	Simulated
Heidfeld, 2021 [9]	Simulated Annealing	Experimental
Tang, 2021 [6]	RL	Experimental
Escoriza, 2021 [14]	Recurrent Neural Network	Experimental
Chen, 2018 [1], Chen, 2019 [15]	BO based on Gaussian & t-Student Process	Simulated
Current Paper	TSBO with t-Student Process	Experimental

when the cost function is non-smooth, so more advanced metaheuristic optimisations have been proposed.

Metaheuristics are procedures that can provide an acceptable solution to an optimisation problem with incomplete information about the cost function. Examples of metaheuristic algorithms are: the already cited downhill simplex algorithm [12], GA [13] or the Multi-Objective GA [5], SA [9] and FFO [4]. A comparison of different metaheuristic algorithms for Kalman filter tuning [4] shows that GA and FFO found the best cost function optimum. This explains why GA is the current state-of-the-art for UKF tuning for vehicle sideslip angle estimation, which is applied in many recent publications [5]. GA is an evolutionary algorithm that tries to emulate Darwin's theory of natural evolution. Due to the high number of generations required, GA can be very time-consuming when the cost function is costly to evaluate.

Data-driven optimisation algorithms such as BO are based on the creation of a surrogate model, Gaussian Processes GP [1], or tSP [15], using as inputs the tuning parameters and as output the value of the cost function. The surrogate model does not only approximate the cost function, but also it is associated with a model probability distribution. The acquisition function takes this information to find the new parameters with the highest probability of being the best new optimum. Standard BO fits the surrogate model on parameters randomly taken from the sample space, but this can lead to a non-optimal surrogate model. Furthermore, it will require many cost function evaluations if the sample space is broad. Other data-driven approaches can be implemented; for instance, Reinforcement Learning RL can be applied to choose the parameters of the process noise, or RL can compensate for the error of the Kalman filter [6] due to not-optimal process noise parameters. A similar concept is the so-called "KalmanNet" [14] which uses a Recurrent Neural Network RNN to compute the Kalman gain, reducing the problems of non-optimal Kalman tuning. Despite the potential of RL and RNN, the amount of data required to train the Neural Net-

works makes them impractical for the tuning/improvement of Kalman filter for sideslip angle estimation.

The main contributions of this paper are twofold. The first is the development of a new TSBO based on a tSP for tuning the process noise parameters of UKF for vehicle sideslip angle estimation. The TSBO in this paper differs from a recent TBSO related to electronic design [7] because the surrogate model is based on a tSP to increase its robustness and ability to reach a better global optimum without increasing the required computational power.

The second contribution is that the developed TSBO reduces the number of simulations required to tune the Kalman filter with respect to the state-of-the-art GA [4], [5] and it improves the accuracy of the tuning respect GA. The tuning performance is tested and validated using an experimental dataset for vehicle sideslip angle estimation.

III. BAYESIAN OPTIMISATION

Filter tuning consists of minimising a nonlinear, non-convex and non-smooth cost function due to the reasons already expressed in section I. A minimisation problem is:

$$\begin{aligned} \min_{q \in Q} J(q) \\ \text{where } Q \subseteq \mathcal{R}^d \end{aligned} \quad (1)$$

where q is the parameter vector of dimension d , $J(q)$ is the cost function, and $Q \subseteq \mathcal{R}^d$ is the solution space. The function $J(q)$ is a "black-box" function because it is not accessible, but its outputs are observed based on some given inputs. Thus, "black-box" stochastic optimisation algorithms are considered. It is time-consuming to perform a dense sampling of the solution space Q , so the algorithm should focus on sampling the subset of Q with the highest probability to contain the cost function optimum q^* . BO can deal with incomplete and sparse knowledge of the solution thanks to its probabilistic approach. It is based on the Bayes' theorem:

$$P(\text{beliefs} | \text{data}) = \frac{P(\text{beliefs}) \times P(\text{data} | \text{beliefs})}{P(\text{data})} \quad (2)$$

where data are the available observations of the function $J(q)$, beliefs are the beliefs of the shape of $J(q)$, and $P(\text{beliefs} | \text{data})$, $P(\text{data} | \text{beliefs})$, $P(\text{beliefs})$, and $P(\text{data})$ are the posterior, likelihood, and the marginal probabilities respectively. Thus, BO aims to find q^* learning the shape of $J(q)$ through Bayesian inference. The TSBO overcomes the standard BO due to its ability to define a better ratio between exploration (function shape learning) and exploitation (approaching the optimum).

BO can be split into two steps: creating a stochastic process of the cost function, called "Surrogate Model", and building an "Acquisition Function" that uses the surrogate model to approach the optimum q^* .

A. Surrogate Model

BO approximates the "black-box" cost function through a stochastic process called the surrogate model. It represents a probabilistic prior over the space of functions. The prior is




Fig. 2. GP and tSP model of the cost function varying the process noise parameter related with the vehicle lateral velocity. The tSP allows more outliers and it can follow the measurements better than the GP. The shadow area is the variance of the tSP and the uncertainty for the GP.

updated as soon as the cost function is evaluated thanks to the Bayes' Theorem. Different priors are eligible for being the surrogate model; a partial list comprises: GP, tSP, Bayesian neural networks and polynomial chaos expansion. GP are the current state-of-the-art because it is a non-parametric model [7], and the Bayesian update step is analytical. Despite this, in this paper a tSP prior is selected for the surrogate model due to two reasons. The first one is that the t-Student distribution allows defining the level of Kurtosis, so it allows much more likely outliers [16], [17] than GP. The second reason is that the observations directly influence the predictive covariance [16], [17]. Both properties are beneficial for the tuning of UKF because the cost function can have discontinuities, for instance, due to the noise. Figure 2 shows the explained differences between tSP and GP modelling in an example that models the cost function varying the process noise parameter related with lateral velocity. It is visible how tSP reduces the influence of the outliers, allowing the mean to have a better fitting.

The tSP is defined as:

$$f(q) = tSP(m(q), k(q, q'), v) \quad (3)$$

where $m(q)$ represents the mean which is assumed equal to constant zero mean function because there is no prior [18], v are the degrees of freedom of the t-Student distribution, which defines the level of Kurtosis, and $k(q, q')$ is the Kernel function between the inputs pair q and q' . The choice of the Kernel is essential to have a reliable model fitting. This paper chooses the automatic relevance determination Matérn 5/2 function, due to its generalisation capabilities and interpretability properties:

$$k_{q, q'} = k(q, q') = \sigma_f^2 \left(1 + \frac{\sqrt{5}r}{\sigma_d} + \frac{5r^2}{3\sigma_d^2} \right) e^{-\frac{\sqrt{5}r}{\sigma_d}} \quad (4)$$

where r is the Euclidean distance between the two points q and q' calculated as:

$$r = \sqrt{(q - q')^T (q - q')} \quad (5)$$

where σ_f and σ_d are two hyperparameters, which are trained to minimize the negative log marginal likelihood of the tSP. It is essential to highlight that tSP should be accurate only when there are higher chances of finding the cost function optimum. Thus, every time a new observation is available, the surrogate model is retrained. tSP are implemented using the GPML Matlab Code version 4.2 [19].

B. Acquisition Function

The Acquisition Function AF is responsible for moving the BO towards the optimum region, choosing the next sample point to be evaluated. Thus, it is responsible for defining the ratio between exploration, moving the search towards the area where the surrogate model is uncertain, and exploiting, moving the search towards the area with a higher chance of improving the current optimum. Ideally, the more focus is given to the exploration stage, the higher the chances to find the global optimum; however, it increases the computational time due to a larger number of cost function evaluations.

The next point to be evaluated (q_{t+1}) is usually obtained through an additional optimisation of the AF. This paper evaluates the AF on some candidate points obtained through the process that will be explained in Section IV to avoid further optimisation. The AF is evaluated on the candidate points [17], [19] according to:

$$\hat{\mu}_{tSP}(q_{t+1}) = k_{q,\tilde{q}}^T K_{\tilde{q},\tilde{q}}^{-1} f(\tilde{q}) \quad (6)$$

where $\tilde{q}$ are the sampling points already evaluated, $\hat{\mu}_{tSP}$ is the posterior mean of a tSP for a new sampling point q and k is defined in Eq. 4.

$$\hat{\sigma}_{tSP}^2(q_{t+1}) = \left(\frac{\mathbf{v} + f(\tilde{q})^T K_{\tilde{q},\tilde{q}}^{-1} J(\tilde{q}) - 2}{\mathbf{v} + |D| - 2} \right) \times (k(q, q) + k_{q,\tilde{q}}^T K_{\tilde{q},\tilde{q}}^{-1} k_{q,\tilde{q}}) \quad (7)$$

where $\hat{\sigma}_{tSP}$ is the posterior covariance, and D is the set of samples $D = [(\tilde{q}_1, J(\tilde{q}_1)), (\tilde{q}_2, J(\tilde{q}_2)), (\tilde{q}_3, J(\tilde{q}_3)), \dots]$.

Two AFs are selected: the Expected Improvement (EI) and the Confidence Bound Minimisation (CBM) [20].

The EI is defined as:

$$EI_{tSP}(q) = \hat{\sigma}_{tSP}(q) \left(\frac{\mathbf{v}}{\mathbf{v} - 1} \right) \left(1 + \frac{z_s^2}{\mathbf{v}} \right) \phi(z_s) + [\hat{y} - \hat{\mu}_{tSP}] \Phi(z_s) \quad (8)$$

where $\phi(z_s)$ and $\Phi(z_s)$ are the probability density function and the cumulative distribution function of an univariate standard t-Student's random variable, z_s , and $\hat{y}$ is the current optimum sampled by the BO.

The CBM is formulated as:

$$CBM_{tSP}(q) = \hat{\sigma}_{tSP}(q) \sqrt{\beta} + [f^* - \hat{\mu}_{tSP}] \quad (9)$$

where the β defines the ratio between exploitation and exploration, and f^* is the best known optimum, representing the prior knowledge of the optimum.

IV. TWO-STAGE BAYESIAN OPTIMISATION

This paper proposes an evolution of the BO to tune the process noise parameters of the UKF. It is called TSBO, code by [7], and aims to reduce the number of function evaluations, and improve the ratio between exploitation and exploration. The trade-off improvement is obtained thanks to the subdivisions of the BO in two stages. The first, "fast exploration", aims to find the sample space region where the optimum is contained. The second, "pure exploitation", aims to fine-tune the restricted sample space obtained from the first stage. The first stage allows starting the surrogate model training from a single evaluation at the centre of the sample space. This allows the user not to evaluate a random set of process noises to fit the first surrogate model.

A. Fast Exploration

The fast exploration aims to narrow the sample space $Q \subseteq \mathcal{R}^d$ where the BO search the optimum as fast as possible. It consists of subdividing the sample space into 2^d hyper-rectangles. Every hyper-rectangle centre point is a candidate point, which the AF evaluates to choose the next sampling point q_{t+1} where the cost function is evaluated. The following sampling point is not obtained as in the standard BO through an auxiliary optimisation of the AF, but it is the candidate point with the maximum value of the AF between the other candidate points. The cost function is only evaluated at points q_{t+1} . A new set of hyper-rectangles is generated starting from the hyper-rectangle enclosing the best candidate point. The two available AFs, EI and CBM, are alternatively used to evaluate the candidate points. After a predefined number of iterations MAX_{AF} , the AF which would have had the greatest gain to the optimisation is selected as final AF.

The fast exploration stage is repeated until the euclidean distance between the new best sampling point $q_{t+1,max}$ and the previous best sampling point $q_{t,max}$ is below a user-defined threshold TR_{FE} for a number of times MAX_{FE} :

$$n = \begin{cases} n+1, & \text{if } \|q_{t+1,max} - q_{t,max}\| < TR_{FE} \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

Every time the cost function is evaluated and a new measurement is available; the surrogate model is retrained to improve the fitting with the black-box function.

Algorithm 1 sums up the fast exploration stage of the TSBO.

B. Pure Exploitation

The pure exploitation stage follows a procedure similar to the fast exploration, but it specialises in refining the optimum position. The first step is the computation of the sample space $\tilde{Q}^d$ around the optimum computed in the fast exploration stage:

$$\tilde{Q}^d = \begin{bmatrix} (1-\alpha) q_1^* & (1+\alpha) q_1^* \\ \vdots & \vdots \\ (1-\alpha) q_d^* & (1+\alpha) q_d^* \end{bmatrix} \quad (11)$$

Algorithm 1 Fast Exploration stage of the TSBO

Require: A sample space Q and an evaluated sampling point at the center of Q
Fit the tSP Surrogate model
Subdivide Q in 2^d hyper-rectangles
while $n \leq MAX_{FE}$ **do**
 Compute candidate points at the center of hyper-rectangles
 Evaluate AF at the candidate points
 Pick the candidate points with the maximum AF
 if $\|q_{t+1max} - q_{tmax}\| < TR_{FE}$ **then**
 $n + 1 \leftarrow n$
 else if $\|q_{t+1max} - q_{tmax}\| \geq TR_{FE}$ **then**
 $n \leftarrow n$
 end if
 Evaluate the cost function in the sampling points
 Train the tSP of the Surrogate with the new observation
 Subdivide the hyper-rectangle which were enclosing the best candidate points in 2^d hyper-rectangles
end while

Algorithm 2 Pure Exploitation stage of the TSBO

Require: $n = MAX_{FE}$
while $n_{iter} \leq MAX_{PE}$, and the output from Algorithm 1 **do**
 Compute candidate points at the center of hyper-rectangles
 Evaluate AF at the candidate points
 Pick the candidate points with the maximum AF
 Evaluate the cost function in the sampling points
 if $n_{iter} \leq MAX_{SM}$ **then**
 Train the tSP of the Surrogate with the new observation
 end if
 Subdivide the hyper-rectangle which were enclosing the best candidate points in 3^d hyper-rectangles
end while

where α is a hyper-parameter that defines how wide the sample space should be, and q_i^* is the optimum sampling point obtained in the fast exploration stage. The new sample space $\tilde{Q}^d$ is divided into 3^d hyper-rectangles to have more candidate points to evaluate. The higher number of candidate points helps the pure exploitation stage to converge faster to the global optimum of the function. The rest of the procedure is the same of the fast exploration. The only difference is that when the number of evaluated sampling points n_{iter} is higher than a predefined threshold MAX_{SM} , the tSP is no more retrained.

Algorithm 2 summarises the pure exploitation stage of the TSBO, which is iterated until n_{iter} is lower than a user-defined threshold MAX_{PE} .

V. APPLICATION: UKF TUNING

The proposed TSBO based on tSP is tested on the tuning of the process noise parameters of a UKF for vehicle sideslip angle estimation. In this paper, the single-track model with

tyre axle forces computed by the Dugoff tyre is chosen as the vehicle model. The states are the longitudinal velocity at the Center-of-Gravity CoG V_x , the lateral velocity at the CoG V_y and the yaw rate $\dot{\psi}$ while the vehicle measurements are the V_x , the lateral acceleration at the CoG a_y and the $\dot{\psi}$. The vehicle sideslip angle β_s is obtained as $\beta_s = \arctan \frac{V_y}{V_x}$:

A. Dataset

All the manoeuvres that compose the training set and test set are taken from a real-world experimental dataset composed of 216 manoeuvres which correspond to 2 hrs of driving. It puts together standard vehicle dynamics manoeuvres, e.g. double lane change, slalom, random steer, J-turn, spiral, braking in a turn, and steady-state circular tests, together with recorded laps at the Papenburg track. All manoeuvres were performed on dry asphalt with tyres inflated according to the vehicle specifications.

All manoeuvres are recorded using the BMW test platform instrumented by the inertial measurement unit, wheel force transducers for all four wheels, GPS and an optical sensor from Corrsys-Datron to measure the sideslip angle. The high-end optical speed sensor is not present in consumer vehicles and will be used as a reference for TBSO tuning and validation.

The UKF should be tuned for different working conditions, so various vehicle maneuvers are jointly considered in the training set. The training set is composed of 8 manoeuvres: 1 braking in a turn, 1 skidpad, 2 J-turn, 2 slaloms and 2 lane change.

A test set composed of 23 manoeuvres is taken to prove the generalization capability of the approach. All 23 maneuvers of the test set are different from the training set even if, some of them have similar characteristics. The manoeuvres are: 2 braking in a turn, 2 skidpad, 5 J-turn, 4 slalom, 4 lane change, 2 random steer, 1 lap track and 3 spiral. The characteristics of the manoeuvres cover a wide range of vehicle driving behaviours and different settings of the Electronic Stability Control.

B. Cost Function Formulation

In this paper, the UKF tuning aims to improve the accuracy of the vehicle state estimation, therefore, this paper focuses on an objective function defined by minimising the Normalised Root Mean Squared Error NRMSE rather than maximising the negative log-likelihood [2].

The minimisation problem is described as:

$$\min_{q \in Q} J(q) \quad (12)$$
$$\text{where } Q = \begin{bmatrix} q_{1,min} & q_{1,max} \\ \vdots & \vdots \\ q_{d,min} & q_{d,max} \end{bmatrix}$$

where J is the cost function defined as:

$$J = W_1 E_\beta + W_2 E_{\dot{\psi}} + W_3 E_{a_y} \quad (13)$$
$$\text{where } E_j = \sqrt{\frac{\sum_{i=1}^{N_{man}} (NRMSE_{j,i})^2 N_{s,i}}{\sum_{i=1}^{N_{man}} N_{s,i}}}$$

where W_1, W_2, W_3 are the weights of the cost function, N_s is the length of the manoeuvre measured in sampling points, N_{man} is the number of manoeuvres considered in the training set, and $NRMSE$ is computed as follow:.

$$NRMSE = \sqrt{\frac{\sum_{k=1}^{N_s} (\hat{X}_k - X_k)^2}{N_s}} \frac{1}{\max(|X|)} \quad (14)$$

where X and $\hat{X}$ are the vectors of the measured and estimated states respectively.

C. Key Performance Indicator

The accuracy of the estimation of sideslip angle is evaluated through the use of the root mean squared error RMSE, maximum absolute error MAE, the RMSE and the MAE when the absolute vehicle acceleration a_y is greater than 4 m/s^2 , which are respectively abbreviated with $RMSE_{NON}$ and MAE_{NON} . The latter is used to analyse how the UKF tuning is especially relevant when the vehicle behaves in a nonlinear way.

VI. RESULTS

The performance in terms of minimisation time and optimum localisation of the TSBO based on tSP is compared with the current state-of-the-art technique GA and with a TSBO based on GP. All optimisations perform in the same sample space Q where all the process noise parameters for V_x, V_y and ψ vary from 10^{-10} to 1.

The optimisation parameters for TSBO based on tSP and GP are summarised in Table II. The parameters for GA are summarized in Table III.

The optimum cost and optimisation time for the TSBO based on tSP, TSBO based on GP and GA are summarised in Table IV and Figure 3.

Figure 3 shows how both TSBOs approach the global optimum faster than GA. Where TSBO converges after 40 evaluations of the cost function, the GA needs at least 220 evaluations, resulting in TSBO being on average 79 % faster than GA. Thus, TSBOs are particularly convenient when the cost function evaluation is computationally heavy,

TABLE II
USER-DEFINED PARAMETERS FOR TSBO OPTIMISATIONS.

Parameters	TSBO - tSP	TSBO - GP
Likelihood function	t-Student's	Gaussian
Inference	Variational Bayes	Exact
Kernel function	Matérn ARD 5/2	Matérn ARD 5/2
Mean function	Constant	Constant
v	15	NA
MAX_{FE}	15	15
MAX_{PE}	40	40
MAX_{SM}	38	38
TR_{FE}	$0.01 \parallel q^* \parallel$	$0.01 \parallel q^* \parallel$
β	0.01	0.01
α	0.15	0.15
W_1, W_2, W_3	5, 1, 1	5, 1, 1

TABLE III
USER-DEFINED PARAMETERS FOR GA OPTIMISATION.

Parameters	GA
Population size	15
Max generations	15
Elite count	0.75
Crossover fraction	0.8
W_1, W_2, W_3	5, 1, 1

TABLE IV
NUMERICAL RESULTS OF THE OPTIMUM COST FUNCTION AND OPTIMISATION TIME USING DIFFERENT ALGORITHMS FOR UKF TUNING.

Algorithm	Total cost [-]	Optim. time [s]
Manual tuning	1.714	NA
GA	0.378	1599
TSBO - GP	0.400	306
TSBO - tSP	0.316	320

for instance, due to numerous manoeuvres considered in the training set. Despite the faster optimisation, TSBO based on tSP reaches an optimum 16.4 % lower than the GA optimum. The proposed algorithm outperforms the state-of-the-art approach both in terms of optimisation time and optimisation accuracy. TSBO based on GP is 4.4 % faster than TSBO based on tSP. The rise in optimisation time is due to the surrogate model's training phase, which is slower for the tSP due to the more optimisation parameters. The better optimisation accuracy, 21 %, of the tSP than the GP is due to the higher Kurtosis of the t-Student's distribution which makes the non-parametric model capable of dealing with outliers. A much higher optimisation accuracy compensates for the slight higher computational time.

Table V presents the RMSE, MAE, $RMSE_{NON}$ and MAE_{NON} of the sideslip angle for the various optimised UKFs. The reported values are the average of all manoeuvres composing the training set. Table V highlights how the cost function's minimisation is related to the KPIs used




Fig. 3. Performance comparison of TSBO based on tSP, TSBO based on GP and GA.

TABLE V

COMPARISON OF THE AVERAGE KPIS FOR THE UKF TUNED BY MANUAL TUNING, GA, TSBO BASED ON GP AND TSBO BASED ON STP. RESULTS OF THE TRAINING SET.

KPI	Manual	GA	TSBO - GP	TSBO - tSP
RMSE [deg]	1.028	0.500	0.634	0.423
MAE [deg]	1.196	0.900	1.252	0.765
RMSE _{NON} [deg]	0.716	0.568	0.855	0.428
MAE _{NON} [deg]	0.818	0.671	0.961	0.587

to compare the estimation accuracy because TSBO based on tSP outperforms the other algorithms. On average, the RMSE of the sideslip angle optimised with tSP based on tSP is 15.4 % lower than the one with GA, and the MAE is 15 % lower. Particularly important is the improvement of the RMSE_{NON} and of the MAE_{NON} because when the $|a_y| \geq 4 \text{ m/s}^2$ the sideslip angle estimation becomes more critical for vehicle stability control. The UKF optimised by the TSBO based on tSP reduces the RMSE_{NON}, and the MAE_{NON} of the UKF optimised by GA of 24.6 % and 12.5 %, respectively. The accuracy of the UKF tuned by GA is higher than the one optimised by TSBO based on GP.

Table VI reports the same information as Table V, but for the test set, to check that the numerical optimisation algorithms do not overfit the training set leading to worse performance in other manoeuvres. Table VI shows that TSBO based on tSP still behaves better than all other algorithms for the 4 KPIS, proving better performance in tuning the process noise UKF parameters. On average, the RMSE of the sideslip angle optimised with tSP based on tSP is 9.9 % lower than the one with GA, and the MAE is 17.6 % lower. The improvement of the RMSE_{NON} and of the MAE_{NON} is particularly significant because when the $|a_y| \geq 4 \text{ m/s}^2$, the sideslip angle estimation becomes more critical for vehicle stability control. The UKF optimised with the TSBO based on tSP reduces the RMSE_{NON}, and the MAE_{NON} of the UKF optimised with GA of 10.6 % and 9.8 %, respectively. The improvement in the test set is lower than the one obtained in the training set but still particularly relevant, especially for what concerns the MAE.

Figure 4 shows the relative improvement of the UKF tuned by TSBO based on tSP and GA in 6 different maneuvers of the test set. The relative improvement is measured with the RMSE of the sideslip angle, a_y and ψ . In 5 out of 6 maneuvers the TSBO reaches higher performance than the GA. The spiral maneuver is the only one which shows a degraded accuracy in the RMSE of sideslip angle of 8 %. Performance degradation in this particular maneuver is considered acceptable due to the various driving conditions of the test set. The relative improvement is significant especially for the sideslip angle and for the a_y , while the




Fig. 4. Relative improvement of the UKF tuned by TSBO based on tSP with respect to the UKF tuned by the state-of-the-art GA. The improvement is measured through the sideslip angle's RMSE for 6 different maneuver chosen from the test set.




Fig. 5. Sideslip angle estimation in a skidpad maneuver.

ψ does not have this. The reason is that the ψ has lower maximum absolute error compared with the other variables, so the relative improvement is less noticeable. Figure 4 shows how TSBO can enhance the performance of the UKF for both transient (J-turn, slalom) and steady state (skidpad) maneuvers. Significant is the improvement in the skidpad, see also Figure 5. The reason is that during a steady-state maneuver the vehicle yaw acceleration is almost null so the difference between estimated and measured tyre forces becomes essential for the V_y estimation. The best way to reduce the tyre force error is to optimise the process noise parameters.

Despite this, the improvement in sideslip angle estimation is noticeable also for a J-turn, see Figure 7, and a slalom

TABLE VI

COMPARISON OF THE AVERAGE KPIS FOR THE UKF TUNED BY MANUAL TUNING, GA, TSBO BASED ON GP AND TSBO BASED ON STP. RESULTS OF THE TEST SET.

KPI	Manual	GA	TSBO - GP	TSBO - tSP
RMSE [deg]	0.577	0.383	0.431	0.345
RMSE [deg]	0.672	0.529	0.689	0.436
RMSE _{NON} [deg]	1.328	0.860	1.030	0.769
MAE _{NON} [deg]	1.047	0.746	0.902	0.673




Fig. 6. Sideslip angle estimation in a slalom maneuver.




Fig. 7. Sideslip angle estimation in a J-turn maneuver.

manoeuvre, see Figure 6. Figure 6 shows how the improvement in the tuning of the UKF is noticeable at the sideslip angle peaks of the slalom maneuver. The improvement at the peaks is more remarkable when the sideslip angle is above 3 deg while when it is between 0 and 3 deg the difference is not visible due to the lower absolute error of both UKFs. Figure 7 shows the sideslip angle estimation in J-turn and it highlights another time how the UKF tuned by TSBO based on tSP reduces the estimation error especially when the sideslip angle is above 4 deg.

VII. CONCLUSIONS

This paper presented a two-stage Bayesian optimisation approach based on t-Student process to tune the process noise parameters of an unscented Kalman filter for sideslip angle estimation. The proposed approach reaches the cost function optimum 79.9 % faster than the state-of-the-art GA and the optimum is 16.4 % better than the one obtained by GA. The obtained results are tested on an experimental test set composed by transient and quasi-steady state maneuvers. Furthermore, two-stage Bayesian optimisation creates a physical representation of the black-box cost function thanks to surrogate model training. The performance of the two-stage Bayesian optimisation based on t-Student process is also compared with the two-stage Bayesian optimisation based on Gaussian process. The algorithm based on t-Student process is 4.4 % slower than the one based on Gaussian process but it reaches a 21 % better optimum. As future research, the plan is to tune not only the process noise parameters of the unscented Kalman filter but also the parameters defining the location of the sigma points.

REFERENCES

- [1] Z. Chen, C. Heckman, S. Julier, and N. Ahmed, "Weak in the knees?: Auto-tuning kalman filters with bayesian optimization," in *2018 21st International Conference on Information Fusion (FUSION)*. IEEE, 2018, pp. 1072–1079.
- [2] P. Abbeel, A. Coates, M. Montemerlo, A. Y. Ng, and S. Thrun, "Discriminative training of kalman filters," in *Robotics: Science and systems*, vol. 2, 2005, p. 1.
- [3] Z. Chen, C. Heckman, S. Julier, and N. Ahmed, "Time dependence in kalman filter tuning," in *2021 IEEE 24th International Conference on Information Fusion (FUSION)*. IEEE, 2021, pp. 1–8.
- [4] M. Acosta and S. Kanarachos, "Optimized vehicle dynamics virtual sensing using metaheuristic optimization and unscented kalman filter," in *Evolutionary and Deterministic Methods for Design Optimization and Control With Applications to Industrial and Societal Problems*. Springer, 2019, pp. 275–290.
- [5] V. Mazzilli, D. Ivone, S. De Pinto, L. Pascali, M. Contrino, G. Tarquinio, P. Gruber, and A. Sornioti, "On the benefit of smart tyre technology on vehicle state estimation," *Vehicle System Dynamics*, pp. 1–26, 2021.
- [6] Y. Tang, L. Hu, Q. Zhang, and W. Pan, "Reinforcement learning compensated extended kalman filter for attitude estimation," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6854–6859.
- [7] H. M. Torun, M. Swaminathan, A. K. Davis, and M. L. F. Bellaredj, "A global bayesian optimization algorithm and its application to integrated system design," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 4, pp. 792–802, 2018.
- [8] H. Heidfeld, M. Schünemann, and R. Kasper, "Ukf-based state and tire slip estimation for a 4wd electric vehicle," *Vehicle System Dynamics*, vol. 58, no. 10, pp. 1479–1496, 2020.
- [9] H. Heidfeld and M. Schünemann, "Optimization-based tuning of a hybrid ukf state estimator with tire model adaption for an all wheel drive electric vehicle," *Energies*, vol. 14, no. 5, p. 1396, 2021.
- [10] S. van Aalst, F. Naets, B. Boulkroune, W. De Nijs, and W. Desmet, "An adaptive vehicle sideslip estimator for reliable estimation in low and high excitation driving," *IFAC-PapersOnLine*, vol. 51, no. 9, pp. 243–248, 2018.
- [11] S. Kerst, B. Shyrokau, and E. Holweg, "A model-based approach for the estimation of bearing forces and moments using outer ring deformation," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 1, pp. 461–470, 2019.
- [12] T. D. Powell, "Automated tuning of an extended kalman filter using the downhill simplex algorithm," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 5, pp. 901–908, 2002.
- [13] Y. Oshman and I. Shaviv, "Optimal tuning of a kalman filter using genetic algorithms," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2000, p. 4558.
- [14] A. L. Escoriza, G. Revach, N. Shlezinger, and R. J. G. van Sloun, "Data-driven kalman-based velocity estimation for autonomous racing," in *2021 IEEE International Conference on Autonomous Systems (ICAS)*, 2021, pp. 1–5.
- [15] Z. Chen, N. Ahmed, S. Julier, and C. Heckman, "Kalman filter tuning with bayesian optimization," *arXiv preprint arXiv:1912.08601*, 2019.
- [16] A. Shah, A. Wilson, and Z. Ghahramani, "Student-t processes as alternatives to gaussian processes," in *Artificial intelligence and statistics*. PMLR, 2014, pp. 877–885.
- [17] B. D. Tracey and D. Wolpert, "Upgrading from gaussian processes to student's t processes," in *2018 AIAA Non-Deterministic Approaches Conference*, 2018, p. 1659.
- [18] H. M. Torun, J. A. Hejase, J. Tang, W. D. Beckert, and M. Swaminathan, "Bayesian active learning for uncertainty quantification of high speed channel signaling," in *2018 IEEE 27th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*. IEEE, 2018, pp. 311–313.
- [19] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [20] C. Clare, G. Hawe, and S. McClean, "Expected regret minimization for bayesian optimization with student's-t processes," in *Proceedings of the 2020 3rd International Conference on Artificial Intelligence and Pattern Recognition*, 2020, pp. 8–12.