

Document Version

Final published version

Licence

CC BY

Citation (APA)

Bertipaglia, A., Alirezaei, M., Happee, R., & Shyrokau, B. (2024). A Learning-Based Model Predictive Contouring Control for Vehicle Evasive Manoeuvres. In G. Mastinu, F. Braghin, F. Cheli, M. Corno, & S. M. Savaresi (Eds.), *16th International Symposium on Advanced Vehicle Control: Proceedings of AVEC 2024 - Society of Automotive Engineers of Japan* (pp. 632-638). (Lecture Notes in Mechanical Engineering). Springer. https://doi.org/10.1007/978-3-031-70392-8_89

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.



A Learning-Based Model Predictive Contouring Control for Vehicle Evasive Manoeuvres

Alberto Bertipaglia¹ (✉) , Mohsen Alirezaei² , Riender Happee¹,
and Barys Shyrokau¹ 

¹ Delft University of Technology, 2628 CD Delft, The Netherlands
{A.Bertipaglia,R.Happee,B.Shyrokau}@tudelft.nl

² Siemens PLM Software, 5708 JZ Helmond, The Netherlands
mohsen.alirezaei@siemens.com

Abstract. This paper presents a novel Learning-based Model Predictive Contouring Control (L-MPCC) algorithm for evasive manoeuvres at the limit of handling. The algorithm uses the Student-t Process (STP) to minimise model mismatches and uncertainties online. The proposed STP captures the mismatches between the prediction model and the measured lateral tyre forces and yaw rate. The mismatches correspond to the posterior means provided to the prediction model to improve its accuracy. Simultaneously, the posterior covariances are propagated to the vehicle lateral velocity and yaw rate along the prediction horizon. The STP posterior covariance directly depends on the variance of observed data, so its variance is more significant when the online measurements differ from the recorded ones in the training set and smaller in the opposite case. Thus, these covariances can be utilised in the L-MPCC's cost function to minimise the vehicle state uncertainties. In a high-fidelity simulation environment, we demonstrate that the proposed L-MPCC can successfully avoid obstacles, keeping the vehicle stable while driving a double lane change manoeuvre at a higher velocity than an MPCC without STP. Furthermore, the proposed controller yields a significantly lower peak sideslip angle, improving the vehicle's manoeuvrability compared to an L-MPCC with a Gaussian Process.

Keywords: Learning-based model predictive control · Student-t process · Evasive manoeuvre · Obstacle avoidance · Limit of handling

1 Introduction

A crucial safety element of automated driving is to prove the capacity to avoid obstacles at the limit of handling. A common solution is based on Nonlinear

The Dutch Science Foundation NWO-TTW supports the research within the EVOLVE project (nr. 18484). European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie actions, under grant agreement (nr. 872907).

Model Predictive Control (NMPC), which optimises the steering angle and the longitudinal force of a vehicle. However, in such scenarios, where longitudinal and lateral dynamics are coupled, the uncertainties and inaccuracies due to the tyre's non-linear behaviour pose a particularly challenging problem [1, 2]. Therefore, we focus on developing a Student-t Process (STP) combined with a Model Predictive Contouring Control (MPCC), which improves the prediction model, reducing the tyre model mismatches and minimising the vehicle lateral state uncertainties. An L-MPCC based on a sparse Gaussian Process (GP) has recently been proposed for lap-time optimisation [4]. It leverages the capacity of the GP to predict the mismatches between the prediction model and the measured vehicle states: longitudinal and lateral velocity and yaw rate. Furthermore, the vehicle is constrained inside the track, tightening the track boundaries with the vehicle position uncertainty. The latter is first computed as the GP's posterior covariance, then open-loop propagated along the prediction horizon using successive linearisation similar to an Extended Kalman Filter (EKF) [4]. However, the propagated position uncertainty can increase exponentially over the prediction horizon, strongly limiting or eliminating the allowable driving area. This results in a very conservative controller. A possible alternative is to consider the uncertainty in the vehicle velocity states rather than in the vehicle position [3]. For instance, the NMPC cost function can be extended by the vehicle lateral velocity and yaw rate variance. Thus, the latter can be minimised to reduce the operating time in the unstable region of the vehicle. However, these variances are computed and propagated by linearising the prediction model with a constant tyre model uncertainty optimised offline. This simplification does not consider the remarkably different accuracy of the tyre model in different operating regions, and therefore, it does not reduce the prediction model mismatches in all regions.

We propose an L-MPCC based on an STP, which predicts the mismatches between the prediction model and the vehicle yaw rate and the lateral tyre forces measured by intelligent (force sensing) bearings [5]. The proposed controller minimises the vehicle's operating time at unstable working points, thanks to the STP posterior covariance of the tyre forces. The latter ones are used to compute and propagate the vehicle lateral state uncertainties along the prediction horizon.

The contributions of this paper are threefold. The first is the development of the L-MPCC based on an STP, which directly reduces the prediction model mismatches in the tyre model rather than in the vehicle velocity states typically used [4]. This results in successfully performing evasive manoeuvres at 8.5% higher velocity than the current state-of-the-art MPCC [1]. The second contribution is related to the STP, which improves the outlier resistance of the state-of-the-art GP. Furthermore, the STP posterior covariance depends on the observed measurements, providing a higher variance than a GP for operating points different than in the training set and a lower one in the opposite situation [8]. The third contribution is improving the vehicle stability by reducing the sideslip angle peak of a 76% during an evasive manoeuvre, thanks to the reduction of the model mismatches and the minimisation of the vehicle lateral

state uncertainties. Thus, the controller decreases the time the vehicle spends in operating points close to the vehicle’s handling limits.

2 Learning-Based Model Predictive Contouring Control

The proposed L-MPCC controller is based on an MPCC for obstacle avoidance at the limit of handling [1]. The prediction model is a nonlinear single-track vehicle model. The Cartesian reference system describes the vehicle kinematics, as the contouring formulation requires. The vehicle dynamics, the longitudinal (v_x) and lateral (v_y) velocity, and the vehicle yaw rate (r) are described as follows:

$$\begin{cases} \dot{v}_x = \frac{(F_{x,F} \cos(\delta) - (F_{y,F} + \Delta F_{y,F}) \sin(\delta) + F_{x,R} - F_{drag})}{m} + (r + \Delta r) v_y \\ \dot{v}_y = \frac{(F_{x,F} \sin(\delta) + (F_{y,F} + \Delta F_{y,F}) \cos(\delta) + F_{y,R} + \Delta F_{y,R})}{m} - (r + \Delta r) v_x \\ \dot{r} = \frac{((F_{y,F} + \Delta F_{y,F}) \cos(\delta) l_f - (F_{y,R} + \Delta F_{y,R}) l_r + F_{x,F} \sin(\delta) l_f)}{I_{zz}} \end{cases} \quad (1)$$

where the road-wheel angle (δ) and the longitudinal force at the front ($F_{x,F}$) and rear axle ($F_{x,R}$). The lateral front and rear tyre forces, respectively $F_{y,F}$ and $F_{y,R}$, are computed using validated Fiala tyre model [2]. However, during an evasive manoeuvre, the prediction model inaccuracies can increase significantly, so the model mismatches of the front ($\Delta F_{y,F}$) and rear ($\Delta F_{y,R}$) lateral tyre forces and the yaw rate (Δr) are computed by an STP.

The proposed L-MPCC cost function is responsible for ensuring path tracking, maintaining the physical feasibility of the control inputs, prioritising obstacle avoidance in case of collision risk, and minimising the uncertainties of the prediction model, thus limiting the operating time in nonlinear regions. The cost function (J) is defined as follows:

$$\begin{aligned} J = \sum_{i=1}^N & \left(\sum_{j=1}^{N_{Obs}} (q_{e_{Obs,j,i}} e_{Obs,j,i}^2) + \sum_{k=1}^{N_{Edg}} (q_{e_{Edg,ki}} e_{Edg,ki}^2) + q_{\delta} \delta_i^2 + q_{\dot{F}_x} \dot{F}_{x,i}^2 \right. \\ & \left. + q_{e_{Con}} e_{Con,i}^2 + q_{e_{Lag}} e_{Lag,i}^2 + q_{e_{Vel}} e_{Vel}^2 \right) + \sum_{j=1}^{N_{Prob}} \left(q_{\sigma_r} \sigma_{r,j}^2 + q_{\sigma_{v_y}} \sigma_{v_y,j}^2 \right) \end{aligned} \quad (2)$$

where N is the length of the prediction horizon, N_{Obs} and N_{Edg} are respectively the number of obstacles and road edges, and $q_{e_{Obs}}$, e_{Obs} , $q_{e_{Edg}}$ and e_{Edg} are used to prioritise obstacle avoidance and keeping the vehicle away from the road edges in case of emergency [1]. The tracking performance is ensured by the minimisation of the contouring (e_{Con}), lag (e_{Lag}) and velocity e_{Vel} errors with their respective cost terms, i.e. $q_{e_{Con}}$, e_{Lag} , $q_{e_{Vel}}$. The lateral velocity and yaw rate uncertainties are minimised through the σ_{v_y} and σ_r parameters of the cost function. The σ_{v_y} and σ_r are propagated open-loop, so they can increase rapidly, potentially overpowering the other elements of the cost function. To avoid the problem, they are propagated not over the entire prediction horizon, $N = 30$, but only for the reduced horizon $N_{Prob} = 20$ [4]. Furthermore, the cost parameters,

$q_{\sigma_{v_y}}$ and q_{σ_r} , are tuned not to exceed the cost related to obstacle avoidance prioritisation. Therefore, the proposed controller pushes the vehicle to work at operating points close to normal driving. However, it allows the vehicle to drive at operating points close to the handling limits if it helps avoid a collision.

3 Student-T Process and Uncertainty Propagation

This study uses an STP as a stochastic process with a Laplace inference to estimate model mismatches and uncertainties for two reasons. First, the Student-t distribution allows the algorithm to define the level of Kurtosis, reducing the influence of the outliers and improving the accuracy of the predictions [7,9]. The robustness to outliers is an essential property for L-MPCC, which relies on online measurements with sensor noise. For instance, a single outlier can highly influence the GP prediction; pushing the posterior means far away from the level of the other observations, which is not the case in an STP [9]. The second reason is that the STP posterior covariance directly depends on the observed measurements and not only on the location of the observed measurements. This implies that the covariance increases when the measurements vary more than expected, i.e., a higher difference between training and test sets, and vice-versa decreases when the difference is lower [7,8]. Once again, this is an essential property for an L-MPCC, which can work in conditions with a high discrepancy from the training scenario. Regarding the kernel selection, this work opts for the automatic relevance determination Matérn 5/2 function due to its generalisation capabilities and interpretability properties. The STP is implemented using the GPML Matlab Code version 4.2 [6].

We predict the $\Delta F_{y,F}$, $\Delta F_{y,R}$ and Δr model mismatches, rather than the errors in v_x and v_y as usually done in literature [4] because it allows the reduction of the mismatches in a proactive way. It also directly targets the source of the errors, i.e., the tyre forces, and not the states that depend on them. This is possible thanks to the availability of tyre lateral forces at the front and rear axle, measured by the intelligent bearings [5], the vehicle yaw rate and the control inputs F_x and δ . Thus, we assume that the discrepancies are independent of the vehicle position [4] and that the velocity states, not measurable through the standard sensor setup for passenger vehicles, are entirely dependent on the tyre forces. The training data is derived from the algebraic difference between the measurements and the nominal model predictions. The training set is composed of 2 double lane changes at respectively 55 km/h and 80 km/h, and a double lane change with collision avoidance prioritisation at 55 km/h. The manoeuvres are selected to identify the mismatches in the linear and nonlinear regions of the vehicle. The test set comprises a double lane change at 60 km/h and the same manoeuvre with collision prioritisation. Despite the high prediction model accuracy and the reduction of the model mismatches, some unmodeled effects can still influence the prediction model. These are captured by the STP posterior covariance for the front and rear lateral tyre force, respectively $\sigma_{F_{y,F}}$ and $\sigma_{F_{y,R}}$. Thus, the tyre forces are not only evaluated as a single point but also as a

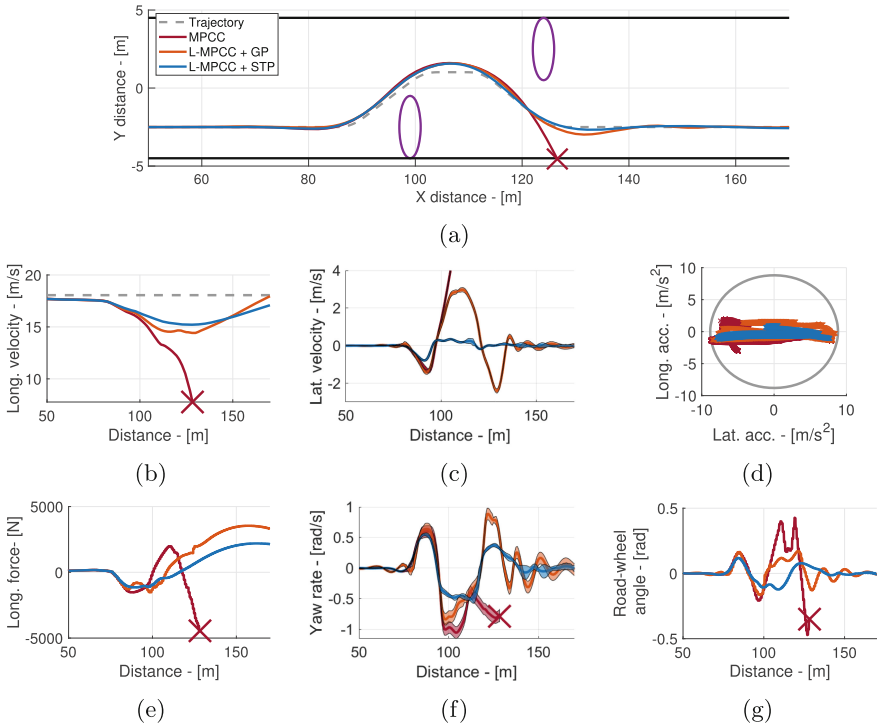


Fig. 1. States and control inputs in a double lane change.

distribution [3]. The STP variance and degrees of freedom are transformed and simplified into a Gaussian variance to be able to propagate it through the prediction horizon and relate them to the lateral states related to vehicle stability, i.e. v_y and r , using a procedure similar to an EKF [3,4].

4 Results

Figure 1 shows the performance of the evaluated controllers in a double lane change with collision avoidance prioritisation at 65 km/h. Figure 1a shows that only the learning-based controllers can successfully avoid the two obstacles and stay inside the road track. Vice versa, the MPCC baseline [1] leaves the track on the right side at 130 m. The improvement is due to reduced model mismatches in the L-MPCC. For instance, the proposed L-MPCC+STP reduces the root mean square of the $F_{y,F}$, $F_{y,R}$ and r mismatches by respectively 33.14%, 24.85% and 60.61% compared to the MPCC and further 20.14%, 11.85% and 61.33% with respect to the L-MPCC+GP. This proves the importance of implementing an STP rather than a GP to predict the model mismatches. The enhanced performance of the proposed controller is also visible from the reduction of the vehicle lateral velocity peak; see Fig. 1c. Furthermore, the L-MPCC+STP can reduce the

lateral velocity and, consequently, the sideslip angle peak by 76%, maintaining an overall higher velocity during the manoeuvre than the L-MPCC+GP. Figures 1c and 1f show that the vehicle sideslip angle reduction and the increased operating time in the stable vehicle behaviour are not only due to the lower model mismatches but also to the reduced uncertainties for the vehicle lateral states. However, it has a lower vehicle longitudinal acceleration when the obstacle avoidance manoeuvre is over to achieve the desired velocity, Fig. 1e.

5 Conclusions

This paper presented a novel Learning-based Model Predictive Contouring Control based on a Student-t Process for evasive manoeuvres with an online minimisation of model mismatches and uncertainties. In a high-fidelity simulation environment, we demonstrate that our proposed controller successfully avoids obstacles, keeping the vehicle stable while driving a double-lane change manoeuvre at a higher velocity compared to a non-learning-based baseline. Furthermore, the proposed controller reduces the peak of vehicle sideslip angle by 76%. The performance enhancement is due to the properties of the Student-t Process, which can further reduce the prediction model mismatches and better capture the vehicle state uncertainties. Future work involves the proposed controller implementation on rapid prototyping hardware to prove its real-time capability.

References

1. Bertipaglia, A., Alirezai, M., Happee, R., Shyrokau, B.: Model predictive contouring control for vehicle obstacle avoidance at the limit of handling. arXiv preprint [arXiv:2308.06742](https://arxiv.org/abs/2308.06742) (2023)
2. Bertipaglia, A., et al.: Model predictive contouring control for vehicle obstacle avoidance at the limit of handling using torque vectoring. arXiv preprint [arXiv:2405.10847](https://arxiv.org/abs/2405.10847) (2024)
3. Brown, M., Gerdes, J.C.: Robust stabilization and collision avoidance through minimizing open-loop velocity uncertainty. In: IEEE Intelligent Vehicles Symposium, pp. 259–264 (2020)
4. Kabzan, J., Hewing, L., Liniger, A., Zeilinger, M.N.: Learning-based model predictive control for autonomous racing. *IEEE Robot. Autom. Lett.* **4**(4), 3363–3370 (2019)
5. Kerst, S., Shyrokau, B., Holweg, E.: A model-based approach for the estimation of bearing forces and moments using outer ring deformation. *IEEE Trans. Industr. Electron.* **67**(1), 461–470 (2020)
6. Rasmussen, C.E., Nickisch, H.: Gaussian processes for machine learning toolbox. *J. Mach. Learn. Res.* **11**, 3011–3015 (2010)
7. Shah, A., Wilson, A., Ghahramani, Z.: Student-t processes as alternatives to gaussian processes. In: Artificial Intelligence and Statistics, pp. 877–885 (2014)
8. Tracey, B.D., Wolpert, D.: Upgrading from gaussian processes to student-T processes. In: AIAA Non-Deterministic Approaches Conference, p. 1659 (2018)
9. Vanhatalo, J., Jylänki, P., Vehtari, A.: Gaussian process regression with student-T likelihood. In: Advances in Neural Information Processing Systems, vol. 22 (2009)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

