

# Fusion of Heterogeneous Friction Estimates for Traction Adaptive Motion Planning and Control

Lars Svensson and Martin Törngren\*

May 17, 2021

## Abstract

Traction adaptive motion planning and control has potential to improve an automated vehicle's ability to avoid accident in a critical situation. However, such functionality require an accurate friction estimate for the road ahead of the vehicle that is updated in real time. Current state of the art friction estimation techniques include high accuracy local friction estimation in the presence of tire slip, as well as rough classification of the road surface ahead of the vehicle, based on forward looking camera. In this paper we show that neither of these techniques in isolation yield satisfactory behavior when deployed with traction adaptive motion planning and control functionality. However, fusion of the two provides sufficient accuracy, availability and foresight to yield near optimal behavior. To this end, we propose a fusion method based on heteroscedastic gaussian process regression, and present initial simulation based results.

## 1 Introduction

One of the remaining challenges in the development of automated driving (AD) and advanced driver assistance systems (ADAS) is fully automated operation in critical situations i.e., situations where an accident is imminent, for example due to unpredictable behavior of other traffic agents or rapid changes in the operational conditions. In critical situations, full utilization of the physical capacity of the vehicle is paramount to minimizing the risk of an accident, since overly cautious behavior will impair accident mitigation performance. This motivates research in motion planning and control at the limits of handling - a field that has seen a lot of progress in the last decade [1–5]. At the limit of tire adhesion, longitudinal, lateral and yaw dynamics are tightly coupled. Therefore, the traditional approach of performing motion planning, lateral- and longitudinal control in isolation is being outperformed by combined optimization approaches, where the lateral and longitudinal motion planning and control is performed by solving a single optimization problem [2, 4].

To further complicate the problem of motion planning and control in critical situations, the physical motion capacity of a vehicle is far from static. Instead, it varies greatly with the operational situation, e.g. road surface and tire conditions. In our previous work, we tackle the problem of motion planning and control under time-varying traction constraints [6, 7]. We show that such an approach improves the vehicle's ability to avoid accident in critical situations, and thus have potential to improve safety performance of AD and ADAS systems. However, the performance increase hinges on run-time availability of a tire-road friction estimate over the planning horizon, i.e., *ahead* of the vehicle. We refer to this concept as a *predictive friction estimate*. One of

---

\*L. Svensson and M. Törngren are with the Department of Machine Design, KTH Royal Institute of Technology, Stockholm, Sweden. {larsvens,martint}@kth.se



Figure 1: Illustration road surface classification within a lane, based on forward looking camera. Classes: orange - dry, purple - wet, blue - snow/ice

the conclusions drawn from our previous work is that in critical situations, accident avoidance performance is negatively affected by both over- and under-estimation in the predictive friction estimate. Over-estimation leads to the planning of dynamically infeasible motions, and under-estimation to poor utilization of the available traction, which can needlessly lead to accidents in e.g., a collision avoidance situation. To summarize, accurate predictive friction estimation is key for autonomous accident mitigation at the limits of handling.

Fortunately, research over the past two decades have presented several viable methods of obtaining such estimates [8,9]. In later years, the rapid development in supervised machine learning and computer vision have greatly improved performance of such functionality, mainly through road surface classification from forward looking camera images [10,11], providing a predictive estimate online with high availability. However, due to the limited number of visually distinguishable road surface classes (e.g., dry, wet, snow/ice), accuracy is poor compared to state of the art local friction estimation methods. Local friction estimation methods [12,13] use vehicle dynamics models and traditional estimation techniques to estimate the friction coefficient locally around the vehicle. Such methods have high accuracy but the estimate is only available in the presence of tire slip [14], and provide no foresight about the conditions ahead of the vehicle. Considering the high level properties of the two paradigms;

- Camera based: high availability, foresight, low accuracy,
- Local: low availability, lack of foresight, high accuracy,

we hypothesize that using either method individually could lead to poor performance in the context of traction adaptive motion planning. However, the indication that the two paradigms are to some extent complementary, motivates investigation in fusion of the two.

Gaussian process (GP) regression [15] is a nonparametric Bayesian approach to regression, that has become an increasingly popular modeling tool in optimization based control. In contrast to the traditional approach of modelling the system to be controlled from first principles, a GP representation of the system does not make strong assumptions of the underlying dynamics, instead the dynamics, including a direct representation of model uncertainty, is inferred from input/output data. Previous approaches uses GP representations either to completely replace, [16–19] or to complement a traditional dynamics model [20–23], by representing residual dynamics and uncertainty.

This work differentiates from the previous efforts by using a GP representation of time varying input constraints, rather than the dynamics. We extend the framework of [7], with a novel, GP regression based method for merging predictive and local friction estimates, and quantitatively

represent uncertainty in the fused estimate. We show that such a fusion strategy is able to exploit the virtues of both camera based and local friction estimation. In summary, the contributions of this paper are:

- Justification for the argument that neither camera based nor local state of the art friction estimation techniques in isolation is sufficient for traction adaptive motion planning and control.
- Preliminary results indicating that fusion of the two complementary estimation paradigms can produce near-optimal behavior.
- A proposed GP regression based method for fusion of predictive and local friction estimates that benefits from the virtues of both paradigms.

## 2 Background

Here follows a summary of the traction adaptive motion planning and control framework which forms a basis of this work. See [7] for a detailed account. At each planning time  $t$ , a constrained finite time optimal control problem

$$\begin{aligned} \min_{u_{0|t}, \dots, u_{N-1|t}} \quad & J(\mathcal{T}_t) \\ \text{s.t.,} \quad & x_{k+1|t} = f(x_{k|t}, u_{k|t}), \\ & u_{k|t} \in \mathcal{U}_{k|t}, \\ & x_{k|t} \in \mathcal{X}_{k|t}, \\ & \forall k \in \{0, 1, \dots, (N-1)\}, \\ & x_{0|t} = x_t, \quad x_{N|t} \in \mathcal{X}_{k|t}, \end{aligned}$$

is solved over a prediction horizon  $k \in \{0, 1, \dots, (N-1)\}$ . The predicted input sequence  $\{u_{0|t}, u_{1|t}, \dots, u_{N-1|t}\}$  is applied to the vehicle model  $f(\cdot, \cdot)$ , yielding the corresponding predicted state sequence  $\{x_{0|t}, x_{1|t}, \dots, x_{N|t}\}$ . We compactly denote a predicted trajectory consisting of state and control sequences as  $\mathcal{T}_t = \{\{x_{k|t}\}_{k=0}^N, \{u_{k|t}\}_{k=0}^{N-1}\}$ . The vehicle model  $f(\cdot, \cdot)$  is a dynamic bicycle model with tire force inputs, expressed in a road aligned coordinate frame (See Appendix B of [7]). The behavior of the resulting optimal trajectory  $\mathcal{T}_t^*$  is dictated by a cost function  $J(\mathcal{T}_t)$  and time varying state and input constraints  $\mathcal{X}_{k|t}, \mathcal{U}_{k|t}$ . State constraints  $\mathcal{X}_{k|t}$  are used to encode lane boundaries as well as static and dynamic obstacles.

We use the time varying input constraint  $\mathcal{U}_{k|t}$  as a mechanism to encode the tire adhesion limit in the presence of local traction variations. For each timestep  $k$  in the prediction horizon, tire force inputs are upper bounded by a convex polygon. The size and shape of each polygon is set based on a tire model, predicted normal forces and a predictive friction estimate  $\hat{\mu}(s)$ .

## 3 Fusing Heterogeneous Friction Estimates through GP Regression

The aim of the proposed fusion method is to obtain a vector  $\hat{\mu}(s)$ , that represents the tire-road friction at a sequence of discrete distances ahead of the vehicle,  $s \in \{0, ds, \dots, s_f\}$ , with  $s_f$  corresponding to the spatial extension of the prediction horizon of the planning/control framework introduced in Section 2. The estimate should be

- (a) Conservative - in the sense that the probability of over-estimation is low. This is a generally applied principle in using online friction estimates for vehicle control, because over-estimation of the friction coefficient may lead to loss of control of the vehicle [24, 25].
- (b) Not overly conservative - to avoid that collision avoidance performance deteriorates, e.g., in the manner described in [7].

Our proposed fusion strategy is to infer a distribution over functions  $f_\star(s)$  at  $s \in \{0, ds, \dots, s_f\}$ , that with high probability encompasses the true predictive friction values  $\mu(s)$ . The distribution is obtained through heteroscedastic GP regression based on input data from predictive *and* local friction estimates, including their associated estimation uncertainties, which are propagated from the individual methods. From such a distribution  $f_\star(s)$  we then select a  $\hat{\mu}(s)$  that satisfies both (a) and (b). See Fig. 2 for a visual example and Appendix A for a brief outline of GP regression preliminaries.

We start by selecting a prior  $f(s) \sim \mathcal{GP}(\eta(s), k(s, s'))$ , with mean function  $\eta(s)$  and kernel function  $k(s, s')$ , based on domain knowledge, i.e. that it is probable that  $\mu(s)$  resides in the interval 0.1 to 1.0, and that some spatial correlation exists, i.e., that positions close to each other are likely to have similar friction coefficients. For  $k(s, s')$ , we use a squared exponential kernel, Equation (2) of Appendix A, and select  $\eta(s)$  and  $\sigma_f$  of (2) such that the interval  $[0.1, 1.0]$  represents a 95% confidence interval for the prior distribution over functions, as illustrated in Fig. 2. The prior is determined offline, in contrast to the subsequent steps that are to be executed online.

Next, we prepare the input data. We combine friction estimates and associated margins of error from predictive and local friction estimation algorithms into two input data vectors

$$\hat{\mu}'(s), m'(s) = \begin{cases} \hat{\mu}_l, m_l & \text{if local available \& } s < s_l, \\ \hat{\mu}_p(s), m_p(s) & \text{otherwise.} \end{cases}$$

Subscripts  $(\cdot)_l$  and  $(\cdot)_p$  indicate local and predictive estimates respectively. The predictive friction estimate  $\hat{\mu}_p(s)$  is set as the mean value of its road surface class along with a margin of error  $m_p(s)$  set such that for all  $s \in \{0, ds, \dots, s_f\}$ ,  $\hat{\mu}_p(s) \pm m_p(s)$  encompasses all friction coefficient values within the road surface class associated with  $s$ . If tire slip is sufficient to make a local estimate  $\hat{\mu}_l$  available, the predictive estimates at  $s \in \{0, ds, \dots, s_l\}$  are replaced by  $\hat{\mu}_l$  and  $m_l$ . The local margin of error  $m_l$  is set as the worst case estimation error of the local estimation method. A tunable threshold  $s_l$  dictates how far ahead the local estimate will influence the input data.

Next, we compute the mean function and covariance matrix of the posterior predictive distribution  $f_\star(s) \sim \mathcal{GP}(\eta_\star(s), \Sigma_\star)$  at locations  $s \in \{0, ds, \dots, s_f\}$ , from the prior and the input data  $\hat{\mu}'(s)$  and  $m'(s)$  using equations (4) and (5) of Appendix A. Here, we interpret the interval  $\hat{\mu}'(s) \pm m'(s)$  as a 95% confidence interval to obtain the standard deviation of the input data  $\sigma'(s) = m'(s)/1.96$ . Finally, we select the fused predictive friction estimate as the lower bound of the 95% confidence region of  $f_\star(s)$ .

$$\begin{aligned} \hat{\mu}(s) &= \eta_\star(s) - 1.96\sigma_\star(s) \quad \text{with} \\ \sigma_\star(s) &= \sqrt{\text{diag}(\Sigma_\star)}. \end{aligned}$$

Fig. 2 shows an example plot of the prior distribution, the input data, the posterior distribution and the final predictive friction estimate for a single planning iteration, in a scenario where the vehicle is approaching a low  $\mu$  section. The resulting estimate benefits both from foresight of the predictive and the accuracy of the local estimate to meet conditions (a) and (b).

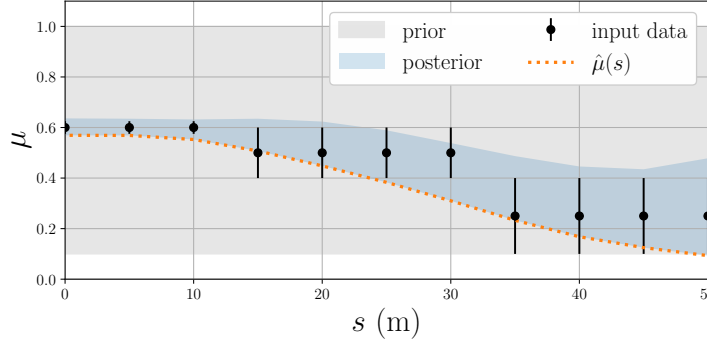


Figure 2: Example showing how the proposed method merges domain knowledge encoded in the prior (gray) with local and predictive friction estimates (black dots) of varying uncertainty (black bars). The resulting estimate is selected as the lower bound of the 95% confidence region of the posterior distribution

## 4 Results and Discussion

In this section, we evaluate the impact of friction estimation performance in the context of traction adaptive motion planning. In a simulated environment, we emulate (i.e., replicate the output of) state of the art predictive and local friction estimation algorithms in terms of their availability, foresight and accuracy. We then compare local only, predictive only and fused configurations to ground truth, in terms of the resulting motion behavior.

### 4.1 Delimitations and Experimental Setup

For the purposes of the evaluation we emulate rather than implement state of the art friction estimation algorithms. This choice is motivated by practical as well as methodological reasons. The focus of this paper is to evaluate various strategies for handling uncertain and not always available friction estimates, in the context of traction adaptive motion planning. Given this purpose, implementing our own and/or a selection of state-of-the-art methods for friction estimation would a) be a substantial time investment and b) introduce an additional risk of biasing the resulting motion behavior through stochastic variations or imperfections in the implementation. Instead, we have opted for parameterizing the performance of state of the art estimation algorithms, and emulating them in a simulation environment. This way we can draw conclusions based on the fundamental properties of the methods. Simulations are performed using an extended version of the Gazebo based FSSIM simulator [26], that have been modified and tuned to represent the REVERE lab’s Volvo FH16 test vehicle, which was used in the experimental evaluation of [7].

Based on the literature summarized in Section 1, we select parameter specifications to represent current state-of-the-art performance in local and predictive road friction estimation in Table 1. Variables  $e_p$  and  $e_l$  denote estimation errors for the predictive and local methods respectively. The parameter  $\lambda_t \in [0, 1]$  denotes the momentary utilization of available tire force.

### 4.2 Evaluated Configurations

We compare four different friction estimation configurations: Ground truth (GT), Local Only (L), Predictive only (P), and Fused (F). Here follows a description of each configuration.

	<b>Accuracy</b>	<b>Availability</b>
Local (L)	$\max\{ e_l \} = 0.025, [14]$	$\lambda_t > 0.5, [14]$
Predictive (P)	$\max\{ e_p \} \approx 0.1 \text{ to } 0.3$ classes: dry, wet, snow/ice [10]	$s \in \{0, ds, \dots, s_f\},$ $s_f = 50\text{m}$

Table 1: Selected performance parameters for state of the art friction estimation methods

The first configuration, GT, is a non-realistic benchmark, representing the optimal behavior on the current road surface. The traction adaptive motion planner/controller is provided the true friction values from the track ahead of the vehicle, i.e.,

$$\hat{\mu}(s) = \mu_{\text{gt}}(s) \quad \text{for all } s \in \{0, ds, \dots, s_f\}.$$

In the local only configuration, L, we emulate a local friction estimate being propagated over the whole prediction horizon, i.e.,  $\hat{\mu}(s) = \hat{\mu}_l$  for all  $s \in \{0, ds, \dots, s_f\}$ , with

$$\hat{\mu}_l = -\max\{|e_l|\} + \begin{cases} \mu_{\text{gt}}(0) + e_l & \text{if local available} \\ \hat{\mu}_0 & \text{otherwise} \end{cases}$$

The estimate is conservative in the sense that we subtract the worst case error before passing the estimate to the planner, so as to prevent over-estimation. If the local estimate is unavailable, the latest available estimate  $\hat{\mu}_0$  is used.

In the predictive only configuration, P, the planner/controller is provided a conservative predictive friction estimate. The estimate is conservative in the sense that it takes the minimum value of the surface class associated with the point  $s$

$$\hat{\mu}(s) = \begin{cases} 0.6 & \text{if } 0.6 < \mu_{\text{gt}}(s) & \text{class: dry} \\ 0.4 & \text{if } 0.4 \leq \mu_{\text{gt}}(s) < 0.6 & \text{class: wet} \\ 0.1 & \text{if } 0.1 \leq \mu_{\text{gt}}(s) < 0.4 & \text{class: snow/ice.} \end{cases}$$

In the fused configuration F, we emulate both local and predictive friction estimation such that the combined information from configurations 2 and 3 is made available.

$$\hat{\mu}'(s), m'(s) = \begin{cases} 0.8, 0.2 & \text{if } 0.6 < \mu_{\text{gt}}(s) & \text{class: dry} \\ 0.5, 0.1 & \text{if } 0.4 \leq \mu_{\text{gt}}(s) < 0.6 & \text{class: wet} \\ 0.25, 0.15 & \text{if } 0.1 \leq \mu_{\text{gt}}(s) < 0.4 & \text{class: snow/ice} \\ \mu_{\text{gt}}(0) + e_l, \max\{|e_l|\} & \text{if local available \& } s < s_l \end{cases}$$

For the predictive estimates, instead of directly providing the lowest value in each class, we extract the mean values  $\hat{\mu}'(s)$  along with the margins of error  $m'(s)$ . With this input data, we obtain a fused estimate  $\hat{\mu}(s)$  using the method described in Section 3.

Notice that configurations L, P and F, all set conservative estimates with respect to the available information. As mentioned in Section 3, this is a generally applied principle in using online friction estimation for vehicle control.

We evaluate the four configurations in two simulated critical situations; Scenario 1, Turn at low  $\mu$ , and Scenario 2, Collision avoidance at high  $\mu$ . For both scenarios, we set adversarial worst case friction estimation errors, i.e., worst case over-estimation in the low  $\mu$  turn scenario (where inability to realize large planned tire forces will lead to accident) and worst case under-estimation in the collision avoidance scenario (where overly restrictive tire force constraints will lead to accident).

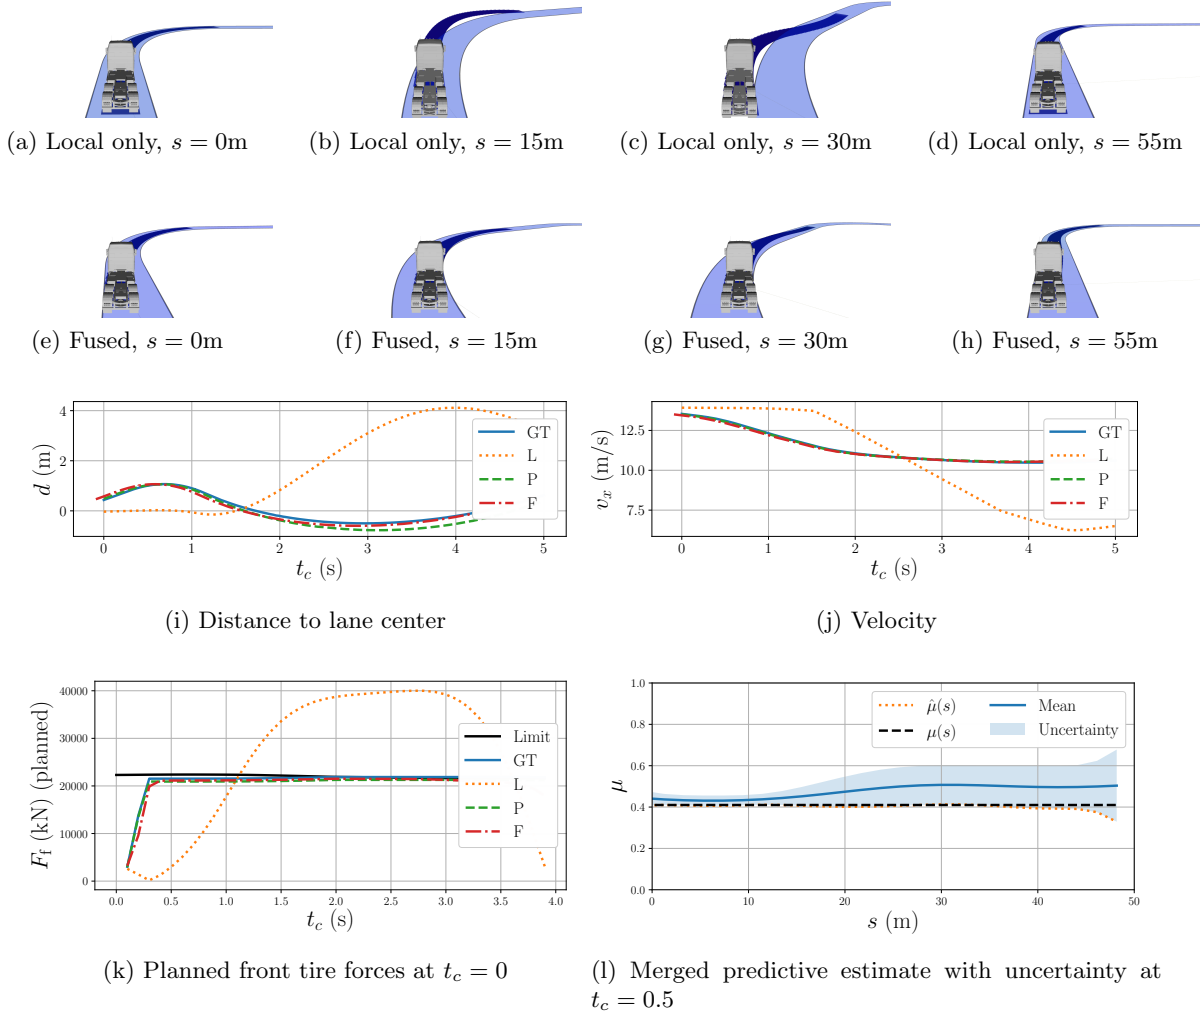


Figure 3: Simulation results for the first critical scenario - a 90 degree turn at low  $\mu$  with worst case over-estimation of friction. Subfigures 3a through 3h show snapshots from the online visualization of the planned motion. The first row (3a - 3d) correspond to the local only configuration and the second row (3e - 3h) correspond to the proposed fused configuration. Subfigures 3i and 3j show the vehicle's distance to the lane center and velocity for the four configurations. Subfigure 3k show planned front tire forces at  $t_c = 0\text{s}$ ,  $s = 0\text{m}$  and the prevailing force limits. Subfigure 3l shows the corresponding friction estimate for the fused configuration, including uncertainty.

### 4.3 Scenario 1 - Turn at low $\mu$

At the start of the first scenario,  $t_c = 0$ ,  $s = 0$ , the vehicle is approaching a 90 degree turn at a velocity of  $12\text{m/s}$ . Its objective is to maintain velocity and stay at the lane center. The traction in the corner is locally reduced,  $\mu(s) = 0.4$  for  $s \geq 0\text{m}$ , from previously good conditions  $\mu(s) = 0.8$  for  $s < 0\text{m}$ . Fig. 3, shows an analysis of the generated motion behavior for the four friction estimation configurations.

At first glance at Figures 3i and 3j, we see that configurations GT, P and F generate similar behavior while L shows distinctly different behavior. Under the L configuration, the vehicle enters the turn at the center of the lane, and plans to stay at the center throughout the turn, Fig. 3a. Upon

entering the corner, Fig. 3b, the vehicle starts slipping toward the outside of the lane, deviating from the planned motion. At the middle of the turn, Fig. 3c, the vehicle has slipped well into the opposing lane, at risk of colliding with oncoming traffic. By now however, the vehicle has managed to reduce its speed such that it is able to recover to the lane, Fig. 3d.

The reason behind this undesired motion behavior is the that the L configuration relies entirely on a local friction estimation, which is only available when tire force utilization  $\lambda > 0.5$ . While travelling on the straight road before the turn at constant velocity,  $\lambda \ll 0.5$ . Thus tire force constraints are set based on the latest available friction estimate, here  $\hat{\mu}_0 = 0.8$ . As is shown in Fig. 3k, the planned tire forces at  $t_c = 0$  for configuration L exceeds the prevailing tire force limit by a wide margin. The moment the vehicle starts turning, Fig. 3b, the planner reduces the tire force constraints and adjusts the planned motion, but at this time, it is physically impossible to avoid veering out into the opposing lane.

In contrast, the other three configurations GT, P and F exhibit the behavior represented by F in Subfigures 3e - 3h. Before entering the turn, Fig. 3e, the vehicle has positioned itself on the outside of the lane, and plans to cut across to the inside of the corner by the middle of the turn and then back to the outside of the lane by the end of the turn, while slightly reducing speed. As is shown in Figures 3f, 3g and 3h, the vehicle does not deviate from the planned motions, and the updated plans are consistent with the previous ones.

The reason behind the difference in behavior is the foresight provided by the predictive friction estimate. This in combination with a conservative selection of tire force limits ensures that the planned tire forces do not exceed the physical limits in either of the three configurations 3k. The traction adaptive planner/controller realizes the reduced tire forces by a coordinated speed reduction and lateral motion in the lane as shown in Figures 3i and 3j.

For the F configuration, because the maneuver is initiated earlier, the local friction estimate becomes available earlier compared to the L configuration. Fig. 3l shows the fused friction estimate at  $t_c = 0.5$ . In this scenario, where the emulated friction estimation error equals the worst case over-estimation, the fused predictive estimate more or less coincides with the ground truth friction value. The same is the case for P. For this reason, the three configurations show very similar planned tire forces in Fig. 3k, and consequently similar motion behavior. In cases where the estimation errors are below the worst case over-estimation, configurations P and F will generate more conservative planned forces, through keeping the same lateral motion pattern, but reducing velocity even more before the turn. Between the two, configuration F will be less conservative due to the higher accuracy of the local estimate. Next, we evaluate the effect of this phenomenon in collision avoidance.

#### 4.4 Scenario 2 - CA at high $\mu$

At the start of the second scenario,  $t_c = 0$ ,  $s = 0$ , the vehicle is cruising at a speed of 20m/s, when an obstacle appears suddenly in the middle of the lane 20m ahead of the vehicle. The local traction on the road surface is high, i.e.,  $\mu(s) = 1.0$  for  $s \geq 0$ m. Fig. 4, shows an analysis of the generated motion behavior for the four friction estimation configurations.

Similar to the result of Section 4.3, three configurations yield similar behavior and one shows a distinct difference in Figures 4h, 4i and 4j. However, contrary to the turn scenario, here it is the behavior for configuration P that is distinctly different, while GT, L and F generate similar behavior. Upon detecting the obstacle under the P configuration, Fig. 4b, the vehicle plans an evasive maneuver. However, under the conservative tire force constraints associated with the friction estimate, green dashed line in Fig. 4j, no collision free trajectory exists. As such, it selects the least violating option and subsequently collides with the obstacle, Fig. 4c, at an impact velocity of



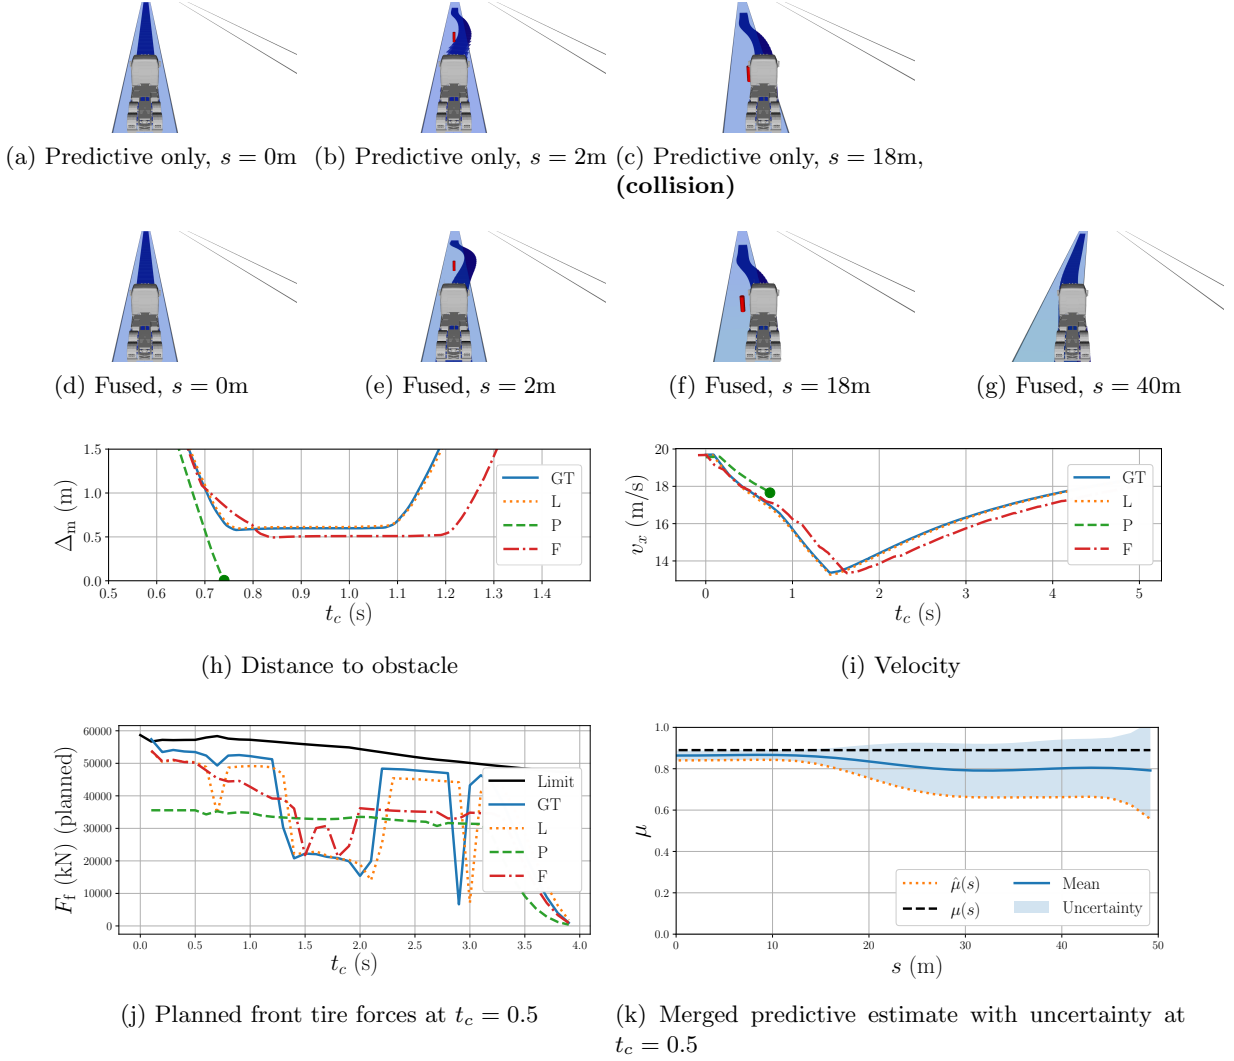


Figure 4: Simulation results for the second critical scenario - collision avoidance at high  $\mu$  with worst case under-estimation of friction. Subfigures 4a through 4g show snapshots from the online visualization of the planned motion. The first row (4a - 4d) correspond to the predictive only configuration and the second row (4d - 4g) correspond to the proposed fused configuration. Subfigures 4h and 4i show the vehicle's distance to the obstacle and velocity for the four configurations. Subfigure 4j show planned front tire forces at  $t_c = 0\text{s}$ ,  $s = 0\text{m}$  and the prevailing force limits. Subfigure 4k shows the corresponding friction estimate for the fused configuration, including uncertainty.

$\sim 17.5\text{m/s}$ .

The reason behind this undesirable behavior is the overly conservative tire force constraints, which in turn stem from the low accuracy of the predictive friction estimate. In order to ensure that tire forces are never over-estimated, the P configuration selects the lowest friction coefficient value of the associated road surface class. Here, in case of the worst case under-estimation, this corresponds to a 32.6% reduction in utilization of available tire force, compared to GT.

In contrast to P, configurations GT, L and F plans a more aggressive evasive maneuver, as shown in Fig. 4h and represented by F in Fig. 4e. The vehicle is able to track the more aggressive plan, Fig. 4f, clears the obstacle with a margin of  $\sim 0.5\text{m}$  and safely recovers to the lane, Fig. 4g.

The reason that the L and F configurations perform better than P here is that they benefit from the high accuracy of the local friction estimate. As soon as the vehicle engages in the evasive maneuver, the tire force utilization exceeds the availability threshold,  $\lambda > 0.5$  such that the local estimate can be utilized in the next planning iteration. Since  $\max\{|e_l|\} \ll \max\{|e_p|\}$ , the resulting tire force constraints are substantially less conservative (5.6% reduction in utilization of available tire force, compared to GT).

The L configuration propagates the local estimate and its worst case error over the whole prediction horizon. This turns out to be a good strategy for scenario 2 but a poor one for scenario 1. The F configuration in contrast performs well in both scenarios. Fig. 3l shows the fused friction estimate at  $s = 0$ . The accurate local estimate influences the initial part of the prediction horizon whereas the more conservative predictive estimate dominates the farther part. This way, the fusion strategy exploits the accuracy of the local estimate, while maintaining the foresight of the predictive estimate.

## 4.5 Discussion

The two scenarios selected for evaluation in this paper represent worst case situations with respect to friction estimation error. Hence performance in terms of accident avoidance will not be worse for any estimation error in the intervals  $[-\max\{|e_l|\}, \max\{|e_l|\}]$ ,  $[-\max\{|e_p|\}, \max\{|e_p|\}]$  provided the underlying assumptions on the algorithms, described in Section 4.1, hold. In scenario 1, in cases where the estimation errors are below the worst case over-estimation, P and F will yield more conservative planned forces, resulting in a larger distance to the traction limit. In scenario 2, in cases where the estimation errors are above the worst case under-estimation, L and F will yield even less conservative planned tire forces, resulting in behavior even closer to that of GT.

## 5 Conclusions and Future Work

In this work we have highlighted the impact of accuracy, availability and foresight of predictive friction estimation functionality for traction adaptive motion planning in critical situations. Initial simulations with emulated state of the art friction estimation algorithms indicate that fusing heterogeneous friction estimates enables exploiting the virtues of predictive (foresight, high availability) and local (high accuracy) friction estimates. Further, we show that the proposed fusion strategy is conservative in that it avoids over-estimation of the friction coefficient while not being overly conservative, even at worst case adversarial estimation errors (e.g., maximum 5.6% reduction in tire force utilization, compared to GT, in the collision avoidance scenario).

Next steps of this research work is to move development of the method and simulation based evaluation to real world validation on a test vehicle, with state of the art local and predictive friction estimation. Longer term future directions include extending the fusion method to include shared

uncertain friction estimates communicated via vehicle-to-vehicle and/or vehicle-to-infrastructure networks.

## Acknowledgements

This work was partially supported by ECSEL PRYSTINE project, grant agreement 783190.

## References

- [1] K. Berntorp, B. Olofsson, K. Lundahl, and L. Nielsen, “Models and methodology for optimal trajectory generation in safety-critical road-vehicle manoeuvres,” *Vehicle System Dynamics*, vol. 52, no. 10, pp. 1304–1332, 2014.
- [2] A. Liniger, A. Domahidi, and M. Morari, “Optimization-based autonomous racing of 1: 43 scale rc cars,” *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, 2015.
- [3] J. Funke, M. Brown, S. M. Erlien, and J. C. Gerdes, “Collision avoidance and stabilization for autonomous vehicles in emergency scenarios,” *IEEE Transactions on Control Systems Technology*, vol. 25, no. 4, pp. 1204–1216, 2017.
- [4] J. K. Subosits and J. C. Gerdes, “From the racetrack to the road: Real-time trajectory replanning for autonomous driving,” *IEEE Transactions on Intelligent Vehicles*, vol. 4, pp. 309–320, June 2019.
- [5] V. Fors, *Autonomous Vehicle Maneuvering at the Limit of Friction*, vol. 2102. Linköping University Electronic Press, 2020.
- [6] L. Svensson, M. Bujarbaruah, N. R. Kapania, and M. Törngren, “Adaptive trajectory planning and optimization at limits of handling,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3942–3948, 2019.
- [7] L. Svensson, M. Bujarbaruah, A. Karsolia, C. Berger, and M. Törngren, “Traction adaptive motion planning at the limits of handling,” *arXiv preprint arXiv:2009.04180*, 2020.
- [8] S. Khaleghian, A. Emami, and S. Taheri, “A technical survey on tire-road friction estimation,” *Friction*, vol. 5, no. 2, pp. 123–146, 2017.
- [9] M. Andersson, F. Bruzelius, J. Casselgren, M. Gäfvert, M. Hjort, J. Hultén, F. Håbring, M. Klomp, G. Olsson, M. Sjö Dahl, *et al.*, “Road friction estimation,” *Saab Automobile AB, Trollhättan, Sweden*, 2007.
- [10] S. Roychowdhury, M. Zhao, A. Wallin, N. Ohlsson, and M. Jonasson, “Machine learning models for road surface and friction estimation using front-camera images,” in *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, July 2018.
- [11] M. Nolte, N. Kister, and M. Maurer, “Assessment of deep convolutional neural networks for road surface classification,” in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 381–386, Nov 2018.
- [12] R. Rajamani, G. Phanomchoeng, D. Piyabongkarn, and J. Y. Lew, “Algorithms for real-time estimation of individual wheel tire-road friction coefficients,” *IEEE/ASME Transactions on Mechatronics*, vol. 17, no. 6, pp. 1183–1195, 2012.

- [13] F. Gustafsson, “Slip-based tire-road friction estimation,” *Automatica*, vol. 33, no. 6, pp. 1087–1099, 1997.
- [14] J. Prokeš, A. Albinsson, and L. Laine, “Quantification of excitation required for accurate friction estimation,” in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2551–2558, IEEE, 2016.
- [15] C. E. Rasmussen, “Gaussian processes in machine learning,” in *Summer school on machine learning*, pp. 63–71, Springer, 2003.
- [16] J. Kocijan, R. Murray-Smith, C. E. Rasmussen, and B. Likar, “Predictive control with gaussian process models,” in *The IEEE Region 8 EUROCON 2003. Computer as a Tool.*, vol. 1, pp. 352–356, IEEE, 2003.
- [17] G. Cao, E. M.-K. Lai, and F. Alam, “Gaussian process model predictive control of an unmanned quadrotor,” *Journal of Intelligent & Robotic Systems*, vol. 88, no. 1, pp. 147–162, 2017.
- [18] S. Kamthe and M. Deisenroth, “Data-efficient reinforcement learning with probabilistic model predictive control,” in *International Conference on Artificial Intelligence and Statistics*, pp. 1701–1710, PMLR, 2018.
- [19] I. D. J. Rodriguez, U. Rosolia, A. D. Ames, and Y. Yue, “Learning unstable dynamics with one minute of data: A differentiation-based gaussian process approach,” *arXiv preprint arXiv:2103.04548*, 2021.
- [20] L. Hewing, J. Kabzan, and M. N. Zeilinger, “Cautious model predictive control using gaussian process regression,” *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2736–2743, 2019.
- [21] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, “Learning-based model predictive control for autonomous racing,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.
- [22] C. D. McKinnon and A. P. Schoellig, “Experience-based model selection to enable long-term, safe control for repetitive tasks under changing conditions,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2977–2984, IEEE, 2018.
- [23] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, “Robust constrained learning-based nmpe enabling reliable mobile robot path tracking,” *The International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.
- [24] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [25] J. Yi, L. Alvarez, and R. Horowitz, “Adaptive emergency braking control with underestimation of friction coefficient,” *IEEE Transactions on Control Systems Technology*, vol. 10, no. 3, pp. 381–392, 2002.
- [26] J. Kabzan, M. I. Valls, V. J. Reijgwart, H. F. Hendriks, C. Ehmke, M. Prajapat, A. Bühler, N. Gosala, M. Gupta, R. Sivanesan, *et al.*, “Amz driverless: The full autonomous racing system,” *Journal of Field Robotics*, vol. 37, no. 7, pp. 1267–1294, 2020.

- [27] P. W. Goldberg, C. K. Williams, and C. M. Bishop, “Regression with input-dependent noise: A gaussian process treatment,” in *Advances in neural information processing systems*, pp. 493–499, 1998.
- [28] M. Lázaro-Gredilla and M. K. Titsias, “Variational heteroscedastic gaussian process regression,” in *ICML*, 2011.

## Appendix A Gaussian Process Regression

A Gaussian Process (GP) is a class of random processes, in which any finite collection of random variables have a multivariate normal distribution. A GP

$$f(x) \sim \mathcal{GP}(\eta(x), k(x, x')) \quad (1)$$

is fully specified by its mean function  $\eta(x)$  and kernel function  $k(x, x')$ , the latter specifying the covariance between pairs of random variables. Since the distribution of a GP is the joint distribution of possibly infinitely many random variables, it can be interpreted as a distribution over continuous functions. GP regression is the process of taking an existing such distribution over functions (a GP prior) and update it with respect to new data to obtain a refined distribution over functions (a GP posterior). Thus, in contrast to the majority of regression methods, the aim of GP regression is not to fit a single function to data, but rather to fit a distribution of functions to data.

A GP prior is specified by selecting  $\eta(x)$  and  $k(x, x')$  in (1), based on prior knowledge about the properties of the function we wish to identify. There exists a multitude of alternative kernel functions [15]. The squared exponential, or radial basis function kernel

$$k(x_i, x_j) = \sigma_f^2 \exp \left( -\frac{1}{2l^2} (x_i - x_j)^\top (x_i - x_j) \right), \quad (2)$$

used in this project, is one example. Parameters  $\sigma_f$  and  $l$  are used to tune the smoothness and vertical variation of the prior.

The posterior is obtained by conditioning the prior on input data. Consider  $n$  input data points  $y(x)$  with uncertainty  $\sigma_y(x)$ , and  $n_\star$  test points  $x_\star$  at which we wish to predict the mean and covariance of the posterior  $f_\star(x_\star)$ . The joint distribution of input data and the predicted function values at  $x_\star$  is given by

$$\begin{bmatrix} y(x) \\ f_\star(x_\star) \end{bmatrix} \sim \mathcal{N} \left( 0, \begin{bmatrix} K(x, x) + \sigma_y(x)I & K(x, x_\star) \\ K(x_\star, x) & K(x_\star, x_\star) \end{bmatrix} \right),$$

where the matrices  $K(\cdot, \cdot)$  represent the covariance of all permutations of the input and evaluation points. Applying standard rules for the conditioning of Gaussians [15] yields the posterior predictive distribution

$$f_\star(x_\star) \sim \mathcal{GP}(\eta_\star(x_\star), \Sigma_\star) \quad \text{with} \quad (3)$$

$$\eta_\star(x_\star) = K(x_\star, x)(K(x, x) + \sigma_y(x)^2 I)^{-1} y(x), \quad (4)$$

$$\Sigma_\star = K(x_\star, x_\star) - K(x_\star, x)(K(x, x) + \sigma_y(x)^2 I)^{-1} K(x, x_\star), \quad (5)$$

where  $\Sigma_\star$  denotes the covariance matrix of the posterior distribution.

Standard GP regression builds on the assumption that all measurements have the same measurement noise, i.e., that  $\sigma_y(x) = \sigma_y$  is constant for all  $x$ . This limitation is lifted by Heteroscedastic GP regression. The extension is described e.g., in [27, 28].