Friction-Aware Speed Planning with Gaussian Process Uncertainty and Chance-Constrained Model Predictive Control

Ektor Karyotakis, Nikolce Murgovski, Derong Yang, and Mats Jonasson

***Abstract*—This paper presents a friction-aware speed-control framework that provides probabilistic braking guarantees on collision avoidance under uncertain road friction conditions. Road friction is modeled as a spatial Gaussian Process (GP) whose nonstationary variance reflects perception confidence, enabling an adaptive representation of friction ahead of the vehicle. Under a Gaussian friction assumption, we derive an analytic one-sided chance constraint that ensures stopping within a prescribed distance at a specified risk level. This chance constraint is incorporated into a spatial-domain model predictive controller (MPC) as a terminal feasibility condition on the distance beyond the prediction horizon, while safety within the horizon is enforced conservatively by constraining acceleration using GP-based lower bounds on available friction. The resulting MPC formulation is computationally efficient and transparent, which facilitates its integration into AEB and ACC functionalities. Validation through 1,000-run Monte Carlo simulations shows that the empirical hit rate (safety-requirement violations) matches the prescribed risk level. The behavior depends on how friction uncertainty is modeled: with spatially uniform uncertainty the controller is mildly conservative, while spatially varying friction uncertainty yields behavior close to the nominal risk. Additional ACC scenarios with a snow–ice transition show that hit rates depend strongly on the assumed lead-vehicle friction model (highlighting the need to model lead braking capability correctly), whereas perception quality and occlusion have relatively minor effects.**

***Index Terms*—gaussian process, chance-constrained control, friction uncertainty, model predictive control, monte carlo, probabilistic braking, longitudinal vehicle control.**

I. Introduction

It is well known that road-surface and weather variability lead to large fluctuations in road friction, which in turn significantly increases braking-related accident risk. Existing in-production driver assistance systems are not designed to handle such uncertainty [1]: automated emergency braking (AEB) typically assumes fixed high-friction, while adaptive cruise control (ACC) prioritizes comfort over safety-critical braking performance. As a result, neither system provides guarantees on whether the vehicle can stop within the available distance under adverse or rare road conditions.

Gaussian Process (GP) models are widely used to represent uncertainty within chance-constrained model predictive control (CC-MPC) formulations for autonomous driving. Most works augment a nominal dynamics model with a GP capturing residual nonlinearities, then propagate the resulting uncertainty using linearization or moment matching to obtain approximately Gaussian state predictions [2]. Chance constraints on states and inputs can be reformulated in closed form under Gaussian assumptions, following classical stochastic MPC methods such as Gaussian-CDF tightening and Boole-based approximations [3]. GPs have also been applied to environment modeling—e.g., road boundaries or obstacles—where the posterior directly informs probabilistic collision-avoidance constraints [4]. However, these approaches do not explicitly model friction as a spatially correlated random field.

Friction-related effects have also been addressed through GP-based residual learning. In [5], a GP is trained offline from a double lane-change maneuver on snow, augmenting a nominal bicycle model with dynamics corrections reflecting low-friction behavior. While effective under fixed conditions, such offline residual models cannot incorporate spatially predictive friction information or provide probabilistic guarantees under varying friction.

An alternative line of work handles dynamics uncertainty due to friction by optimizing over multiple discrete hypotheses. Contingency MPC [6] computes both a nominal trajectory and a contingency trajectory anticipating events such as a sudden transition to low friction, with both trajectories sharing the first control input to maintain preparedness for the adverse case. Similarly, [7] optimizes over several friction hypotheses, enabling aggressive behavior on straight segments while remaining conservative in curves. These scenario-based approaches, however, rely on a finite set of pre-selected friction values and cannot represent continuous, spatially varying friction uncertainty.

Friction-adaptive control has also been demonstrated in autonomous racing, where repeated laps allow systems to learn spatial friction patterns for the track surface [8]. Likewise, [9] employed an ensemble of online GPs, each corresponding to a different friction level, yielding improved tracking when the correct model is identified. Such methods rely on strong and repeated excitation—a realistic assumption in racing but seldom achievable in everyday driving, where available tire-force excitation is limited [10].

E. Karyotakis is with Vehicle Energy and Motion Control, Volvo Car AB, SE-405 31 Gothenburg, Sweden and is also with the Department of Mechanical Engineering, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden (e-mail: ektor.karyotakis@volvocars.com).

N. Murgovski is with the Department of Electrical Engineering, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden (e-mail: nikolce.murgovski@chalmers.se).

D. Yang is with Vehicle Energy and Motion Control, Volvo Car AB, SE-405 31 Gothenburg, Sweden (e-mail: derong.yang@volvocars.com).

M. Jonasson is with the Department of Mechanical Engineering, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden (e-mail: mats.jonasson@chalmers.se).

This work was supported by (a) Volvo Car AB (Volvo Cars) and (b) Fordonsstrategisk forskning och innovation (FFI) [Strategic Vehicle Research and Innovation], part of Verket för Innovationssystem (VINNOVA) [2020-05169], the Swedish Governmental Agency for Innovation Systems.

Manuscript received February 25, 2026

Road-surface classification (RSC) provides another avenue for friction assessment. LiDAR-based RSC [11] correlates reflectivity, distance, and incident angle with surface type, though performance degrades with range and automotive-grade sensors suffer from sparse point clouds beyond 50 m [12], [13]. Cameras offer longer range [14] but also exhibit distance-dependent degradation, while surface classes must still be converted into usable friction intervals for use in planning. Consequently, perception-based friction estimates remain uncertain, range-dependent, and must be handled explicitly in the control layer. This limitation is also highlighted in [15], which shows that neither RSC nor local slip-based estimation alone is sufficient for traction-adaptive planning, and that fusing the two yields substantially better performance in critical scenarios.

A related effort in [16] used inflated perception-based friction bounds in a robust NMPC scheme for ACC, which resulted in markedly conservative behavior. To the best of our knowledge, existing work has not provided closed-form probabilistic stopping guarantees under spatial friction uncertainty, which would help reduce unnecessary conservatism while retaining formal safety margins.

Despite substantial progress in GP-based residual learning, CC-MPC, multi-hypothesis planning, and perception-based surface classification, there is currently no *computationally lightweight* longitudinal controller that simultaneously (i) models friction as a spatially correlated random field, (ii) provides closed-form probabilistic guarantees on stopping feasibility, and (iii) operates consistently across both AEB and ACC scenarios. This paper addresses these limitations.

To address these limitations, the main contributions of this work are:

1) A friction-aware longitudinal speed-planning framework (Fig. 1) for both AEB and ACC with a *mode-agnostic* MPC that only needs the available stopping distance; whether this comes from a stationary obstacle or a lead vehicle is handled outside the MPC.
2) A perception-driven spatial friction uncertainty model that maps road-surface classes to friction envelopes and embeds them in a GP with distance-dependent (nonstationary) variance to capture degrading sensing confidence.
3) An analytic one-sided stopping guarantee using the work–energy relation on a GP friction field, yielding a closed-form *global* chance constraint on speed based on the remaining energy over the post-horizon *tail*.
4) A friction-aware spatial-domain MPC that combines this *global* chance constraint on speed with conservative *local* stagewise feasibility on friction-limited acceleration via the GP-derived lower friction envelope.

II. Preliminaries

A. Paper Outline

The remainder of this paper is organized as follows. Section III provides the problem statement. Section IV defines the particle dynamics. Section VI formulates an analytic chance-constraint for stopping feasibility. Section V presents the friction uncertainty modeling. Section VII formulates a CC-MPC problem that guarantees braking within a prescribed distance at a specified risk level. Section VIII presents a simple rule-based baseline controller for comparison with the CC-MPC. Section IX verifies the proposed friction-aware framework via Monte Carlo simulations under spatially uniform friction uncertainty, where we assess the vehicle's ability to reduce its speed below a specified terminal velocity at a designated position (AEB mode). Finally, Section X presents a sensitivity study and compares the controller with a baseline under a more demanding scenario: ACC operation under spatially varying friction uncertainty. The notation is gathered in Table V.

B. Assumptions

The following assumptions are adopted throughout the analysis:

a) Upstream environment information and friction envelope validity: An upstream Road & Environment Condition Manager (RCM) provides (i) road-condition information in the form of an RSC class together with corresponding spatial friction envelopes and a confidence profile (a friction prior), and (ii) situation awareness ahead (e.g., stationary obstacles or lead-vehicle state). We assume the RSC class is correctly identified and that the realized friction lies within its class-associated envelopes. Thus, we evaluate only the width and conservativeness of these envelopes (their uncertainty), ignoring other sensing errors (e.g., misclassification, out-of-range friction, traffic-detection errors). The upstream perception and data-fusion that generate the RCM outputs are outside the scope; given these outputs, we focus on GP-based friction uncertainty modeling and its integration into CC-MPC.

b) Longitudinal centerline abstraction: Friction and vehicle dynamics are modeled only along the path centerline; lateral motion and lateral dynamics are neglected, as are road slope and rolling resistance. This abstraction is sufficient for longitudinal speed planning and braking feasibility assessment on straight roads.

c) Valid space-domain: The MPC is formulated in the spatial domain and therefore assumes that the vehicle speed remains strictly positive during predictions. In practice, a small minimum speed is enforced for numerical robustness.

d) Uncertainties included and excluded: The primary source of uncertainty modeled in this work is road friction. Other uncertainties—such as sensing errors in the lead vehicle's position or velocity—are not included in the ACC formulation. Further, the lead vehicle is assumed to utilize the same friction coefficient as the ego, which is typically not the case between different vehicle and tyre types.

e) Model validity and conditional guarantees: The probabilistic guarantees depend on the friction uncertainty model accurately representing reality. Specifically, the stated risk levels assume the GP mean and uncertainty either match or conservatively bound the true friction statistics.

f) Idealized stopping model used for analysis: The stopping condition is derived using an idealized friction-limited braking model without drag or actuator dynamics. In implementation, actuator lag is compensated by reducing the available distance by a buffer, which requires reaching the

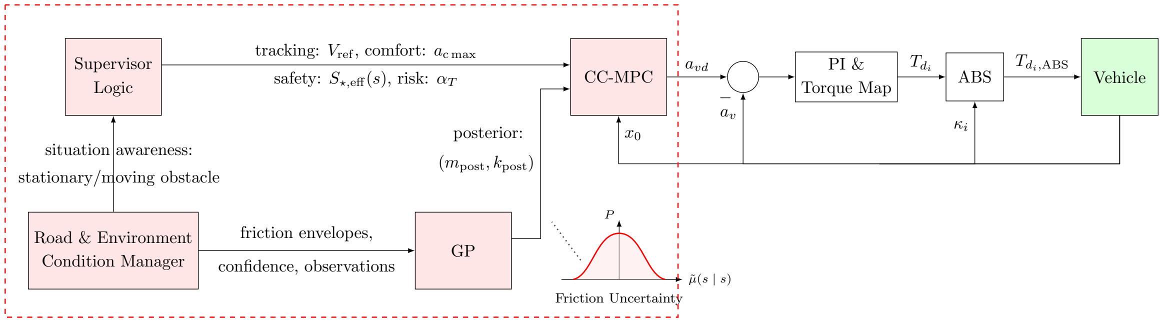

Fig. 1. Block diagram of the simulation architecture including uncertainty-aware friction Gaussian Process (GP) estimation and chance-constrained MPC (CC-MPC). The dashed red outline highlights the proposed friction-aware planning and control framework, consisting of the Road & Environment Condition Manager (RCM), GP estimator, supervisory logic, and CC-MPC.

same terminal speed over a shorter distance and therefore causes earlier deceleration than the idealized model.

g) Definition of failure event: A failure event in this work is defined differently in AEB and ACC but always arises from insufficient braking capability under the realized friction. In AEB, a failure ("Hit") occurs if the ego vehicle does not reduce its speed below the target terminal speed by the designated stopping point. In ACC, a failure occurs if the ego vehicle violates the required standstill gap to the lead vehicle at any point along the trajectory. Other failure modes (e.g., lateral instability, perception dropouts, or incorrect object tracking) are outside the scope of this study.

C. Simulation environment

Figure 1 illustrates the simulation architecture together with the proposed friction-aware framework, comprising (i) an upstream Road & Environment Condition Manager (RCM), (ii) a GP-based friction uncertainty estimator, (iii) a supervisory logic, and (iv) the CC-MPC. The remaining blocks (vehicle plant, torque conversion, PI, ABS) belong to the simulation environment and are used only to evaluate closed-loop performance.

Specifically, the RCM provides a prior description of the road friction ahead—as spatial symmetric friction envelopes and a confidence profile—which is passed to the GP-based uncertainty estimator. The resulting GP posterior is provided to the CC-MPC for planning under friction uncertainty. In parallel, the RCM supplies situation-awareness outputs such as the presence and state of stationary obstacles or a lead vehicle (e.g., relative distance and speed). Based on this information, the supervisory logic generates the CC-MPC references and requirements, including the reference speed, comfort limits, the available stopping distance, and an acceptable risk level.

The CC-MPC outputs an acceleration command, which is applied to a double-track vehicle model of a Volvo EX40 with pitch, wheel, and brake actuator dynamics and a Magic Formula Tire (MFT) model; 5-DoF body modeling details are given in [17]. To track the commanded acceleration, a torque conversion followed by a PI controller is used, together with a simple ABS logic that limits wheel torque to prevent excessive slip.

III. Problem Statement

Consider an ego vehicle traveling along the lane centerline with longitudinal position s and speed $V(s)$, and a desired reference speed $V_{\text{ref}}(s)$. Road friction ahead is uncertain and modeled as a spatial GP $\tilde{\mu}(s)$. The objective is to compute a longitudinal speed control policy that tracks V_{ref} subject to comfort and actuation limits, while guaranteeing–with a prescribed risk level–that the ego can reduce its speed to a target terminal speed $V_f > 0$ within the *available stopping distance* ahead.

Let $L > 0$ denote the available stopping distance from the current position, provided by situation awareness. In AEB, L is computed from the distance to a stationary obstacle, whereas in ACC it is computed from an assumed lead-vehicle stopping distance. Thus, the controller formulation is identical in both cases and depends only on L (i.e., the MPC is *mode-agnostic*). Given current speed V_0 and risk level $\varepsilon \in (0,1)$, the following probabilistic stopping condition is required

$$\mathbb{P}\big(V(L) \le V_f\big) \;\ge\; 1-\varepsilon, \tag{1}$$

under the friction GP $\tilde{\mu}(\cdot)$. The above requirement must be achieved without unnecessarily aggressive braking, i.e., while respecting friction-limited actuation, state/input constraints, and comfort.

At each control update, we compute a lookahead speed/acceleration profile by solving a continuous optimal control problem over a spatial horizon of length $S > 0$. Let $\rho \in [0, S]$ denote the distance ahead of the current position (so the absolute position is $s+\rho$), and let $x(\rho)$ and $u(\rho)$ denote the predicted state and control input along the horizon. The problem is posed as

$$\min_{x(\cdot),u(\cdot)} \quad \int_0^S J(x(\rho), u(\rho), V_{\text{ref}}(s+\rho))\,\mathrm{d}\rho \;+\; J_f(x(S)) \tag{2a}$$

$$\text{s.t.} \quad \frac{\mathrm{d}x}{\mathrm{d}\rho} = f(x(\rho), u(\rho))\,, \qquad \rho \in [0, S], \tag{2b}$$

$$x_{\min} \le x(\rho) \le x_{\max}, \quad u_{\min} \le u(\rho) \le u_{\max}, \tag{2c}$$

$$\text{(local feasibility)} \quad |u(\rho)| \le g\, m_{\text{lower}}(s+\rho), \tag{2d}$$

$$\text{(global feasibility)} \quad \mathbb{P}\big(V(L) \le V_f\big) \ge 1-\varepsilon, \tag{2e}$$

where $J(\cdot)$ penalizes speed tracking error, control effort, and comfort violations. Here, (2d) enforces instantaneous (stagewise) friction feasibility using the GP lower friction bound $m_{\text{lower}}(\cdot)$, while (2e) imposes a probabilistic stopping-feasibility requirement over the available distance L. In Section VI, the latter is reformulated into an analytic one-sided chance constraint via the work–energy relation under the GP friction model. We next introduce the longitudinal particle model that serves as the prediction model for the MPC.

IV. Friction-limited Particle Model at Path Centerline

The particle's longitudinal dynamics is given by (simplified from [18])

$$\begin{aligned} \dot{V}(t) &= a_v(t) - k_d V^2(t), \\ \dot{a}_v(t) &= \frac{a_{vd}(t) - a_v(t)}{\tau_{av}}, \end{aligned} \tag{3}$$

where V is the vehicle's velocity [m/s], k_d is the aerodynamic coefficient [1/m] modeling the drag force acting on the vehicle, a_v [m/s^2] is the tangential (longitudinal) component of acceleration along V, a_{vd} is the acceleration demand, and τ_{av} is the actuator time (lag) constant. The above nonlinear system has the state and input vectors $x = \begin{bmatrix} V & a_v \end{bmatrix}^\intercal,\ u = a_{vd}$.

The dynamic system (3) may also be written in compact form as

$$\dot{x}(t) = f(x, u), \tag{4}$$

where double parentheses, e.g., $f(x(t), u(t)) \to f(x, u)$, are here avoided, for didactic reasons.

A. Space domain formulation

Since friction uncertainty is defined along the road ahead, it is natural to express dynamics in the spatial domain for MPC prediction. The transformation between time and space is given by

$$S_p(t) = \frac{\mathrm{d}t}{\mathrm{d}s} = \frac{1}{V(t)}, \tag{5}$$

provided the speed remains positive, $V > 0$. Here, $s(t)$ [m] denotes the distance traveled along the path centerline. The dynamics (4) can then be expressed in the space domain as

$$\frac{\mathrm{d}x(s)}{\mathrm{d}s} = f(x, u)\, S_p. \tag{6}$$

B. Kinetic energy

Define the *kinetic energy per unit mass* $[\mathrm{m}^2/\mathrm{s}^2]$

$$K(s) \triangleq \tfrac{1}{2} V^2(s), \tag{7}$$

as a change of variable $V(s)$. This energy representation provides the basis for the stopping-feasibility analysis, which relies on a probabilistic model of the spatial friction field described next.

V. Modeling Road Friction and Uncertainty

Road-tire friction varies spatially with surface class and environmental conditions, and its estimation is subject to perception uncertainty. Because friction strongly influences braking performance, this uncertainty must be explicitly modeled in motion planning to ensure safety. We model the friction coefficient as a spatial GP informed by a perception-based RSC estimator (within the upstream RCM), which provides both a surface class (e.g., snow, wet asphalt, dry asphalt) and a measure of confidence.

The model reflects two key properties of road friction. First, friction along asphalted roads is typically continuous and spatially correlated, motivating a Matérn-kernel GP that enforces smooth, once-differentiable variations [19]. Second, each surface class is associated with characteristic friction levels and physical bounds, typically expressed as lower/upper envelopes. Because the GP is symmetric, the prior mean naturally falls at the midpoint of these envelopes, while the envelope width determines the prior variance.

A. Perception-based friction envelopes

The perception module provides a surface class that is matched to friction envelopes along the look-ahead distance. These envelopes encode the expected friction range for the detected surface class and are denoted by $\mu_{\text{envL}}(s)$ and $\mu_{\text{envU}}(s)$ (in this work obtained through a double-sigmoid fit of camera output).

From these envelopes, the GP prior is constructed as follows. The prior mean is taken as the midpoint

$$m(s) = \tfrac{1}{2}\big(\mu_{\text{envL}}(s) + \mu_{\text{envU}}(s)\big). \tag{8}$$

Define the envelope half-width

$$h(s) = \tfrac{1}{2}\big(\mu_{\text{envU}}(s) - \mu_{\text{envL}}(s)\big), \tag{9}$$

and set the prior amplitude

$$a_0(s) = \frac{h(s)}{z_{0.975}}, \tag{10}$$

so that a two-sided 95% Gaussian confidence interval matches the envelope half-width; $z_{0.975}$ is the 97.5 percentile Z-score. Here 95% is assumed as a typical algorithmic confidence of a camera-based system.

For a camera-based system, the number of pixels covering the road decreases with distance, making near-field friction estimates more reliable than far-field ones. To conservatively capture this loss of epistemic confidence, we inflate the GP variance with distance, acknowledging that the surface class may reasonably change beyond the locally observed region. This affects only *planning*; not the physical variability of friction. To reflect this, we adopt a linear confidence profile

$$c_{\text{cam}}(s) = 1 - \frac{s}{S_{\text{view}}}\big(1 - c_{\text{view}}\big) \in (0, 1], \quad s \in [0, S_{\text{view}}], \tag{11}$$

where c_{view} is the confidence at view distance S_{view}. The resulting distance-dependent amplitude is

$$a(s) = \frac{a_0(s)}{c_{\text{cam}}(s)}, \tag{12}$$

which defines the standard deviation of the prior and determines how perceived uncertainty grows with distance ahead.

B. GP Matérn friction model

Let $\tilde{\mu}(s)$ denote the uncertain friction coefficient along the path centerline, modeled as

$$\tilde{\mu}(s) \sim \mathcal{GP}\big(m(s),\, k(s,s')\big), \tag{13}$$

where $m(s)$ is the perception-informed (prior) mean from (8), and $k(s,s')$ is the covariance between friction values at positions s and s', which is based on the Matérn-$\frac{3}{2}$ kernel

$$k_0(s,s') = \left(1+\frac{\sqrt{3}}{\ell}|s-s'|\right)\exp\left(-\frac{\sqrt{3}}{\ell}|s-s'|\right). \tag{14}$$

with length-scale $\ell > 0$. To incorporate the distance-dependent uncertainty derived from perception, we scale the kernel amplitude as

$$k(s,s') = a(s)\,a(s')\,k_0(s,s'), \tag{15}$$

where $a(s)$ is given by (12). This produces a nonstationary GP whose variance increases with distance, reflecting reduced perception confidence.

a) Posterior GP: Given friction observations (e.g., current position: from wheel slip; farther ahead: cloud data) standard GP regression yields a posterior GP

$$\tilde{\mu}(s) \mid \mathcal{D} \sim \mathcal{GP}\big(m_{\text{post}}(s),\, k_{\text{post}}(s,s')\big). \tag{16}$$

In regions without friction observations—other than the perception-derived envelopes— which is the typical case ahead of the vehicle, the GP posterior reduces to the prior. In this work, we will only focus on this case; in the following sections, however, all chance constraints use the posterior functions $(m_{\text{post}}, k_{\text{post}})$ for completeness.

b) Application of GP on data: The GP regression with a stationary kernel was applied to offline real-world data acquired from a vehicle performing ABS braking on the Hällered test track, and the results are presented in Fig. 2. For the observations, the longitudinal deceleration was averaged over 5 m spatial bins during ABS activation. The prior bounds were assumed to be ± 0.1 around the expected friction level associated with the known surface class. The resulting posterior indicates that the GP yields an accurate representation of the data when the prior mean appropriately captures the underlying trend. In practice, friction observations ahead of the ego vehicle may be sparse, in which case the GP naturally reverts toward the prior, thereby producing conservative predictions in unobserved portions of the path.

c) Hyperparameters choice: The GP length-scale ℓ was set to 15 m, reflecting the spatial transition observed in braking data from the Hällered test track. The signal variance is inferred directly from the perception-derived friction envelopes, which match typical friction data ranges for the corresponding surface classes. These hyperparameters are not tuned for performance but provide a conservative and physically interpretable prior, and the MPC behavior was found to be insensitive to moderate variations in their values.

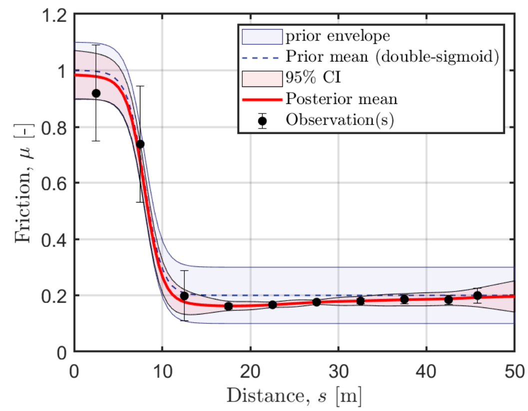

Fig. 2. GPR (Matérn-3/2 s) with double-sigmoid prior mean applied on offline braking data from Hällered test track.

VI. Chance-constrained stopping from energy relation

We now derive a probabilistic stopping condition corresponding to (1). The derivation uses the work–energy relation in the spatial domain under an idealized friction-limited braking model. For analytical tractability, we neglect actuator dynamics in (3) so that the stopping analysis depends linearly on the friction field and the resulting friction-work integral remains Gaussian under the GP model. Even though the aerodynamic drag enters linearly in the kinetic energy, it is also omitted for simpler derivations.

a) Kinetic energy and work done by friction: Neglecting aerodynamic drag and actuator dynamics in (3), and assuming friction-limited braking ($a_v(s) = -g\,\tilde{\mu}(s)$), the kinetic energy (7) satisfies

$$\frac{\mathrm{d}K}{\mathrm{d}s} = -g\,\tilde{\mu}(s). \tag{17}$$

Integrating over $[0, L]$ with $L > 0$ yields

$$K(L) = K(0) - W_f(L), \quad W_f(L) \triangleq g\int_0^L \tilde{\mu}(u)\,\mathrm{d}u, \tag{18}$$

where $W_f(L)$ is the *work done by friction per unit mass* [J/kg] accumulated over distance L. Intuitively, braking is feasible once the work done by friction exceeds the required reduction in kinetic energy.

b) Stopping requirement: Define the required kinetic-energy drop

$$K_{\text{req}} \triangleq K(0) - K_f = \tfrac{1}{2}\left(V_0^2 - V_f^2\right). \tag{19}$$

Reaching the target speed V_f by distance L is equivalent to

$$K(L) \le K_f \iff W_f(L) \ge K_{\text{req}}. \tag{20}$$

c) Gaussianity of the work done by friction: Since $\tilde{\mu}(\cdot)$ is modeled as a GP posterior $\mathcal{GP}(m_{\text{post}}, k_{\text{post}})$, the integral $\int_0^L \tilde{\mu}(u)\,\mathrm{d}u$ is a linear functional of a GP and therefore Gaussian. Hence the work done by friction $W_f(L) = g\int_0^L \tilde{\mu}(u)\,\mathrm{d}u$ is also Gaussian

$$W_f(L) \sim \mathcal{N}\big(m_W(L),\, \sigma_W^2(L)\big), \tag{21}$$

with mean [J/kg] and variance [J^2/kg^2]

$$m_W(L) = g \int_0^L m_{\text{post}}(u)\,\mathrm{d}u, \tag{22}$$

$$\sigma_W^2(L) = g^2 \int_0^L \int_0^L k_{\text{post}}(u,u')\,\mathrm{d}u\,\mathrm{d}u'. \tag{23}$$

For stationary Matérn–$\frac{3}{2}$ kernels with constant amplitude, $\sigma_W^2(L)$ admits the closed form given in Appendix A; otherwise (23) is evaluated numerically.

d) Chance-constrained stopping condition: The chance-constrained stopping requirement

$$\mathbb{P}\big(W_f(L) \geq K_{\text{req}}\big) \;\geq\; 1-\varepsilon, \tag{24}$$

is equivalent to the one-sided Gaussian quantile constraint

$$m_W(L) \;-\; z_{1-\varepsilon}\,\sigma_W(L) \;\geq\; K_{\text{req}}, \tag{25}$$

where $z_{1-\varepsilon}$ is the $(1-\varepsilon)$-quantile of the standard Normal distribution. Equivalently,

$$V_0^2 \;\leq\; V_f^2 + 2\Big(m_W(L) - z_{1-\varepsilon}\sigma_W(L)\Big) \tag{26}$$

Here, if the right-hand side becomes non-positive, the condition is infeasible for the given distance L, meaning that no admissible initial speed V_0 can satisfy $V(L) \leq V_f$ at the chosen risk level. The resulting one-sided chance constraint forms the global stopping requirement, which will be embedded as the terminal feasibility condition in the CC-MPC formulation presented next.

VII. Chance-Constrained MPC for Stopping under Friction Uncertainty

Friction uncertainty impacts braking feasibility, which is accounted for in two conditions: terminal speed and friction bound. Specifically, we enforce a terminal chance constraint that guarantees reaching a nonzero target speed within a chosen distance. At the same time, acceleration is constrained by friction, which is kept at a lower bound for conservative planning.

A. Stopping requirement distance for AEB/ACC

Let $S_\star(s)$ denote the *available* distance from the current position s to reach the target speed $V_f > 0$. This quantity is obtained from situation awareness. For AEB, $S_\star(s)$ is the remaining distance to a slow-moving or stationary obstacle. For ACC, $S_\star(s)$ corresponds to the available distance for the ego vehicle to stop safely behind a potentially hard-braking lead vehicle. A standstill gap can be added in both cases. No separate AEB/ACC terminal logic is needed inside the MPC; only $S_\star(s)$ changes.

a) AEB mode: In Autonomous Emergency Braking, a stationary obstacle at position s_{obs} yields

$$S_\star^{\text{AEB}}(s) = (s_{\text{obs}} - s) - \Delta d_{\min}, \tag{27}$$

where $\Delta d_{\min}$ is a standstill gap. The distance $S_\star(s)$ may shrink over time as the vehicle approaches the obstacle.

b) ACC mode: In Adaptive Cruise Control, the ego vehicle must guarantee that it can brake safely behind a lead vehicle that may decelerate abruptly. Let (s, V) and $(s_{\text{lead}}, V_{\text{lead}})$ denote the positions and speeds of the ego and lead vehicle, respectively, and let $\Delta d_{\min}$ be the required standstill gap.

We model braking for the lead using the same idealized work–energy relation as in (18). Let $m_{\text{lead}}(s)$ denote a conservative friction estimate for the *lead* vehicle (e.g., the GP-derived upper envelope ahead of the lead gives the worst-case lead braking). Over the braking distance S_{lead}, the lead's speed satisfies

$$\begin{aligned} V_{\text{lead}}^2(S_{lead}) &= V_{\text{lead},0}^2 - 2\,W_{f,\text{lead}}(S_{lead}), \\ W_{f,\text{lead}}(S_{lead}) &= g \int_0^{S_{lead}} m_{\text{lead}}(u)\,\mathrm{d}u, \end{aligned} \tag{28}$$

and the distance S_{lead} is obtained numerically by evaluating the cumulative integral of $m_{\text{lead}}(\cdot)$ over a spatial grid and inverting the above work–energy relation. This provides a deterministic estimate of the lead's stopping distance. When occlusion prevents friction estimation ahead of the lead, a constant bound $m_{\text{lead}}(s_{\text{lead}})$ may be used as a simplification.

The ego vehicle's available braking distance is then

$$S_\star^{\text{ACC}}(s) = \big((s_{\text{lead}} + S_{\text{lead}}) - s\big) - \Delta d_{\min}, \tag{29}$$

which is recomputed at each control update.

B. Actuator-lag distance buffer

The chance-constraint (25) is derived for idealized instantaneous friction-limited deceleration. To preserve tractability while accounting for actuation lag, we conservatively reduce the available stopping distance by a buffer ΔS_{act}, which can be chosen based on the actuator's time constant and current speed, e.g. $\Delta S_{\text{act}} = V\,\tau_{av}$. Define the *effective* stopping distance

$$S_{\star,\text{eff}}(s) \;=\; S_\star(s) - \Delta S_{\text{act}}. \tag{30}$$

C. Tail chance constraint at the end of the control horizon

Let S denote the MPC control horizon length, and let $\rho \in [0, S]$ denote the local distance along the horizon measured from the current position s. The perception system provides friction estimates over a longer look-ahead interval $[0, S_{\text{view}}]$ with coordinate ρ_v. Both ρ and ρ_v represent forward distance along the path, but the friction posterior $(m_{\text{post}}, k_{\text{post}})$ is defined on the perception coordinate ρ_v and evaluated at absolute positions $s + \rho_v$.

Define the *tail interval*

$$\mathcal{T}(s) = [S,\; S_{\star,\text{eff}}(s)], \qquad S_{\star,\text{eff}}(s) > S, \tag{31}$$

i.e., the segment beyond the MPC horizon over which stopping feasibility must still be guaranteed. Let V_S denote the predicted speed at the end of the horizon. Applying the chance-constrained stopping condition (25) to the tail interval yields the terminal condition

$$V_S^2(s) \;\leq\; V_f^2 + 2\Big(m_{W,\text{tail}}(s) - z_{1-\alpha_T}\,\sigma_{W,\text{tail}}(s)\Big), \tag{32}$$

where $\alpha_T \in (0,1)$ is the chosen *terminal risk level*, and $m_{W,\text{tail}}$ and $\sigma_{W,\text{tail}}$ are the mean and standard deviation of the friction work accumulated over the tail

$$m_{W,\text{tail}}(s) = g \int_{\mathcal{T}(s)} m_{\text{post}}(s+\rho_v)\, d\rho_v, \tag{33}$$

$$\sigma^2_{W,\text{tail}}(s) = g^2 \int_{\mathcal{T}(s)} \int_{\mathcal{T}(s)} k_{\text{post}}(s+\rho_v,\, s+\rho_v')\, d\rho_v\, d\rho_v'. \tag{34}$$

For stationary Matérn–$\frac{3}{2}$ kernels with constant amplitude, the closed-form expression in (47) may be evaluated with the corresponding tail length; otherwise (34) is computed numerically.

Constraint (32) ensures that, with probability at least $1-\alpha_T$, the friction work available from the end of the MPC horizon to the stopping point exceeds the remaining kinetic-energy reduction required to reach V_f.

D. Discrete MPC formulation

The spatial horizon $[s, s+S]$ is discretized into N intervals of length $\delta s = S/N$ with nodes $\rho_i = i\,\delta s$, $i = 0, \ldots, N$. At node $s+\rho_i$, the MPC state and input are $x_i = [V_i \;\; a_{v,i}]^\top$ and $u_i = a_{vd,i}$. The dynamics are discretized using RK4 over δs, denoted by $x_{i+1} = F(x_i, u_i)$.

1) Acceleration/Braking Comfort: Comfort is handled by a blended constraint that allows the controller to exceed the comfort limit $a_{\text{c max}}$ only when the slack $\epsilon_{a,i}$ is activated

$$|a_{v,i}| \;\leq\; a_{\text{c max}} \;+\; \Big(m_{\text{lower},i}\, g - a_{\text{c max}}\Big)\, \epsilon_{a,i}, \tag{35}$$

where $a_{\text{c max}} = 2\,\text{m/s}^2$, and $0 \leq \epsilon_{a,i} \leq 1$; the slack is penalized in the cost. Friction feasibility is enforced through a hard bound on the commanded acceleration

$$|u_i| \;\leq\; g\, m_{\text{lower},i}, \tag{36}$$

using the GP *lower* bound

$$m_{\text{lower}}(\cdot) = m_{\text{post}}(\cdot) - z_{1-\alpha_T}\sqrt{k_{\text{post}}(\cdot,\cdot)}, \tag{37}$$

for conservative planning. The corresponding *upper* bound, used for the lead vehicle's braking, is defined symmetrically as

$$m_{\text{upper}}(\cdot) = m_{\text{post}}(\cdot) + z_{1-\alpha_T}\sqrt{k_{\text{post}}(\cdot,\cdot)}. \tag{38}$$

Here, the same risk level α_T is used as in the tail chance constraint (32) for consistency.

2) Terminal chance constraint: When an effective tail interval exists, i.e., $S_{\star,\text{eff}}(s) > S$, terminal feasibility is enforced via the tail chance constraint (32). If the required stopping point lies within the prediction horizon, $S_{\star,\text{eff}}(s) \leq S$, no tail segment remains and we instead impose the deterministic in-horizon stopping condition $V(\rho^\star) \leq V_f$ with $\rho^\star = S_{\star,\text{eff}}(s)$ (implemented at the nearest grid node at $\rho^\star$).

When the remaining distance is too short (e.g. $S_{\star,\text{eff}} <$ 10 m), the controller switches to deterministic full braking $u_i = -g\, m_{\text{lower},i}$ until the target speed V_f is reached to avoid infeasible MPC solutions. Although one could instead employ the upper friction bound to intentionally exceed the friction limit and trigger ABS for potentially higher deceleration, this strategy is not pursued in this work.

3) Speed constraints: The speed is bounded between a lower speed $V_{\min} > 0$ to avoid the space-domain singularity and a maximum speed $V_{\max,i} > V_{\text{ref},i}$ that may be dynamic

$$V_{\min} \leq V_i \leq V_{\max,i}. \tag{39}$$

4) Cost and MPC formulation: The cost function balances reference speed tracking, control effort, and comfort. Tracking performance is weighted by W_V, control effort by W_U, and comfort violations by W_a through the slack variables $\epsilon_{a,i}$. A jerk (regularization) term with weight W_j promotes smooth control variations along the horizon. The discrete MPC formulation is given by

$$\begin{aligned}
\min_{\{x_i, u_i, \epsilon_{a,i}\}} \quad & J = W_V (V_N - V_{\text{ref},N})^2\, \delta s \\
& + \sum_{i=0}^{N-1} \Big[W_V (V_i - V_{\text{ref},i})^2 + W_a\, \epsilon_{a,i}^2 + W_U\, u_i^2 \Big] \delta s \\
& + \sum_{i=0}^{N-2} W_j\, (u_{i+1} - u_i)^2 / \delta s && \text{(40a)} \\
\text{s.t} \quad & x_{i+1} = F(x_i, u_i), && \text{(40b)} \\
& V_{\min} \leq V_i \leq V_{\max,i}, && \text{(40c)} \\
& |a_{v,i}| \leq a_{\text{c max}} + \Big(m_{\text{lower},i}\, g - a_{\text{c max}}\Big) \epsilon_{a,i}, && \text{(40d)} \\
& |u_i| \leq g\, m_{\text{lower},i}, && \text{(40e)} \\
& V_N^2 \leq V_f^2 + 2\Big(m_{W,\text{tail}} - z_{1-\alpha_T} \sigma_{W,\text{tail}}\Big) && \text{(40f)} \\
& x_0 = x_{\text{init}}. && \text{(40g)}
\end{aligned}$$

VIII. Baseline controller in ACC mode

This section presents a computationally-simpler controller that corresponds to today's ACC in emergencies. The baseline controller is constructed under the assumption of constant friction for both the lead and ego vehicles, which reduces (29)–(30) to the condition $s_{\text{lead}} - s \geq \Delta s_{\min}$, with

$$\Delta s_{\min} = \frac{1}{2g}\left(\frac{V^2(s)}{\bar{m}_{lower}(s)} - \frac{V^2_{lead}(s_{lead})}{m_{lead}(s_{lead})} \right) + V(s)\tau_{av} + \Delta d_{\min}. \tag{41}$$

This equation is similar to the Responsibility-Sensitive Safety (RSS) formulation in [20], with the distinction that the abstract braking levels of the ego and lead vehicles are replaced by friction bounds. In this context, $\bar{m}_{\text{lower}}(s)$ represents the spatial average of the lower friction bound defined in (37) over the interval $[s, s_{\text{lead}}]$ for the ego vehicle. Furthermore, it is assumed that, possibly due to occlusion, the lead vehicle continues to travel under the same friction level as currently perceived at s_{lead}, denoted by $m_{\text{lead}}(s_{\text{lead}})$. This value is taken either as the upper friction bound (38) (yielding a worst-case scenario) or as the posterior mean (more optimistic).

Define the relative distance and speed errors

$$e_{\Delta s} = \Delta s - \Delta s_{\min}, \quad \Delta s = s_{\text{lead}} - s \tag{42}$$

$$\dot{e}_{\Delta s} = V_{\text{lead}} - V. \tag{43}$$

A simple proportional-derivative controller is chosen as the baseline, which regulates relative distance through the rule

$$a_{vd}^{rss}(s) = \begin{cases} \max\left[k_1 e_{\Delta s} + k_2 \dot{e}_{\Delta s},\ -\bar{m}_{\text{lower}}(s)g\right], \\ \qquad\qquad \text{if } \Delta s < \Delta s_{\min}; \\ 0, \qquad\quad \text{otherwise}; \end{cases} \tag{44}$$

where $k_1,\ k_2$ are tuning parameters.

IX. Verification using Monte Carlo in sppatially uniform friction uncertainty

In this section, we verify the proposed friction-aware framework through Monte Carlo (MC) simulations under a braking scenario with spatially uniform friction uncertainty. Because friction is *constant along distance* in this test case, the terminal chance constraint reduces to a *scalar* quantile condition parameterized by the risk level α_T. In each simulation, a random friction curve is sampled while the MPC plans using the corresponding $(1-\alpha_T)$ lower-quantile friction bound (scalar here). The goal is to reach the target speed within the given distance; failure occurs when the realized friction falls below the assumed bound. We report the *Hit Rate*, the fraction of simulations in which the vehicle fails to reach the target speed. Over many MC trials, the Hit Rate is expected to approach the specified risk level α_T, up to finite-sample variation, control–plant and actuator-model mismatches, as well as the influence of the MPC horizon and cost terms.

We use 1,000 Monte Carlo runs, which provides sufficient statistical resolution for validations (empirical std for 10% risk level $\sigma_{\text{MC}} = \sqrt{\frac{p(1-p)}{N}} \approx \pm 0.9\%$). The scenario conditions are: AEB mode, specified braking distance $S_\star = 30\,\text{m}$, horizon distance $S = 10\,\text{m}$ with step $\delta s = 0.5$ m, reference speed $V_{\text{ref}} = 70\,\text{km/h}$, initial speed $V_0 = 60\,\text{km/h}$, and a damp road with $\mu_{\text{mean}} = 0.8$, $\mu_{\text{lower}} = 0.6$, $\mu_{\text{upper}} = 1$. Braking begins at 50 m to allow initial conditions to settle, and the vehicle targets $V_f = 5\,\text{km/h}$ at 80 m, i.e., 30 m ahead. Weights are set as: $W_V = 1$, $W_U = 0.25$, $W_a = 10$, and $W_j = 10\,\delta s^2$. Further, the distance buffer in (30) is set to zero so that the Hit Rate includes the effect of ignoring the actuator and vehicle body reaction times.

The simulations are implemented in MATLAB using CasADi [21] for nonlinear optimization and IPOPT [22] as the solver, and run on an Intel Core Ultra 7 165U processor (1.70 GHz) and 32 GB of RAM, with the average solve time per MPC iteration in 20–30 ms.

Table I reports Monte Carlo results showing that the prescribed risk level aligns with the Hit Rate, with some conservatism, under the examined MPC settings. Additional conservatism can be achieved by (1) increasing the prediction horizon, (2) reducing the weight of comfort costs, and (3) enlarging the actuator distance buffer in (30). These mechanisms are discussed in detail in Section XI.

The table also lists the average cruising speed prior to initiating braking, defined from here onward as the pre-braking speed. Its magnitude reflects both the assumed friction bound and the cost weights used in the MPC objective. For the displayed friction realization, the difference in pre-braking speed between the 1% and 10% risk settings is approximately

TABLE I
Monte Carlo results for a damp asphalt braking scenario

Risk α_T	Hit Rate	Pre-braking Speed [km/h]	Friction Bound [-]	Impact Speed [km/h]
1%	0.7%	57.5	0.56	9.9
5%	3.2%	59.5	0.63	13.1
10%	6.2%	60.5	0.67	13.7
50%	35.3%	64.5	0.8	16.1
99%	94%	69	1.04	22.4

3 km/h. Finally, the average impact speed for the hits is included (reported at the stopping point), which relates directly to crash severity and can be mapped to published injury-risk curves by interested readers.

For reference, estimating the stopping risk via Monte Carlo by sampling realizations of the same stochastic friction $\tilde{\mu}$ and checking the analytic stopping event $W_f(L) \geq K_{\text{req}}$ shows that 70 km/h corresponds to an approximate 3% stopping risk over 30 m for the considered uncertainty profile. This reflects the idealized friction-limited constant-deceleration assumption underlying the work–energy model. In contrast, finite-horizon MPC with comfort penalties and applied to an advanced vehicle plant model slightly reduces the pre-braking speed, yielding mildly conservative Hit Rates.

An example of a failed case to reach the target speed is presented in Fig. 3. The vehicle starts to brake at the specified distance of 30 m before the target, but due to out-of-bound true friction it cannot reach the intended speed. While braking, the vehicle initially remains on the safe speed envelope (red dashed), but when the friction assumption is violated, the current speed (blue line) surpasses this envelope. The speed before braking is determined from the predictions (yellow lines). At each control iteration, the speed at the end of the horizon should reach the terminal speed (red dash-dotted line). The braking profile is shaped by the comfort- and control effort-related costs, which influence both the pre-braking speed and the aggressiveness of the braking phase.

X. Validation case in ACC mode and spatially varying friction uncertainty

In this scenario, we model a transition from snow to ice in ACC mode. The lead vehicle brakes at a predetermined position using the true friction limit, while the ego vehicle brakes according to the conservative lower friction bound prescribed by the MPC design. First, we make a sensitivity study of three factors that affect the controller's performance and evaluate the Hit Rate. Second, we compare the MPC to a simpler rule-based controller. In this ACC setting, a Hit is defined when the relative distance falls below the required standstill gap at any time.

The simulation conditions are: 1,000 Monte Carlo runs, ACC mode, horizon length $S = 10\,\text{m}$, reference speed $V_{\text{ref}} = 70\,\text{km/h}$, ego initial speed $V_0 = 70\,\text{km/h}$, initial ego–lead distance 50 m, and lead speed $V_{\text{lead}} = 60\,\text{km/h}$. Road friction follows a snow–ice transition with a mean profile $\mu_{\text{mean}}(s) : 0.4 \to 0.2$ and symmetric bounds of ± 0.1 around the mean. As before, the distance buffer in (30) is set to zero so

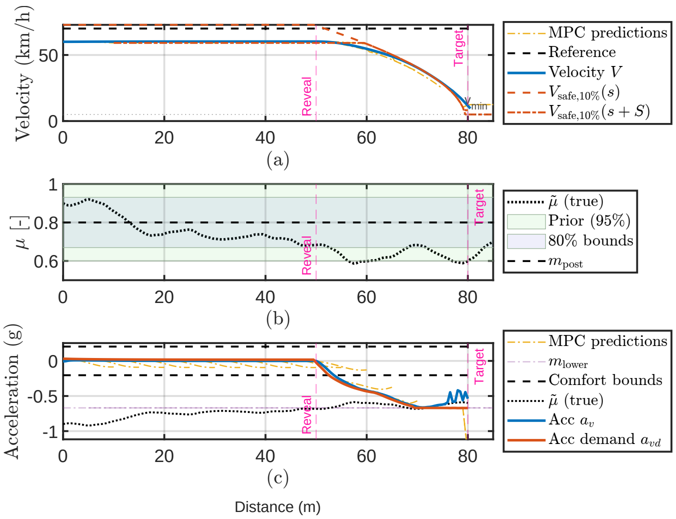

Fig. 3. Simulation of a vehicle attempting to brake within a specified distance under uncertain friction. In this case, the vehicle fails to reach the target speed due to friction falling below the assumed bound. (a) Velocity profile and speed limits; (b) Friction profile along the route; (c) Longitudinal acceleration. MPC predictions are plotted every 10 m; risk level is 10%.

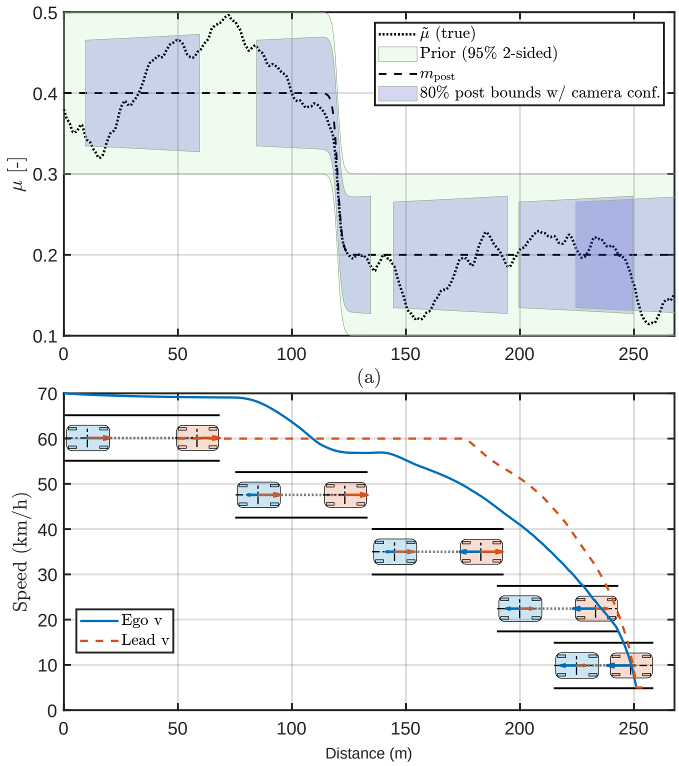

Fig. 4. Simulation example in ACC mode. Camera-based friction envelopes and vehicle states plotted at five positions along the route. (a) True friction bounds together with the GP posterior envelopes and the camera-based confidence profile (shown over a 50 m segment for clarity). (b) Ego (blue) and lead (red) vehicle states at the same five positions. Vehicle outlines are scaled for visualization and do not represent actual size. Blue arrows indicate longitudinal force vectors and red arrows denote velocity vectors

that the observed Hit Rate reflects purely the effect of friction uncertainty (actuator and body dynamics are intentionally not compensated here). The ego risk level is set to 10%. The lead vehicle follows the particle model time dynamics (3), begins braking at $175\,\text{m}$, and stops at $V_f = 5\,\text{km/h}$. A simulation example of a successful stop is presented in Fig. 4.

A. Sensitivity

We consider three categorical factors: camera confidence, occlusion, and the assumed lead-vehicle friction bound, yielding $2^3 = 8$ scenarios. In this context, occlusion is defined as the situation in which friction information is available only up to the position of the lead vehicle; consequently, the most recently observed friction value is assumed to remain valid for the road segment ahead.

Table II reports the Monte Carlo Hit Rates, grouped by camera confidence, occlusion, and the assumed lead braking friction. Under the posterior-mean (here the mean-profile due to no observations) friction assumption for the lead vehicle, the MPC yields a high Hit Rate. This happens because assuming the mean smoothens out occasional low-friction patches, such patches, however, can greatly increase stopping distance and are often what causes the rare but critical braking shortfalls. The braking guarantee is thus conditional and applies only when the assumed friction bound is valid or more conservative. In contrast, when the lead braking friction is modeled using the conservative GP upper friction bound, the MPC Hit Rate drops close (but slightly above) the intended 10% risk level. Adding a small actuator distance buffer, the Hit Rate remains below the prescribed 10% level, as evaluated in the next subsection. Concluding, spatially-varying friction uncertainty increases the risk and makes the controller less conservative compared to spatially-uniform friction uncertainty.

The pre-braking relative distance (reported when the lead starts braking) and the relative impact speed (reported at the first violation of the standstill gap) are also reported and illustrate the conservativeness and risk of the controller. Larger distances before initiating braking generally correspond to lower Hit Rates and larger impact speeds.

Although not designed under a specific jerk limit, for this examined case, the controller achieves an RMS jerk about 1 m/s^3 and the highest 95th percentile lies around 3 m/s^3 which fulfills industry standards for an ACC.

To highlight qualitative trends, Table III summarizes the directional effect of the three categorical factors on Hit Rate. Overall, the dominant factor affecting safety is the assumed lead-vehicle friction; using a conservative (upper) bound is essential for maintaining the highest probability of avoiding accidents.

B. Comparison to baseline controller

The baseline parameters $[k_1, k_2] = [15, 0.2]$ were tuned on a high-friction scenario to achieve a Hit Rate ≈ 0, using $\mu = 1 \pm 0.1$, as $\mu = 1$ is the typical friction assumption in current industry-standard AEB functionality. With this tuning the controller is expected to be safe but brake aggressively on low-friction surfaces, provided friction is correctly estimated and conservative.

For this comparison, the risk is set to 10%, the camera confidence to 100%, occlusion is imposed so that the same point estimate of the lead-vehicle state is used by both controllers,

TABLE II
HIT RATE FROM 1000 MONTE CARLO SIMULATIONS IN ACC MODE UNDER A SNOW–ICE TRANSITION SCENARIO, GROUPED BY THREE FACTORS: CAMERA CONFIDENCE, OCCLUSION, AND LEAD VEHICLE BRAKING LIMIT ASSUMPTION

Occlusion	m_{lead} assumed bound	Hit Rate	Pre-braking relative distance [m]	Relative impact speed [km/h]
Camera confidence $= 100\%$				
No	Upper	10.8%	33.6	8.1
No	Mean	89.7%	28.7	16
Yes	Upper	11.2%	34.1	8
Yes	Mean	90.2%	29.1	16
Camera confidence at $S_{\text{view}} = 60\%$				
No	Upper	4.1%	35.2	6.3
No	Mean	80.4%	29.3	14.3
Yes	Upper	4.9%	35.3	6.7
Yes	Mean	81%	29.8	14.4

TABLE III
QUALITATIVE EFFECT OF SCENARIO FACTORS ON HIT RATE (SNOW–ICE TRANSITION)

Factor	**Hit Rate Effect**	**Interpretation**
Occlusion	↑	Less information about friction ahead; slight increase in Hit Rate.
Lead friction assumption	(Mean) ↑↑ (Upper) ↓	Optimistic assumption; underestimates lead stopping ability; greatly increases Hit Rate. Conservative assumption; keeps Hit Rate close to the risk level.
Camera confidence decreases with distance	↓	GP variance grows ahead; MPC becomes more conservative; reduction in Hit Rate.

and the lead vehicle braking is modeled at its upper friction bound. Further, the actuator distance buffer ($V\tau_{av}$) is included both for the baseline and the MPC for completeness. Under these assumptions, the main difference between the controllers lies in their adaptability to risk under different conditions: the baseline needs to be tuned for every surface and risk level, while the CC-MPC incorporates it into its design.

The results for both controllers are presented in Table IV. The baseline controller is markedly conservative in Hit Rates although keeping longer pre-braking distances. At the same time, the addition of the distance buffer reduces Hit Rates by 3.5% (compared to the 3rd case in Table II) while other metrics changed slightly. More importantly, the outcome is highly sensitive to tuning: for example, increasing the risk level from 10% to 20% altered the Hit Rate from 0.6% to 6%, illustrating that the baseline controller does not regulate risk in a predictable way. Nevertheless, the baseline controller demonstrates the benefit of explicitly modeling friction in ACC design and serves as a useful sanity check, whereas the CC-MPC provides a principled, systematically tunable approach for regulating braking risk under spatially varying friction.

TABLE IV
MPC VS. BASELINE RSS-BASED CONTROLLER COMPARISON IN ACC MODE. SCENARIO FACTORS: CAMERA CONFIDENCE 100%, OCCLUSION, $m_{\text{lead}} = m_{\text{upper}}$, RISK LEVEL 10%.

Controller	Hit Rate	Pre-braking relative distance [m]	Relative impact speed [km/h]
CC-MPC	7.7%	34.4	7.7
baseline	0.6%	36.0	0.1

XI. DISCUSSION

The proposed friction-aware framework (Fig. 1) provides a structured way to reason about friction uncertainty in longitudinal control. Key insights follow:

a) Adaptive risk tuning in AEB: Adjusting the specified braking distance and risk level in AEB mode naturally yields different speed profiles adapted to upcoming friction and traffic context. For example, around school areas one would impose short braking distances and near-zero risk, whereas on a country road a larger stopping distance and a nonzero risk level may be acceptable.

b) MPC cost influences braking behavior: Control-effort and comfort penalties consistently make the controller more conservative, leading to earlier and milder braking and lower cruising speeds. The comfort-slack penalty influences the timing of braking more subtly: with a low weight, braking begins earlier and smoothly, whereas a high weight delays braking but makes it more abrupt. The prediction horizon also shapes closed-loop behavior via preview. Longer horizons let the optimizer anticipate the terminal (tail) feasibility requirement and spread deceleration over the horizon, typically lowering cruising speed and improving robustness. Shorter horizons shift more braking to the tail segment, and therefore, tend to increase the Hit Rate. Together, cost structure and horizon length tune the trade-off between smoothness, comfort, and safety.

c) Extension to lateral motion: Although this work focuses on friction variation along the path centerline, the framework can be extended to lateral maneuvers by augmenting the vehicle model to two spatial dimensions (s, e) as in [18] and defining a GP over $\mathbb{R}^2$ (e.g., a Matérn kernel), yielding a spatially varying $\tilde{\mu}(s, e)$ suitable for combined braking and steering.

This study is limited to:

d) Failure-mode coverage: Only longitudinal infeasibility is quantified; lateral failure modes are not accounted for, except indirectly through speed reduction before curves (cf. [16]).

e) Perception-to-friction mapping: The coupling between the RSC output and the friction envelopes used in this work is based on offline exploration and has not yet been validated in real-world driving. Nevertheless, the proposed approach is generic: it remains applicable to any friction-range specification and any sensor-confidence model, provided that these can be represented as spatial envelopes for the GP prior.

f) Conditional probabilistic guarantees: As with any chance-constrained scheme, the risk bounds hold only under the assumed GP prior and noise model. Unmodeled perception errors (e.g., misclassification) or friction distributions outside the prior are not explicitly handled.

g) Computational aspects: While the MPC solve times are lightweight for the scenarios tested, further work is needed to benchmark the framework under hardware-embedded implementations.

h) Closed-loop friction estimation: The framework assumes no friction observations ahead of the vehicle. Incorporating online friction updates in closed loop, and propagating their effect through the chance constraints, is left for future work.

XII. Conclusion

This paper presented a friction-aware longitudinal speed-planning framework for AEB and ACC that regulates stopping feasibility under uncertain road friction at a prescribed risk level. Perception-derived road-surface information is mapped to spatial friction envelopes and embedded in a spatial Gaussian Process model, while stopping feasibility is expressed as a one-sided chance constraint derived from the work–energy relation and integrated into a spatial-domain chance-constrained model predictive controller through a terminal tail constraint.

The framework was validated in Monte Carlo simulations by comparing the empirical Hit Rate to the prescribed risk level. Under spatially uniform friction uncertainty, the Hit Rate closely matched the intended risk level, with mild conservatism attributable to finite-horizon effects and comfort/effort penalties. Under spatially varying uncertainty, Hit Rates increased indicating less conservative behavior which can be mediated in several ways as by adding a distance buffer.

In ACC scenarios with a snow–ice transition, several qualitative effects were examined and the dominant factor affecting safety was the assumed braking capability of the lead vehicle: using a conservative (upper) friction bound yielded Hit Rates close to the specified risk, whereas a mean-friction assumption was overly optimistic and led to frequent violations. Reduced camera confidence increased conservatism and lowered Hit Rates by tightening the GP-based friction bounds, while occlusion had a comparatively minor effect in the examined scenarios.

We conclude that the proposed global chance-constraint on the kinetic energy–work relation combined with local GP lower bound stagewise braking offers a transparent approach for friction-aware longitudinal control, and can serve as a foundation for future extensions including steering maneuvers and closed-loop friction estimation.

Appendix

For stationary Matérn-$\frac{3}{2}$, set $k(s,s') = \kappa(|s-s'|)$, $\phi = |s-s'|$, (14) becomes

$$\kappa(\phi) = \sigma^2(1+\iota\phi)\exp(-\iota\phi), \qquad \iota = \frac{\sqrt{3}}{\ell} \tag{45}$$

For the work done by friction $W_f(L) = g\int_0^L \tilde{\mu}(u)\,\mathrm{d}u$, the variance is

$$\sigma_W^2(L) \;=\; g^2\, 2\int_0^L (L-\phi)\,\kappa(\phi)\,\mathrm{d}\phi. \tag{46}$$

Carrying out the convolution yields the closed form

$$\sigma_W^2(L) \;=\; g^2\,\frac{2\sigma^2}{\iota^2}\Big[2\iota L - 3 + (\iota L + 3)e^{-\iota L}\Big]. \tag{47}$$

References

[1] T. Rahman, A. Liu, D. Cheema, V. Chirila, and D. Charlebois, "Adas reliability against weather conditions: Quantification of performance robustness," in *27th International Technical Conference on the Enhanced Safety of Vehicles (ESV) National Highway Traffic Safety Administration*, no. 23-0306, 2023.

[2] 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.

[3] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear model predictive control with chance constraints – a review," *Journal of Process Control*, vol. 44, pp. 53–67, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0959152416300130

[4] A. D. Bonzanini, A. Mesbah, and S. Di Cairano, "Perception-aware chance-constrained model predictive control for uncertain environments," in *2021 American Control Conference (ACC)*, 2021, pp. 2082–2087.

[5] R. R. Arany, H. Van der Auweraer, and T. D. Son, "Learning control for autonomous driving on slippery snowy road conditions," in *2020 IEEE Conference on Control Technology and Applications (CCTA)*, 2020, pp. 312–317.

[6] J. P. Alsterda, M. Brown, and J. C. Gerdes, "Contingency model predictive control for automated vehicles," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 717–722.

[7] R. K. Aggarwal and J. C. Gerdes, "Friction-robust autonomous racing using trajectory optimization over multiple models," *IEEE Open Journal of Control Systems*, vol. 5, pp. 16–30, 2026.

[8] L. Hermansdorfer, J. Betz, and M. Lienkamp, "A concept for estimation and prediction of the tire-road friction potential for an autonomous racecar," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 1490–1495.

[9] T. Nagy, A. Amine, T. X. Nghiem, U. Rosolia, Z. Zang, and R. Mangharam, "Ensemble gaussian processes for adaptive autonomous driving on multi-friction surfaces," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 494–500, 2023.

[10] M. Acosta, S. Kanarachos, and M. Blundell, "Road friction virtual sensing: A review of estimation techniques with emphasis on low excitation approaches," *Applied Sciences*, vol. 7, p. 1230, 11 2017. [Online]. Available: http://www.mdpi.com/2076-3417/7/12/1230

[11] H. Hongyu, T. Minghong, G. Fei, B. Mingxi, and G. Zhenhai, "Road surface friction estimation based on lidar reflectivity for intelligent vehicle," *IEEE Transactions on Instrumentation and Measurement*, vol. 74, pp. 1–12, 2025.

[12] D. Yang and V. Bogren, "Lidar-based road weather condition detection for increased traffic safety," FFI, Tech. Rep., 2024.

[13] V. Le and P. Johansson, "Lidar-based road state estimation," Ph.D. dissertation, Chalmers University of Technology, 2023.

[14] H. Karunasekera and J. Sjöberg, "In search for better road surface condition estimation - using non-road image region," *IEEE Intelligent Vehicles Symposium, Proceedings*, pp. 3205–3211, 2024.

[15] L. Svensson and M. Törngren, "Fusion of heterogeneous friction estimates for traction adaptive motion planning and control," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 424–431.

[16] M. Crnic and S. Koutsoftas, "Road friction aware adaptive cruise control using robust nonlinear model predictive control with uncertainty quantification," Master's thesis, Chalmers University of Technology, 2025.

[17] L. Nielsen and B. Olofsson, *Optimal Vehicle Maneuvers*. Wiley-IEEE Press, 2026, ISBN: 9781394428168.

[18] E. Karyotakis, "Braking distance minimization on roads with varying friction," Licentiate Thesis, Chalmers University of Technology, 2025.

[19] M. L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*, ser. Springer Series in Statistics. New York: Springer, 1999.

[20] T. Yu, Y. Tang, R. Chen, and S. Zhao, "Optimization of adaptive cruise control strategies based on the responsibility-sensitive safety model," *Vehicles*, vol. 7, no. 2, p. 28, 2025.

[21] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, pp. 1–36, 2019.

[22] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, pp. 25–57, 2006.

EKTOR KARYOTAKIS received his Diploma of mechanical engineering from Aristotle University of Thessaloniki, Greece, in 2016 and his M.Sc. in automotive engineering from Chalmers University of Technology, Gothenburg, Sweden, in 2019.

From 2019 to 2021, he was a research assistant at Chalmers University. Since 2021, he has been with Volvo Cars and Chalmers University as an Industrial Ph.D. student. His research focuses on enhancing vehicle motion planning through the use of predictive road friction information.

NIKOLCE MURGOVSKI received his M.S. degree in software engineering from University West, Trollhättan, Sweden, in 2007, the M.S. degree in applied physics and the Ph.D. degree in systems and control from the Chalmers University of Technology, Gothenburg, Sweden, in 2007 and 2012, respectively.

He is currently a Professor with the Department of Electrical Engineering at Chalmers University of Technology. His research interests include modeling and simulation of dynamical systems and their use in optimal control with topics in electromobility, autonomous driving, and active safety in automobiles.

DERONG YANG received her M.Sc. in automotive engineering and her Ph.D. degree in vehicle dynamics and control from Chalmers University of Technology, Gothenburg, Sweden, in 2009 and 2013, respectively.

She then joined Volvo Cars, where she is currently a Technical Expert in vehicle motion control. Her research interests include developing motion control functions for both manual and autonomous vehicles, focusing on optimizing safety, comfort, and energy efficiency.

MATS JONASSON received his M.Sc. in electrical and electronic engineering from Chalmers University of Technology, Gothenburg, in 1997 and the Ph.D. degree from the KTH Royal Institute of Technology, Stockholm, Sweden, in 2009.

From 2009 to 2016, he worked at Volvo Cars in motion control. He is currently a Senior Researcher in the Division of Vehicle Engineering and Autonomous Systems at Chalmers University of Technology. His research interests include vehicle motion control, electrification of vehicles, and automated driving functions.

TABLE V
NOTATION USED THROUGHOUT THE PAPER

Symbol	**Unit**	**Meaning**
Dynamics and spatial domain		
s	m	Position along lane centerline
ρ	m	Local spatial coordinate along MPC horizon
$V(s)$	m/s	Ego speed at position s
V_0	m/s	Initial ego speed
V_f	m/s	Target terminal speed
$V_{\text{ref}}(s)$	m/s	Reference speed
a_v	m/s^2	Longitudinal acceleration state
a_{vd}	m/s^2	Acceleration demand (MPC input)
τ_{av}	s	Actuator lag time constant
k_d	1/m	Aerodynamic drag coefficient
S_p	s/m	Time–space conversion factor
$x = [V \;\; a_v]^\top$	m/s, m/s^2	State vector
$u = a_{vd}$	m/s^2	Control input
Energy and stopping feasibility		
$K = \frac{1}{2}V^2$	J/kg	Kinetic energy per unit mass
K_{req}	J/kg	Required energy drop to reach V_f
$W_f(L)$	J/kg	Friction work over distance L
$m_W(L)$, $\sigma_W(L)$	–	Mean and Std. of $W_f(L)$
g	m/s^2	Gravitational acceleration
Stopping distances		
L	m	Generic stopping distance (problem statement)
$S_\star(s)$	m	MPC available stopping distance
Δd_{min}	m	Standstill gap
ΔS_{act}	m	Actuator distance buffer
$S_{\star,\text{eff}}$	m	Effective stopping distance
S	m	MPC horizon length
$\mathcal{T}(s)$	m	Tail interval for terminal constraint
GP friction model		
$\tilde{\mu}(s)$	–	Random friction coefficient
$\mu_{\text{envL}}, \mu_{\text{envU}}$	–	Perception-derived envelopes
$m(s)$	–	GP prior mean
$a_0(s)$, $a(s)$	–	GP prior amplitude / distance-scaled amplitude
$c_{\text{cam}}(s)$	–	Camera confidence profile
$k(s, s')$	–	GP covariance (Matérn-$3/2$)
ℓ	m	GP length-scale
m_{post}, k_{post}	–	Posterior mean/covariance
m_{lower}	–	GP lower friction bound
m_{upper}	–	GP upper friction bound
Risk levels		
ε	–	Generic stopping-risk requirement
α_T	–	MPC risk parameter (experiments use $\alpha_T = \varepsilon$)
$z_{1-\alpha}$	–	Standard Normal $(1-\alpha)$-quantile
Discrete MPC		
N	–	Number of spatial intervals
δs	m	Spatial step, $\delta s = S/N$
x_i, u_i	–	State and input at node i
$\epsilon_{a,i}$	–	Comfort constraint slack
J	–	objective function value
W_V, W_U, W_a, W_j	–	MPC cost weights
V_S, V_N	m/s	Terminal speed at end of horizon (continuous: V_S; discrete MPC: V_N)