Transportation Research Interdisciplinary Perspectives 39 (2026) 102141

Contents lists available at ScienceDirect

Transportation Research Interdisciplinary Perspectives

journal homepage: www.elsevier.com/locate/trip

Navigating uncertainty: Safety solutions for autonomous vehicle integration into mixed-mode mobility[☆]

J.R.E. Cooke [a,*], S.N. Lingam [b,f,*], P. Alavian [c], C. Collast [d], G.C.H.E. de Croon [e], M. Franssen [f], J. Hagenaars [e], J. Kober [e], P. Morsink [g], K. Liu [a], Y-H. Liu [h], E. Matthews [i], J. Mejias [c], A. Palffy [j], C.M.A. Pennartz [c], W. Sanberg [k], L. Selen [a], T. Sun [l,o], N. van Nes [m], R. Verstegen [f], M. Martens [f,n], J. Alsono-Mora [e], W.P. Medendorp [a], S.M. Bohte [c,o]

[a] Radboud University, Nijmegen, The Netherlands
[b] Royal Netherlands Aerospace Centre (NLR), Amsterdam, The Netherlands
[c] University of Amsterdam, The Netherlands
[d] VVN, Amersfoort, The Netherlands
[e] Delft University of Technology, Delft, The Netherlands
[f] Eindhoven University of Technology, Eindhoven, The Netherlands
[g] RDW, Zoetermeer, The Netherlands
[h] Holst Centre, IMEC, Eindhoven, The Netherlands
[i] 2getthere, Utrecht, The Netherlands
[j] Perciv.ai, Delft, The Netherlands
[k] NXP, Eindhoven, The Netherlands
[l] Maastricht University, Maastricht, The Netherlands
[m] SWOV, Den Haag, The Netherlands
[n] TNO, Helmond, The Netherlands
[o] CWI, Amsterdam, The Netherlands

ARTICLE INFO

Keywords:
Uncertainty
Autonomous vehicles
Risk
Neuroscience
Human–automation interaction

ABSTRACT

Uncertainty fundamentally shapes human perception and decision-making, a factor that is becoming increasingly critical as autonomous vehicles (AVs) begin to share physical and social spaces with humans. This perspective paper synthesizes insights from neuroscience, robotics, and behavioral science represent uncertainty and act under it. We identify a significant disconnect between low-level mathematical concepts of uncertainty and risk, the uncertainty-aware prediction and planning methods used in robotics, and the high-level psychological uncertainty experienced by humans. Here, we conceptualize an integration pathway for mixed-mode mobility: we show that all three treat control under partial information as the same two-stage problem – probabilistic *inference* of the current and future state, followed by a *risk*-weighted choice of action – and we make this shared structure explicit. We thus argue that AVs should quantify and calibrate their own uncertainty, select risk metrics that reflect both individual and collective safety, model how their actions shape human uncertainty and behavior, and communicate intent and confidence in ways that support predictable interaction. This synthesis positions psychological uncertainty as a behaviorally relevant variable for AV design and outlines a research agenda for safer, more understandable transport systems.

☆ This article is part of a Special issue entitled: 'AI in Transportation' published in Transportation Research Interdisciplinary Perspectives.

* Corresponding authors.
E-mail addresses: james.cooke@donders.ru.nl (J.R.E. Cooke), Nischal.Lingam@nlr.nl (S.N. Lingam), s.p.alavian@uva.nl (P. Alavian), c.collast@vvn.nl (C. Collast), G.C.H.E.deCroon@tudelft.nl (G.C.H.E. de Croon), m.franssen@tue.nl (M. Franssen), j.j.hagenaars@tudelft.nl (J. Hagenaars), J.Kober@tudelft.nl (J. Kober), PMorsink@rdw.nl (P. Morsink), joyu.liu@donders.ru.nl (K. Liu), Yao-Hong.Liu@imec.nl (Y.-H. Liu), emi@2getthere.eu (E. Matthews), j.f.mejias@uva.nl (J. Mejias), andras.palffy@perciv.ai (A. Palffy), C.M.A.Pennartz@uva.nl (C.M.A. Pennartz), willem.sanberg@nxp.com (W. Sanberg), luc.selen@donders.ru.nl (L. Selen), tao.sun@cwi.nl (T. Sun), nicole.van.nes@swov.nl (N. van Nes), r.verstegen@tue.nl (R. Verstegen), m.h.martens@tue.nl (M. Martens), J.AlonsoMora@tudelft.nl (J. Alsono-Mora), pieter.medendorp@donders.ru.nl (W.P. Medendorp), s.m.bohte@cwi.nl (S.M. Bohte).

https://doi.org/10.1016/j.trip.2026.102141
Received 14 January 2026; Received in revised form 26 June 2026; Accepted 6 July 2026
Available online 17 July 2026

1. Introduction

Uncertainty pervades virtually every aspect of human experience, thereby fundamentally shaping how we perceive, decide, and act in our daily lives (Alchian, 1950; Busemeyer and Townsend, 1993; Pezzulo and Friston, 2019). Whether we are crossing a busy intersection, interpreting another person's intentions, or making financial investments, our behavior is constantly modulated by varying degrees of uncertainty. This universal human experience of uncertainty has taken on new significance as we enter an era where autonomous vehicles (AVs) and systems increasingly share our physical and social spaces (Rasouli and Tsotsos, 2019; Schwarting et al., 2019).

Consider a pedestrian approaching a crosswalk as an autonomous vehicle nears the intersection. The pedestrian's decision-making process involves multiple layers of uncertainty: Will the vehicle stop? Has it detected my presence? What happens when I cross anyway? These questions highlight a critical challenge in the deployment of autonomous vehicles: understanding and managing not just the technical uncertainties inherent in the systems themselves, but also the human assessment of uncertainty about their behavior (Dragan et al., 2013). This also applies vice versa: to interact, the autonomous vehicle has to determine the set of likely human behaviors in a given context, up to the point of modifying its own behavior in order to influence human behavior and reduce human uncertainty (Sadigh et al., 2016).

This issue extends beyond private autonomous cars as mixed-mode mobility includes buses, shuttles, delivery robots, drones, cyclists, pedestrians, and human-driven vehicles operating in shared infrastructure. In such transport systems, uncertainty shapes crossing behavior, yielding, merging, route choice, acceptance, trust, and the perceived safety of public space. A transportation perspective therefore requires both objective uncertainty management by automated systems and attention to the subjective uncertainty experienced by occupants, bystanders, and other road users.

This perspective thus addresses a fundamental question at the intersection of human behavior and autonomous vehicles: what conceptual and design requirements follow when AVs are treated as agents whose uncertain actions are interpreted by uncertain humans? Here, we offer an interdisciplinary synthesis that identifies shared computational structure, clarifies where current models do not yet meet, and outlines how uncertainty-aware AV design could better account for human behavior. We outline the respective treatments of uncertainty in neuroscience, robotics, and the behavioral science of human–automation interaction, show that they share a common computational backbone, and synthesize where the disciplines fail to connect.

We first examine the foundations of uncertainty, distinguishing between risk and uncertainty (Knight, 1921), and explore how these concepts manifest in both perception – the processing of sensory information into meaningful and coherent understandings of the world, and action – the execution of goal-directed movement or sequences of movements. In classical decision theory, risk refers to decision situations in which outcomes are uncertain but probabilities and losses can be specified; in real traffic, those probabilities are themselves estimated, learned, and socially influenced. We therefore treat risk as a decision-oriented transformation of uncertainty: a way of mapping uncertain outcomes to action-relevant costs.

We investigate uncertainty from multiple perspectives: from how the human brain processes uncertainty, how current autonomous vehicles handle uncertain situations, and how human assessment of uncertainty influences the behavior of autonomous vehicles and of humans and also their interactions. We integrate these insights around the computational challenges that arise when modeling complex uncertain environments, as modeling almost any relevant environment quickly becomes computationally intractable (Kaelbling et al., 1998), and any autonomous agent by necessity has to make simplifying assumptions to arrive at decidable constructs.

We find that despite significant advances in each of these fields, an important gap remains in our understanding: as we argue, low-level mathematical concepts of uncertainty connect poorly to high-level behavioral models that are both computationally tractable and capture human decision-making, even though they solve similar problems. This paper presents our view on and progress in addressing this challenge, with the ultimate goal of developing more effective and intuitive autonomous vehicles that can seamlessly interact with humans in environments full of uncertainties (Simon, 1955).

Inference, then risk. Across neuroscience, robotics, and behavioral psychology, acting under partial information can be described in two stages: *inference* followed by a *risk-weighted decision*. An agent maintains a probabilistic model of the world, combines prior knowledge with incoming evidence to form a posterior belief about the current state, and propagates this belief into a distribution over the future consequences of candidate actions. A *risk functional* then collapses this distribution of uncertain outcomes onto a single preference, ranging from cautious (risk-averse) to assertive (risk-seeking). In this view, uncertainty becomes a control variable that calibrates behavior. We use this two-stage lens where inference leads to risk to compare otherwise disparate literatures and, in the synthesis (Section 5), pinpoint where they diverge.

Contribution and scope. The contribution of this perspective is integrative: we (i) make explicit the shared inference-then-risk structure underlying uncertainty in the three fields; (ii) use it to compare how each field measures uncertainty and converts it into action; (iii) identify that mechanistic neuroscientific models and formal robotic risk models do not yet connect to the subjective uncertainty that shapes human responses; and (iv) propose metacognition as a concrete research pathway for closing it (Section 5.2). As such, this perspective suggests new directions for operational integration.

2. Uncertainty from a neuroscience perspective

Uncertainty is referenced in a wealth of neuroscience literature, spanning from studies on low-level visual processing (Knill and Pouget, 2004) to motor control (van Beers et al., 2004) and high-level perception and cognitive functions (Griffiths et al., 2024). The concept of uncertainty and the role it plays in the understanding of brain and behavior is most effectively articulated through a combination of information theory (Cover and Thomas, 2006) and Bayesian decision theory (Wolpert and Landy, 2012). Information theory allows us to formalize uncertainty through probabilistic methods and Bayesian decision theory provides a method to account for this uncertainty to make optimal decisions. Optimal decisions are derived from a posterior distribution, which combines prior knowledge, existing evidence, and the potential outcomes associated with various actions (Wolpert and Landy, 2012).

For example, (Fig. 1A) to drive a car effectively as humans, we must integrate signals provided by our noisy and delayed visual, vestibular, somatosensory and proprioceptive systems (Fig. 1B) to estimate our speed relative to the other vehicles on the road (Fig. 1C). Additionally, we must assess our task performance in terms of risk – such as overtaking another car while avoiding collisions – to determine the appropriate motor responses, like pressing the accelerator to accelerate and turning the steering wheel (Fig. 1C) (Nash and Cole, 2020; Liu et al., 2025). Very similar issues emerge when controlling a flying contraption, like an airplane or helicopter.

2.1. Sources of uncertainty

Uncertainty arises at the lowest level of neural processing itself, which is endowed with intrinsic random noise (see Faisal et al., 2008 for a review). For instance, the human visual perception is affected by noise generated during the transduction process at the level of retinal cells. This entails that even basic tasks related to visual orientation, position, and motion perception are complicated by the noise inherent in neural signals. Likewise, our motor system is subject to noise not only in the central planning of a movement (van Beers et al., 2004) but also during its execution. During the latter, the motor system exhibits significant signal-dependent noise; specifically, the variability in force production tends to increase with the desired force (Jones et al., 2002).

In addition to the noise introduced at the various stages of neural processing, there are also substantial transmission delays from the peripheral sensory organs to the brain and between different brain areas (Franklin and Wolpert, 2011), which can extend up to 100 ms for visual processing (Schlag and Schlag-Rey, 2002). Such delays contribute significantly to uncertainty in time-varying systems, as it hampers our ability to track rapidly varying sensory signals (Liu et al., 2025).

Finally, a large source of uncertainty reflects the inherent dynamic world in which we live. Most behaviors require navigating ambiguous selection requirements, as outcomes are uncertain, and probabilities are unknown.

In the real world, the statistics of sensory inputs can change continuously depending on the actions we take. It is crucial that we can distinguish between a change in sensory input related to our own movements and changes in the world around us. For example, when driving, visual signals from our surroundings passing by us due to self-motion induces optic flow (Brucklacher et al., 2025). However, changes in optic flow could also occur due to visual objects that move on their own (e.g., a car passing by, or pedestrians walking). Some of them are more important than others, e.g., objects on a potential collision or interception course. This holds for both humans and autonomous vehicles alike. Hence, a crucial challenge is to disambiguate which parts of the optic flow are due to self-motion and which parts come from other objects in the world. Mathematically, this is an ill-posed problem, which inflicts substantial ambiguity in how optic flow information should be integrated into perception (Noel et al., 2023). Notably, separating self-motion from the independent motion of other agents is the same disambiguation problem that an autonomous vehicle's perception and prediction stack must solve (Section 3.2), and is both formalized in terms of probabilistic inference and optimal control.

2.2. Accounting for uncertainty

Bayesian decision theory serves as the foundational framework for a dominant computational approach in sensorimotor control, known as optimal feedback control (OFC) theory (Todorov and Jordan, 2002; Todorov, 2005; Crevecoeur et al., 2011). OFC employs Bayesian principles to control a dynamic system, such as the human body, using time-varying feedback controllers to minimize the expected cost over time. Typically, this cost contains terms penalizing motor effort (e.g. torque) and task error, encouraging solutions which are low effort but still lead to good task performance. More formally, OFC combines sensory input and motor commands in a Bayesian manner to provide precise state estimates and applies a controller to minimize overall cost to this state estimate.

Both Bayesian decision theory and OFC theory generally adopt a model-based approach to facilitate optimal decision-making for perception and action amid various uncertainties. Is there evidence of their application by humans? For example, one prediction is that human and animal perceptual uncertainty can be reduced by combining inputs from multiple sensory systems and prior information learned over time. Indeed, by now, multiple studies have demonstrated that humans and animals combine different sensory modalities to reduce perceptual uncertainty (e.g. weighting by their relative precisions) (Ernst and Banks, 2002; Fetsch et al., 2010; van Beers et al., 1999). Additionally, humans utilize prior knowledge across various tasks, from visual perception (Knill and Richards, 1996) to sensorimotor control (Körding and Wolpert, 2004; Liu et al., 2025). Importantly while Bayesian integration is effective in reducing the effects of sensory noise, addressing motor noise requires computations by the controller as part of the Bayesian integration process. Methods from stochastic optimal control provide ways to compute control policies which are optimal even in the presence of signal-dependent noise (Todorov and Jordan, 2002; Moore et al., 1999). Accounting for signal-dependent noise helps to explain many kinematic properties of human movements (Harris and Wolpert, 1998; Todorov, 2005; Braun et al., 2009).

The effect of noise is exacerbated in the presence of sensory and sensorimotor delays. How does the brain compensate for these delays? Since we cannot eradicate these delays, the best strategy is to account for them in the state estimation processes. To this end, most computational approaches rely on a predictive model which takes into account the delays associated with the human sensorimotor system (Kasuga et al., 2022; Crevecoeur et al., 2011; Todorov and Jordan, 2002) and uses this model to optimally account for them.

Because the world around us is dynamic, the sensorimotor system must be able to learn, update and switch between different control and estimation policies when necessary. Bayesian causal inference presents a method to account for model ambiguity by extending inference to also include the causal structure. There is substantial evidence that the brain performs such computations for multisensory integration, allowing it to flexibly combine sensory information based on the likelihood of a common source (Körding et al., 2007; de Winkel et al., 2018; Acerbi et al., 2018; Rohe et al., 2019; von Hünerbein et al., 2026).

However, this causal inference aspect has received less attention in biological control theory, where traditional approaches often presume a known environmental and body model (Todorov, 2005; Kasuga et al., 2022; Crevecoeur et al., 2011) (but see Braun et al., 2009; Izawa et al., 2008). When the parameters or the model of the system are uncertain, this creates an adaptive control problem (Stengel, 1994), in which observers must both accomplish the task at hand while simultaneously estimating the unknown parameters or model. This presents a fundamental trade-off for learning systems in a complex world: should the agent focus on model learning or performing the task? A specific illustration is seen in active learning, where actions are deliberately generated to acquire information about the environment. For example, humans plan their saccades (eye-movements) to maximize information gathering under the constraints of the sensorimotor system (Yang et al., 2016). Conversely, individuals may choose to focus on task performance even in the face of uncertainty, employing strategies (e.g. increasing limb stiffness) to improve task performance at the cost of an increase in effort (Mitrovic et al., 2010; Berret et al., 2024). Currently, we do not know of an overall framework which integrates all these factors into a single computational solution.

2.3. Risk

Just as Bayesian decision theory allows us to define uncertainty, we can leverage it to provide derived quantitative definitions of concepts such as *risk*. Intuitively, the notion of *risk* describes the decision-making response: the cost, trade-offs, or safety implications of choosing one action versus another given those uncertainties. Risk is about consequences. Typically Bayesian decision theory (Wolpert and Landy, 2012) defines risk as the expected loss over outcomes, and thus a Bayes optimal decision maker selects an action to minimize this. This is precisely the same problem as addressed by the optimal feedback control framework (Todorov and Jordan, 2002; Todorov, 2005): the optimal action under OFC is one that minimizes risk in the statistical decision sense. As we already discussed, there is extensive evidence that humans select actions that take cost and/or risk into account (Todorov and Jordan, 2002; Wolpert and Landy, 2012). In some cases, neuroscientists use the meaning of risk in a slightly different way, inspired by work in control theory (Whittle, 1981), using risk to refer to a controller that is sensitive to the variability in the loss as well as the expected value (Nagengast et al., 2010). This work has demonstrated that observers employ risk-sensitive and risk-seeking strategies that

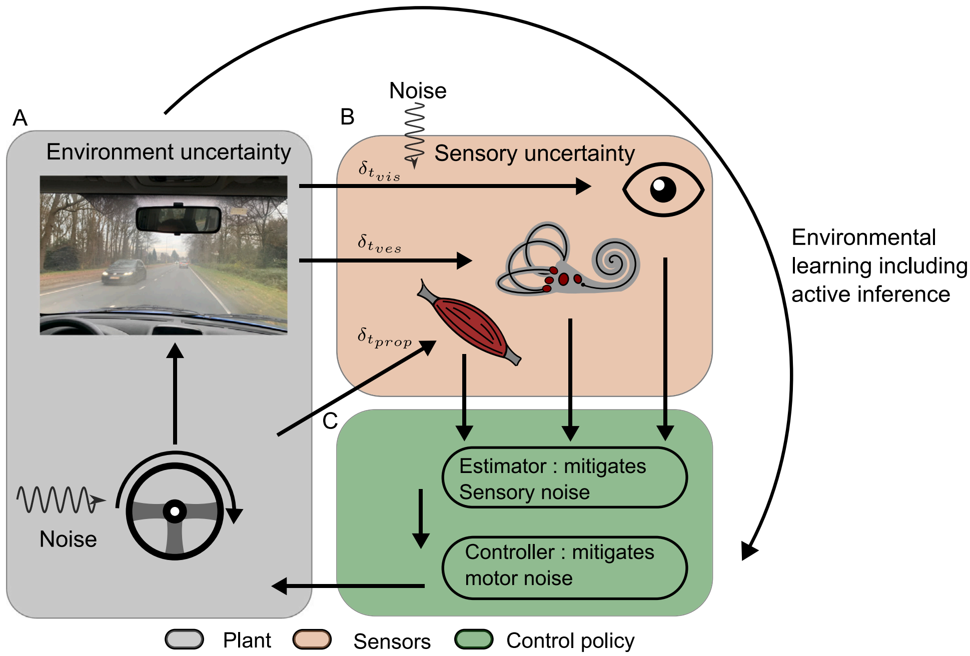

Fig. 1. Overview of the sources of uncertainty and overall computational processes during driving.

vary between individuals and thus take both the expected loss but also the variance in loss into account to plan actions. Notably, the same distinction appears in AV planning (Section 3.3), where the chosen risk metric determines whether behavior is cautious, neutral, or aggressive under uncertainty; the neuroscience and robotics notions of risk thus coincide, differing mainly in whether the metric is inferred from human behavior or chosen as a design parameter.

2.4. Neural mechanisms

How are the computations outlined above realized in neural circuits? The brain is generally argued to have plausible neural mechanisms for each step of the inference-then-risk pipeline. We give some details here and refer to Haefner et al. (2024) for more in depth considerations. Broadly, three ingredients recur in various viewpoints. *First*, sensory measurements must be encoded in a way that preserves their uncertainty: low-level features such as orientation are represented by the joint activity of neural populations (Tanabe, 2013), and exactly how such population activity encodes uncertainty – through probabilistic population codes (Ma et al., 2006), doubly distributional codes (Sahani and Dayan, 2003), or neural sampling (Orbán et al., 2016; Echeveste et al., 2020) (see Haefner et al., 2024) – remains debated, though decoding studies confirm that early sensory areas carry uncertainty about their inputs (Walker et al., 2020; van Bergen et al., 2015). There has also been significant work demonstrating that artificially constructed neural networks (NN) can accomplish a variety of behavioral tasks (Mante et al., 2013; Song et al., 2016; Wierda et al., 2025). *Second*, this information must be combined across modalities for state estimation, a function associated with parietal cortex (Medendorp and Heed, 2019; Scott, 2004; von Hünerbein et al., 2026); notably, artificial neural networks trained on such tasks develop usable uncertainty representations without being given explicit probabilistic machinery (Orhan and Ma, 2017; Echeveste et al., 2020). *Third*, the resulting estimate must be turned into a motor command, plausibly by motor cortex acting as a general feedback controller that integrates visual (Stavisky et al., 2017) and proprioceptive (Pruszynski et al., 2011) feedback (Cross et al., 2024; Scott, 2004), consistent with models of goal-directed reaching (Kalidindi et al., 2021; Berniker and Penny, 2019). Finally, the brain uses mechanisms to account for and utilize uncertainty within individual cells, synapses and small circuits, where it can enable useful dynamical regimes and improve signal detection through stochastic resonance (Douglass et al., 1993; Longtin, 1993; Torres et al., 2011).

2.5. Neural frameworks for high-level perception and planning under uncertainty

Beyond low-level sensorimotor control, neuroscience has proposed several computational frameworks to explain how the brain plans complex, goal-directed behaviors in the face of uncertainty. These models conceptualize decision-making as a form of probabilistic inference, where the brain actively maintains and updates beliefs about the world to guide actions.

Extending on the Bayesian notions of low-level perception and control, Bayesian Reinforcement Learning (BRL), formalizes the problem of learning and acting in an unknown environment by explicitly representing uncertainty about environmental parameters (e.g., transition probabilities and rewards) using probability distributions. This allows for a principled resolution of the exploration–exploitation trade-off: actions are chosen not only to maximize immediate reward but also to reduce uncertainty and gather information that can lead to better long-term outcomes (Ghavamzadeh et al., 2015). The neural arbitration between an exploratory, goal-directed system and a more rigid, habitual one is thought to be managed by an uncertainty-based competition between the PFC and the dorsolateral striatum, respectively (Daw et al., 2005), that also involves sensorimotor cortical areas, ventral striatum and medial caudate (Rusu and Pennartz, 2020; Pezzulo et al., 2014).

The theory of *predictive coding* unifies perception, action, and cognition (Lee and Mumford, 2003; Friston, 2005; Rao and Ballard, 1999) into a single framework where the brain actively predicts upcoming sensory stimuli. Predictive coding can be formalized as an approximation to variational Bayesian inference, particularly under assumptions of Gaussian distributions over hidden causes and prediction errors (Friston, 2005). Then, the brain minimizes variational free energy to estimate the posterior distribution over the hidden causes of sensory input. In this view, the brain maintains an internal generative model of the world, continuously generating predictions about incoming inputs, comparing them with the actual sensory data, and using the resulting prediction errors to improve both inference (updating) and learning (adjusting weights in the model itself).

Within the notion of Predictive Coding, the Active Inference framework suggests that for acting, the brain operates as a generative model that constantly tries to minimize the discrepancy between its predictions and sensory inputs (i.e., minimize "free energy") (Parr et al., 2022). Planning then involves evaluating policies, or sequences of actions, based on their expected free energy. This expectation balances extrinsic value (reaching desired goals) with an intrinsic, information-seeking drive to reduce uncertainty about future states. The balance between deliberative, goal-directed planning and automatic, habitual responses is thought to be controlled by the precision of beliefs about policies. High precision favors deliberation, a process likely involving the dorsolateral PFC (dlPFC) and ventromedial PFC (vmPFC), while low precision allows for the dominance of habitual actions mediated by the basal ganglia and related areas (Pezzulo et al., 2014, 2018; Rusu and Pennartz, 2020).

In both approaches, it remains an active topic of research how "desired goals" are defined and mechanistically arrived at, and how low-level mechanistic neural implementations can be scaled up to such high-level computations.

Taken together, neuroscience treats uncertainty as a principled feature of perception and action: it arises from noisy, delayed sensorymotor channels and a changing, ambiguously caused world; it is typically quantified with Bayesian posteriors; and it is managed through multisensory integration, causal inference, and stochastic optimal feedback control, with neural population codes and predictive coding being promising candidate implementations. Generally, the brain is believed to maintain and update state beliefs, weigh information by its precision, plan actions that trade off performance and information gain, and adapt controllers as models change (Pearson et al., 2021). Next, we show how these same ingredients are operationalized in robotics, in particular in autonomous vehicles, via probabilistic perception, sensor fusion, Bayesian filtering, and risk-aware planning and control that explicitly reasons about and acts under uncertainty.

3. Uncertainty from a robotics perspective

With the aim of achieving a desired high-level outcome, notions of uncertainty in robotics typically concern high-level aspects of the static environment and the plans and potential reactions of other actors. For a car-driver to turn "left" on an intersection the presence and potential interference of other traffic participants has to be determined. Uncertainty enters for example from visibility, which may be limited, the state-of-mind (attentive/distracted) of other road users, and the anticipated adherence to traffic rules. Humans routinely plan over long horizons despite uncertainty about others' intentions and the environment; AVs must perform an analogous task using explicit prediction, planning, and control algorithms.

Despite decades of progress in optimal decision-making (operations research, numerical optimization, and control), planning under uncertainty remains challenging in robotics. Below, we give an informal account of motion prediction and planning for autonomous driving, emphasizing how *posterior* uncertainty and *risk metrics* interact to determine behavior. These considerations mostly carry over to autonomous vehicles in general, like drones.

3.0.0.1. Motivating example. Consider a scenario of an autonomous vehicle (AV) approaching an intersection alongside two human-driven vehicles as shown in Fig. 2. In this scenario, even including the presence or absence of signaling, each human driver exhibits one of two possible intent choices; the vehicle approaching from the left may either turn left or continue straight, while the vehicle moving downward can either yield to the AV or turn left. The evolution of these interactions introduces two sources of uncertainty:

1. Intent ambiguity which is known as *modes* depicted in blue and yellow for each human-driven vehicle.
2. Given these modes, the state of each vehicle (e.g. position and orientation) is subject to uncertainty. These uncertainties are visualized at each time step by the colored ellipsoids surrounding the predicted trajectories.

3.1. Challenges of motion planning for autonomous robots

One of the main challenges of motion planning under uncertainty for AVs stems from the computational intractability of the underlying mathematical programs against the infinitely many possible future outcomes. Imagine extending the previous example to an AV planning to navigate a Shibuya crossing in Tokyo; Ideally, the vehicle should first *predict* the intended trajectories of all the surrounding cars and pedestrians, then *plan* contingent collision-free trajectories based on those predictions. Such a scenario poses a huge computational burden from a numerical optimization point of view for both prediction and planning. As the number of agents and their modes increase in the scene, the number of decision variables of the optimization problem increases exponentially along the prediction time horizon. This might take the AV hours to find an optimal solution (if found) considering all possible modes, resulting in the solution becoming outdated.

Moreover, dealing with the two sources of uncertainties in the motivating example raises the following questions: How do we model the distributions of those uncertainties? What happens when those distributions change based on different pedestrian behaviors and driving laws from one region to another? Answering those questions is a model design choice that will be addressed in Section 3.2.

Importantly, planning under uncertainty requires satisfying a set of constraints along with a desired objective according to some *risk metrics* associated with achieving the task. Consider the previous driving scenario. The AV can exhibit different decision-making behaviors based on how it assesses risks. This is generally referred to as either *Risk-seeking*, where the AV may assume the most likely scenario—that the human driver will yield and proceed through the intersection without adjusting its speed; *Risk-averse*, where the AV may hedge against all possibilities by coming to a full stop, waiting for the drivers' intentions to become clear; and *Risk-neutral*, where the AV would neither assume the worst-case scenario (risk-averse) nor the most likely one (risk-seeking). Risk connects planning under uncertainty to actual decision-making.

Given these types of risk-mediated behaviors, an immediate question is raised: how do we assess such metrics in the presence of those dynamic uncertain environments? For instance, should the AV consider satisfying all collision scenarios if it is risk-averse? Should it consider not colliding on the average case? Notice that the design choice of a risk measure is yet another one that produces a different safety outcome, which is intertwined with the optimality of achieving a task. It is also a function of how uncertainties are modeled. This is not an easy decision, since an overestimation of future uncertainties can lead to conservative behavior that deteriorates the task objective of a robot. At the same time, a very relaxed one can hinder safety by violating the planning constraints.

Take the classic example of an AV on a highway that keeps skipping all exits on its way to the airport based on a worst-case risk measure to ensure a large separation distance to avoid colliding with the other vehicles. Although this risk measure achieves perfect safety, it fails its main task of getting to its destination on time (i.e. deteriorating the task performance) which might as well end up with the car standing still to satisfy the constraint. On the other hand imagine the complement of this scenario based on some average risk measure, where the vehicle arrives at its destination on time after making a couple of collisions with its surroundings to achieve its objective. In Section 3.3, we report some of the risk metrics currently used in motion planning in robotics to answer the above questions (Majumdar and Pavone, 2020).

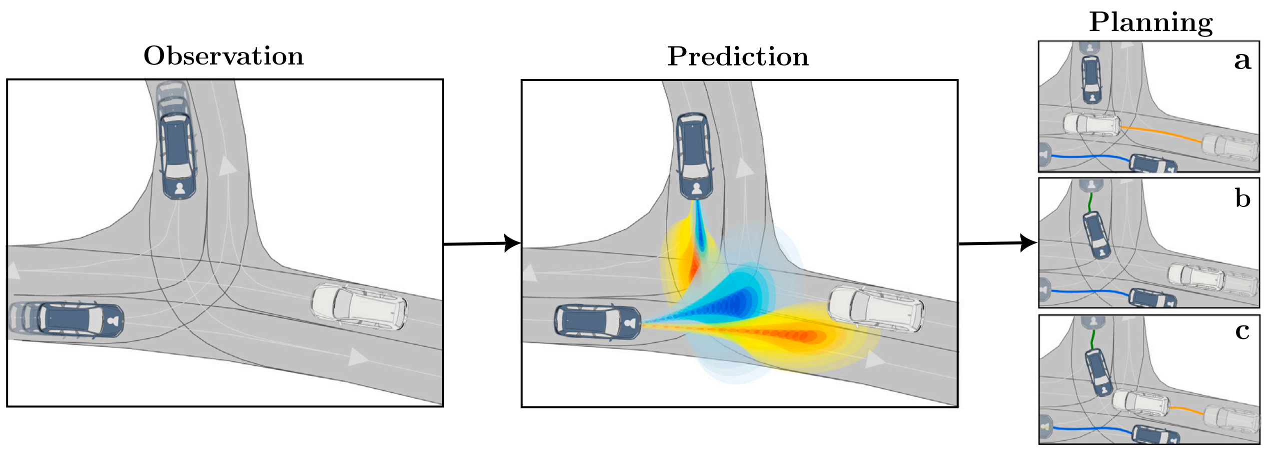

Fig. 2. Example of the prediction–planning pipeline on an intersection scenario. In the first step, the AV (white) observes two human-driven vehicles (dark blue) approaching the intersection. The AV keeps track of the other vehicles' history as well as the layout of the intersection as its *observation*. The observation is then used by the *prediction* model to predict the intents of the human driven vehicles—for simplicity we only depict two. The predicted trajectories are shown with associated uncertainty depicted with the colored ellipsoids, where intent uncertainty determines possible motion modes (shown in blue and yellow) and trajectory uncertainty reflects variation within each mode. The prediction is then used in the *planning*. Depending on whether the AV is (a) risk-seeking, (b) risk-averse, or (c) risk-neutral the resulting plans can differ from each other, with the AV speeding up in the intersection expecting the human-driven vehicle to yield for it as in the first case, being too cautious and stopping outside of the intersection in the second case, or following a strategy based on the expected outcome that may result in collision as in the final case.

3.2. Motion prediction under uncertainty

Operating in dynamic human environments is a challenging problem for AVs. One fundamental challenge arises from the inherent uncertainty in human behavior, otherwise known as aleatoric uncertainty. Aleatoric uncertainty refers to the intrinsic randomness of a process, such as the variability in human actions due to different cognitive states or external influences (Ferro and Sornette, 2020). This type of uncertainty cannot be reduced with more data, making it particularly important for AVs to accurately model the aleatoric uncertainty. Take the previous example of the driving scenario in Fig. 2; relying on merely deterministic predictions or having inaccurate uncertainties can easily lead to catastrophic safety outcomes. The main task of prediction models is to inform the motion planner about the uncertain future outcomes of the surrounding agents. In this paradigm, prediction models use information about the past of the surrounding agents and often the static environment as well to condition their predictions on a given situation. The predictions aim to capture the future states of the agents in the scene which can be modeled as random variables, with their variability depicted as ellipsoids as shown in Fig. 2. One can think of the variability of the agents' states coming from slight variations in velocity and steering in similar situations. Accurately predicting such uncertainty can in theory lead to better decisions for the motion planner of the AV.

Technically, a random variable is an uncertain quantity defined on a probability space consisting of: (i) a sample space of all possible outcomes, (ii) a collection of events, and (iii) a probability rule that assigns each event a number between 0 and 1. In the context of our motivating example, this random variable captures the uncertainty associated with the discrete modes and the continuous trajectory variations given the mode.

Prediction Models. There are two types of methods for capturing the aleatoric uncertainty over people's future behavior—methods that provide an implicit uncertainty over the predicted samples, and methods that provide these uncertainties explicitly. While implicit uncertainty methods are common in trajectory prediction (Gupta et al., 2018; Kawasaki and Seki, 2021; Wang et al., 2022; Niedoba et al., 2023), they only generate samples of possible futures. A majority of uncertainty-aware planners, however, need access to the probability of those futures to evaluate how likely a given outcome may be and plan accordingly. To avoid extracting this information from samples generated by implicit methods, one can directly learn models that explicitly provide it. Explicit uncertainty methods strive to provide either the likelihood of individual samples or a closed-form description of the full probability distribution.

A simple way of capturing aleatoric uncertainty in this manner is to only focus on the high-level decisions (Zhi et al., 2021; Narayanan et al., 2021; Shi et al., 2023). This approach consists of providing a representative trajectory for each of the modes and a corresponding classification score indicating how likely each mode is out of a fixed number of predicted modes. The main issue with this is that the variability within the modes is entirely disregarded. This variability is generally important for adequately informing motion planners and providing safety guarantees. To address this issue, one can instead use methods which learn continuous distributions. One way of doing this, is explicitly learning the parameters of multi-modal distributions such as those of a Gaussian Mixture Model (Salzmann et al., 2020; Shi et al., 2022; Varadarajan et al., 2022; Nasr Azadani and Boukerche, 2023) or of a mixture of Laplace distributions (Zhou et al., 2022). However, these approaches are often limited to assuming a specific type of distribution (often Gaussian) as the components of the multi-modal distribution, which can lead to inaccuracies if the assumption is incorrect. To combat this, Normalizing Flows (NFs) (Sun et al., 2021; Schöller and Knoll, 2021; Mészáros et al., 2024) can be used to capture multi-modal distributions without assuming the form of the distribution. To be able to provide an explicit likelihood of a drawn sample, NFs are built on the principle of bi-directional transformation functions which transform a high-dimensional, multi-modal distribution into a Normal Gaussian distribution. Since Gaussian distributions are easy to sample from and the likelihood of drawn samples is known, together with the transformation functions of the NF it is possible to obtain samples and their likelihoods in the multi-modal distribution.

Uncertainty Quantification in Prediction Models. Although the above methods aim to achieve the best possible fit over the distribution of people's future actions, they do not provide information on how good the predicted distributions are. This is important since it is difficult to ensure a perfect fit even on the training data. It is even more difficult to guarantee perfect generalization to situations that the models have not seen in the training data, particularly when these new situations differ substantially from the training data. Take for example a model trained in the UK, where people drive on the left, and deployed in the Netherlands, where people drive on the right. Or imagine a model trained in the US, where urban environments are often well structured and the main form of transportation is cars, and deploying this model in Vietnam which is governed by chaotic scooter traffic. Even everyday events such as changing weather conditions, wildlife on the road, roadworks, etc. can place a model in previously unseen situations which can negatively impact its predictions.

Since inaccuracies in the prediction's risk negatively impact the motion planner's perception of the future trajectories of the AV's surrounding agents, it can be valuable to quantify these inaccuracies using uncertainty quantification (Ghanem et al., 2017) to provide a form of confidence bound around the predicted trajectories. A possible way of achieving this is by leveraging conformal prediction (Lindemann et al., 2023; Dixit et al., 2023). This approach generates bounds for a predicted variable which guarantee that the true outcome will lie within these bounds with a specified level of confidence. This means that if a confidence of 90% is specified, the true value will on average fall within the confidence bounds 90% of the time. The current application of these approaches has thus far been limited to uni-modal predictions and has yet to be extended to multi-modal predictions. Another way of obtaining this form of uncertainty quantification is through a model that provides both aleatoric and epistemic uncertainty, where the epistemic uncertainty is an indicator of how much new situations differ from the model's training data. The epistemic uncertainty can then be used to assess the predicted aleatoric uncertainty (Neumeier et al., 2024; Suk and Kim, 2025).

With many models increasingly being obtained as neural networks trained by deep learning, it bears special consideration how uncertainty is quantified in these types of models. Roughly, these approaches can be divided into two groups: deterministic methods and Bayesian methods (Gawlikowski et al., 2021). In a deterministic deep network, each predictive distribution is estimated by a single forward propagation through the network after which a softmax function is applied—the outputs of the softmax are then interpreted as predicted probabilities. While feasible, these methods tend to be overconfident (Lakshminarayanan et al., 2017; Gawlikowski et al., 2021). A Bayesian network directly outputs probabilities by marginalizing the likelihood of the input with the estimated posterior distribution; this however is generally intractable. Tractable approximation methods have been demonstrated to work well, and include Monte-Carlo-dropout (MC-dropout) (Gal and Ghahramani, 2016), which repeatedly assesses each sample in perturbed versions of the network, and deep ensembles (Lakshminarayanan et al., 2017), where a collection of networks is trained and each sample is inferred in each network. For energy-efficient hardware, analogs have been developed for spiking neural networks suitable for implementation in neuromorphic hardware (Sun et al., 2023; Sun and Bohté, 2025).

3.3. Risk metrics for motion planning

To assess risk in motion planning, we require a cost function that captures the utility of the AV's task objectives. From a numerical optimization perspective, the motion planner attempts to minimize a cost that depends on the AV's decision variables and on uncertainty. Common terms in such a cost include a weighted distance between the car's current and target positions, a safe collision distance with the surrounding agents, and a penalty on control effort to reduce fuel consumption.

In Section 3.0.0.1, the uncertainty corresponds to the trajectories of the surrounding agents in the scene. A *risk metric* then maps this uncertain cost to a single number that represents risk. A simple example is the average (expected) cost, which reflects average task performance.

Several risk metrics are widely used in operations research and recently robotics—that shape motion planning for autonomous driving along safety (risk-averse vs. risk-seeking) and conservativeness. *Worst-case (robust) optimization* minimizes the maximum cost over bounded uncertainty, commonly via *Robust Model Predictive Control (RMPC)* as a minimax formulation (Gandhi et al., 2021; Tolksdorf et al., 2023; Zhou et al., 2025); by considering only the support (e.g., two-sigma bounds) rather than full distributions, it can guarantee safety within assumptions but is often conservative. *Risk-neutral* planning instead minimizes expected cost, favoring computational efficiency and unbiased behavior but remaining vulnerable to rare events (Nyberg et al., 2021). *Risk-averse* methods leverage distributions to allow controlled violation probabilities: *Value at Risk (VaR)* sets a threshold α such that costs stay below it with probability at least $(1-\alpha)$, enabling chance-constrained trade-offs between safety and performance (de Groot et al., 2021; Nair et al., 2022; Mustafa et al., 2023; Ren et al., 2024); *Conditional Value at Risk (CVaR)* further penalizes tail severity by minimizing a threshold plus the mean excess in the worst (α) outcomes, aiding heavy-tailed settings and yielding more conservative behavior; CVaR-based chance constraints can be embedded in receding-horizon optimization to bound, for example, distance to safe regions within the (α) quantile (Hakobyan et al., 2019; Chen et al., 2022; Lew et al., 2023; Yin et al., 2022). Finally, *risk-sensitive* metrics expose a tunable parameter to span risk-seeking to neutral to risk averse behaviors: *mean–variance* minimizes expected cost plus (γ) times its variance, where $\gamma < 0$ encourages variance (risk-seeking), $\gamma = 0$ recovers risk-neutral, and $\gamma > 0$ reduces variance (risk-averse) (Sobel, 1982). *Entropic risk* uses an exponential utility with sensitivity (λ) to summarize all moments and exponentially weigh losses, generalizing mean–variance and emphasizing tails, particularly under non-Gaussian, heavy-tailed uncertainty (Marthe et al., 2025).

3.4. Collective risk in motion planning

Motion planning for AVs has a strong impact on the distribution of risk among road users. While AV passengers may favor risk-minimizing strategies that prioritize their own safety, such behavior can increase risk for vulnerable road users (VRUs) such as pedestrians and cyclists, raising ethical concerns about fairness in risk allocation (Kirchmair and Paulo, 2023; Gros et al., 2025). To address this, researchers have explored ethical risk allocation frameworks that balance self-preservation with *collective safety*. Approaches like in Geisslinger et al. (2021, 2023) incorporate ethical principles into risk–cost functions, ensuring AVs consider their impact on surrounding agents. Similarly, Schwarting et al. (2019) model AV interactions as a dynamic game using Social Value Orientation (SVO) to quantify selfishness or altruism, enabling AVs to predict human behavior and generate socially aware control policies that balance individual and collective risk.

Beyond ethical considerations, predictability is key to mitigating uncertainty-induced risk in motion planning. AVs that behave in a legible and socially consistent manner reduce the ambiguity experienced by human drivers, fostering smoother interactions and improving coordination. Studies in Human–Robot Interaction (HRI) highlight the importance of clearly communicating intent to improve coordination. This principle has been applied to autonomous vehicles (AVs) through various motion planning frameworks. For example, some approaches encourage AVs to signal their intentions early by initiating avoidance maneuvers sooner (Geldenbott and Leung., 2024), while others focus on reducing uncertainty in collision avoidance by dynamically adjusting goal regions around nearby agents (Bastarache et al., 2023). Predictability-based trajectory costs have been introduced to ensure that AV motions align with human expectations, minimizing the risk of unexpected behavior (Brüdigam et al., 2018). Furthermore, Gil et al. (2025) introduce a generalizable approach where agents predict their own behavior to align with human expectations, fostering implicit social conventions without explicit communication.

In all, safe, effective decision-making in autonomous vehicles requires planning under uncertainty through appropriate notions of risk. We find that the meaning of uncertainty between neuroscience and robotics is essentially the same, building on probability theory given some model. In robotics however, computational tractability and measures of risk are of prime consideration. As enumerating all agents, modes, and futures quickly becomes computationally intractable, the planner must rely on prediction models that estimate uncertainty and derived risk metrics to act sensibly in real time. Beyond individual safety, collective considerations have to be included, like how risk is allocated across road users and how legible, predictable AV behavior may reduce others' uncertainty (Dey et al., 2020). Thus, while capabilities are rapidly increasing, open challenges include real-time effective solutions, which is particularly complicated and important when human agents are present.

4. Psychological uncertainty in human–automation interaction

Rowe (1994) conceptualizes uncertainty as a state of not knowing. Beyond the direct measures of uncertainty when planning with noisy and limited sensory information, behavior is also shaped by a more meta-cognitive perspective: autonomous agents and people mutually shape each other's behavior, with human plans and actions that could be influenced by the uncertainty they feel based on noisy perceptions and fluctuating internal states, including when interacting with technology.

In this context of Human–automation interaction (HAI), we define human psychological uncertainty as, "a perceived inadequacy, during the interaction, to form an accurate prediction about their role as a user or the intention of the automated agent". This inadequacy can have significant implications for human decision-making during HAI (Lindley, 2013), as individuals rely on their expectations of the agent's actions to guide their responses. For example, when there is a mismatch between expectations and the behavior of an AV, increased uncertainty can affect the driver's behavior, leading to hesitation, greater risk, or even crashes (Victor et al., 2018). Beyond the scope of user's safety, experiencing uncertainty could lead to negative emotions or undesirable feelings. Such uncertainty takes many forms, for example about future outcomes like delays, and, as a trait-like construct, therefore term *psychological uncertainty* is used to emphasize the subjective and psychological nature of the construct (Windschitl and Wells, 1996). It can mediate human responses (Kappes et al., 2018; Windschitl and Wells, 1996) and is associated to negative emotions (Anderson et al., 2019; Morriss et al., 2022; Franssen et al., 2024b).

In HAI, psychological uncertainty influences trust (Lee and See, 2004; Hoff and Bashir, 2015; Schaefer et al., 2016); for example, users may distrust a delivery drone if they feel uncertain about its flight behavior (Lingam et al., 2025c). Psychological uncertainty likely affects acceptance and perceived risk, and may contribute to errors (Lingam et al., 2024a). Addressing it is critical for natural human–agent interaction (Muthugala and Jayasekara, 2016) and public acceptance. In Human–Drone Interaction (HDI), a new and developing mode of interaction compared to autonomous cars, experts argue that managing psychological uncertainty is essential for deploying drones in public spaces (e.g., emergency response, delivery, surveillance) (Lingam et al., 2024a). Lingam et al. (2025c) define psychological uncertainty in HDI as "a state of doubt experienced by humans when drone interactions deviate from expectations, leading to a loss of understanding of the drone's intentions or next actions". Previous research further emphasize that agent understandability and predictability are core to this construct (Lingam et al., 2025c). Franssen et al. (2026) further define psychological uncertainty through a systematic review as *"the mental experience of a lack of knowledge about the past, present, or future, which makes it impossible to accurately predict outcomes or determine which of multiple possibilities will be true, and which in turn shapes emotional responses and behavioral output"*.

Psychological uncertainty in HAI is typically measured as user-perceived uncertainty. Prior work has used subjective ratings towards an agent or its elements (Franssen et al., 2024a; Lingam et al., 2025c). Franssen et al. (2024a) assessed responses to static AV design elements with a 7-point Likert scale ("Absolutely Uncertain" to "Absolutely not Uncertain"), showing uncertainty varies across individuals in similar conditions. Lingam et al. (2025c) introduced a continuous Likert scale to track perceived uncertainty across interaction stages with a delivery drone, revealing stage-dependent variation. To probe implications for mental states, researchers correlate uncertainty with emotions to explain choices (e.g., travel) and explore alternative measures (Franssen et al., 2024b). Because emotions such as fear, frustration, and anxiety are associated with psychological uncertainty (Grupe and Nitschke, 2013; Morriss et al., 2022), mapping these patterns can refine measurement and guide the design of automated systems (e.g., drones, AVs) towards more transparent and predictable interactions.

4.1. *Factors affecting psychological uncertainty in human–automation interaction*

Given the variety of human–automation scenarios, many factors can influence psychological uncertainty. We focus on automated vehicles (AVs) and drones—domains where trust, acceptance, and physical risk are acute and important.

Automated vehicles. AVs differ in form, context, and capability, see for instance (SAE International, 2021). Psychological uncertainty arises for occupants and other road users. Market deployment proceeds incrementally, with benefits tempered by unanticipated failures: drivers can be surprised by AV behavior, leading to errors (de Winter et al., 2023), especially when systems communicate poorly (Sarter et al., 1997). Transparency supports trust (Choi and Ji, 2015), which influences perceived risk (Choi and Ji, 2015) and relates to psychological uncertainty (Mayer et al., 1995; Mitchell, 1999). Although trust, acceptance, and safety are often studied, psychological uncertainty itself is not. One of the few studies on psychological uncertainty in an AV-context found that passengers of an AV can experience uncertainty about the design and functioning of the automated vehicle in traffic. For example, passengers report that interior features such as opaque windows or removing the steering wheel increase uncertainty (Franssen et al., 2024a).

From outside the vehicle, diverse road users, such as pedestrians, cyclists, and manual drivers, face their own uncertainties (Rasouli and Tsotsos, 2019; Dey et al., 2020; Faas et al., 2020a). Factors related to automated vehicles could lead to uncertainty, as AV driving styles can differ from human drivers' (Bazilinskyy et al., 2021) which complicates anticipation of future actions. The absence of a human interlocutor inside the AV may further increase uncertainty in situations requiring explicit communication or after incidents. Faas et al. (2020b) found that pedestrians face greater uncertainty when interacting with self-driving vehicles because they can no longer rely on traditional social cues from a human driver, such as eye contact or gestures. This can lead to uncertainty to be present in potential space sharing conflicts, which can be solved through acquiring more information about the situation or by a road user changing their behavior (Markkula et al., 2020). One way to provide such additional information to those outside of the vehicle, is by adding external human–machine interfaces (eHMI) on an AV (Faas et al., 2020a; Dey et al., 2021, 2020). Such eHMIs can then influence human's interactions with AVs, as demonstrated by Faas et al. (2020b) who found that when eHMIs provided clear information that the vehicle intended to yield, pedestrians initiated crossing earlier (shorter crossing onset times), reported higher perceived safety, and evaluated the interaction more positively. Like eHMIs, road infrastructure can also be used to influence human interaction, both by providing information about regulations (e.g., right of way) and by influencing vehicle kinematics (Kalantari, 2023; Rasouli and Tsotsos, 2019).

Drones. once primarily used for aerial photography, are now increasingly being utilized in delivery, surveillance, entertainment, and warfare. As they enter everyday public spaces, the importance of both physical risks and *psychological uncertainty* increases. A human-centered interview and design study identifies several drivers of uncertainty during public-space deliveries: the degree of human involvement (e.g., teleoperation vs. autonomy), the dynamics and crowding of the setting, people's familiarity with the technology, and the criticality of the task (Lingam et al., 2025d). Uncertainty often differs between direct users, such as package recipients, and non-users, such as bystanders, largely due to knowledge gaps (Lingam et al., 2025d). Responsibilities (who must act or yield), proximity to the drone, and environmental and situational factors, such as landing-site geometry, birds, children, and pets, further influence uncertainty (Lingam et al., 2025d).

Design and behavior play a central role. Echoing (Franssen et al., 2024a), specific design choices can raise or reduce uncertainty for both users and non-users. Flight behavior, proximity during approach/hover, acoustic signature, and visual appearance all influence judgments (Lingam et al., 2024a, 2025c, 2024b, 2025b). For instance, straight-line approaches were perceived as less predictable than curved paths, increasing uncertainty and reducing trust (Lingam et al., 2025c). Across studies, a consistent driver of uncertainty is insufficient information about drone intentions; HDI experts (Lingam et al., 2024a) and participants emphasize the need for visual interfaces and behaviors that explicitly communicate state and intent (e.g., approach path, landing point, next action) (Lingam et al., 2025d,a).

Causality and inherent uncertainty. Establishing causality for factors that contribute to uncertainty in the context of automation can be difficult. Features like tinted AV windows (Franssen et al., 2024a) or nearby children around a delivery drone (Lingam et al., 2025d) may heighten uncertainty, yet similar effects could occur without automation. Thus, not all uncertainty-inducing factors are unique to automated agents.

4.2. Preliminary model of psychological uncertainty in human–automation interaction with vehicles and drones

To explore the impact of psychological uncertainty on decision-making and responses within HAI, we propose a preliminary model that outlines the key factors contributing to this psychological uncertainty in the context of autonomous agents (Fig. 3). We expect this model to serve as a foundational step towards understanding and mitigating uncertainty in HAI, ultimately improving the design of more intuitive and user-trusted automated systems.

The proposed model is inspired from the models of human-information processing by Parasuraman et al. (2000) and Wickens et al. (2021). The model comprises four primary stages in a linear manner: human sensory processing, human perception, human decision-making, and human response execution (including no response). These processes provide feedback that can influence the automated systems (Wickens et al., 2021). In line with Wickens et al. (2021), the attentional resources stage is incorporated to account for the finite cognitive capacity humans possess for processing and acting on information. Given the perceptual nature of psychological uncertainty and its dependence on available attentional resources, this construct is positioned between perception and attention in the model. At the start of the interaction, the 'Sensory processing' stage involves acquiring and registering multiple sources of information from the automated agents. This includes positioning and orienting sensory organs, processing inputs, filtering the information cognitively, and selectively attending to relevant stimuli from automated agents. The 'Perception' stage centers on conscious perception and the manipulation of processed and retrieved information within working memory. When there is a lack of or difficulty interpreting the information in 'Perception', 'Psychological uncertainty' arises. 'Attentional resources' are used to select information, that helps with interpretation and potentially reducing 'Psychological uncertainty.' In the 'Decision making' stage, decisions are formed, and responses are selected based on prior cognitive processes. Finally, the 'Response selection' stage involves the implementation of these responses, which may translate to various forms of actions or inactions with respect to the automated vehicle. The system then has the chance to respond to this action. This could be a responsive action, by the system, aimed at reducing the psychological uncertainty of its user or an action that is counterintuitive to the user thereby increasing uncertainty.

Compared to the classical models of Parasuraman et al. (2000) and Wickens et al. (2021), which treat information processing as a largely feed-forward chain and do not represent subjective uncertainty as a distinct construct, here psychological uncertainty is given an explicit function between perception and attentional resources, where it both arises from and modulates information processing. Additionally, we close the loop: the human's response feeds back to the automated system, which can then act to raise or lower the user's uncertainty. This bidirectional coupling thus connects the otherwise human-internal model to the AV's own prediction-and-planning loop (Section 3.2).

4.3. Implications of psychological uncertainty on human behavior

As discussed in the previous section, psychological uncertainty plays a critical role in shaping human actions, and thus their behavior. According to both experts (Lingam et al., 2024a) and public users (Lingam et al., 2025c,d), such uncertainty – particularly when stemming from a lack of awareness about the automated system's intentions or behavior – can elevate perceived risk during interaction. This often leads to discomfort and reduces trust in automation. In response, humans tend to hesitate approaching the autonomous vehicle and will maintain physical distance from the system (Wojciechowska et al., 2019; Yeh et al., 2017). Schaefer et al. (2016) found that trust becomes especially important when autonomous systems operate in uncertain, unplanned circumstances. The degree of psychological uncertainty – ranging from mild hesitation to high-risk avoidance – can lead to perceived risk and varying behavioral responses, such as slower approach, pausing, keeping distance, or complete disengagement. While making the system's intentions explicit (e.g., through visual interfaces) can reduce uncertainty and encourage approach behaviors or safe crossing ahead of the system (Dey et al., 2021; Hetherington et al., 2021), other forms of uncertainty persist. For example, when a drone clearly communicates its drop-off intentions, recipients may approach the delivery spot while experiencing lingering uncertainty and perceived risk about whether it is safe to do so (Lingam et al., 2026). The effects of psychological uncertainty on human behavior may also depend on contextual factors, such as the urgency of the task (e.g., emergency response vs pizza delivery), prior experiences with automation, the human role or social norms in public spaces. Future research should investigate how different types (e.g., uncertainty about system motion vs. safe to approach) and intensities of psychological uncertainty shape human behavior in real-world human–automation interactions.

These behavioral signals have a direct, yet underexploited, consequence for the technical pipeline described in Section 3.2: hesitation, slowed approach, pausing, and last-moment changes of mind widen and multiply the motion *modes* that a prediction model must represent (Kalantari, 2023; Tian et al., 2023), and they inflate the per-mode trajectory-uncertainty ellipsoids that the planner ultimately consumes (Fig. 2). Effectively, a bystander's psychological uncertainty, seemingly subjective, feeds back into the objective distribution over human motion that the AV must predict and plan against (Markkula et al., 2020). This gives a concrete yet qualitative link between the two perspectives: the measurements of psychological uncertainty point to those situations in which human motion becomes harder to predict, and thus to situations in which the AV should expect wider predictive uncertainty and may need to plan more conservatively or behave more legibly to actively reduce uncertainty.

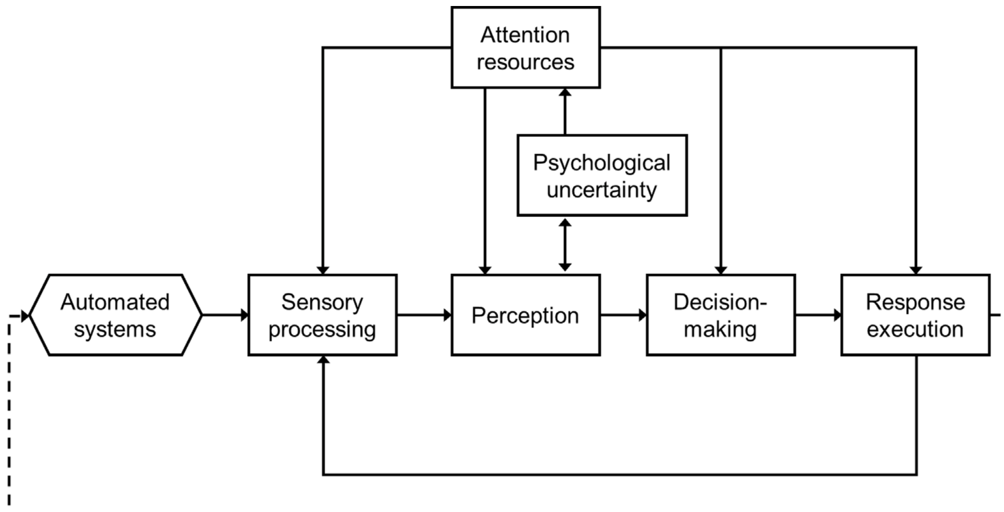

Fig. 3. The role of psychological uncertainty in the human information processing model for interactions with automated systems. The solid lines illustrate how automated systems influence elements of the human information processing model and psychological uncertainty. The dotted lines indicate the potential for human responses to impact automated systems during the interaction. Cognitive components like working memory, planning and long-term memory are taken as part of perception and decision-making elements. Note that in the brain, decision making and response selection are an integrated and competition-based process (Cisek, 2007).

Summarizing, in the presented model, psychological uncertainty and perception likely influence each other bidirectionally. Observing an automated agent can elicit uncertainty, reallocating attention to seek additional sensory input (Dieterich et al., 2016; Rose et al., 2022); in turn, the notion of uncertainty shapes how subsequent information is processed (Gupta and Solo, 2010). Thus, we argue that in HAI, uncertainty is not a mere byproduct of perception but an active, dynamic driver.

5. Synthesis

All three perspectives frame uncertainty as a fundamental and pervasive aspect influencing decision-making, perception, and action, whether in biological systems, autonomous agents, or human-automation interactions, as summarized in Fig. 4. Here we study in particular how entities (humans, AVs) make decisions and select actions despite incomplete or noisy information and an unpredictable environment: decision-making under uncertainty.

To solve this, each field relies on models (Bayesian, computational, cognitive) that infer current states, predict future states/behaviors, and understand the causes of observations. Actions are then selected by minimizing a risk functional of the induced cost. In neuroscience, this reads as precision-weighted prediction errors and uncertainty-guided control; in behavioral psychology, as choice-making under risk undergoing uncertainty calibration; in robotics, as risk-aware planning under uncertainty from posterior predictive scene futures. In all three, uncertainty is *a critical design parameter*: it quantifies and calibrates risk, and thus shapes behavior. Whether it is the brain selecting an action, an autonomous vehicle selecting a plan, or a human adapting to an AV, all perspectives share an underlying goal to perform effectively and safely by accounting for uncertainty.

Yet, most neuroscience models focus on the brain's mechanisms and algorithms for representing, processing, and mitigating uncertainty at neural and cognitive levels (e.g., neural noise, optimal feedback control, causal inference) while using a cost function to determine action choices. The robotics perspective and emphasizes computational tractability, safety, all explicitly in the context of specific risk metrics. HAI then examines factors like trust, predictability, and understandability to determine the factors influencing human perception and behavior at a more meta-cognitive level. Importantly, human decision making is also driven by elicited emotion and motivation, a phenomenon that an AV often does not take into account when making choices.

5.1. A cross-cutting example: a pedestrian at the crosswalk

Rather than three separate problems, we find that the three perspectives describe one coupled system: consider a pedestrian deciding whether to cross in front of an approaching AV through each view (Fig. 5).

Neuroscience. The pedestrian must infer the AV's state and intent from noisy, delayed visual input: its distance and closing speed from optic flow, and whether it is decelerating. This is a causal-inference problem (has it detected me? will it yield?) whose posterior is multi-modal: "it will stop" versus "it will not.". The pedestrian then applies a risk functional to that posterior, as the cost of being wrong about "it will stop" is severe, which rationalizes cautious, risk-averse behavior such as waiting or edging forward to gather more evidence (Kalantari, 2023; Tian et al., 2023), an information-seeking, active-inference action.

Robotics (the AV). The AV faces the mirror-image problem. From the pedestrian's observed kinematics it must predict a distribution over motion modes (cross now/wait / hesitate), with within-mode trajectory uncertainty (Section 3.2). It then selects a risk metric (Section 3.3) that maps this distribution onto a risk-aware plan, such as risk-averse yielding, risk-neutral nudging, etc. Ideally, this also accounts for the impact of that plan on the pedestrian (Section 3.4).

HAI. The two inference-then-risk loops are dependent: each agent's action is the other's evidence. If the pedestrian is psychologically uncertain, for instance because he/she is unsure whether the AV has seen them, they hesitate, which (as argued in Section 4) widens the very motion-mode distribution the AV must predict, making the AV less certain in turn. Conversely, an easily interpretable AV action reduces the pedestrian's psychological uncertainty, sharpens their behavior into a single predictable mode, and so reduces the AV's predictive uncertainty as well. The same scenario thus exhibits the gap: the AV's formal risk machinery and the pedestrian's subjective uncertainty are tightly coupled in reality but are usually modeled separately.

5.2. Metacognition as a potential bridge

The clearest pathway from the shared inference-then-risk structure to an actual integration runs through *metacognition*, defined as the monitoring and control of one's own cognitive processes (Fleming and Daw, 2017). Under different jargon, Metacognition is studied in all

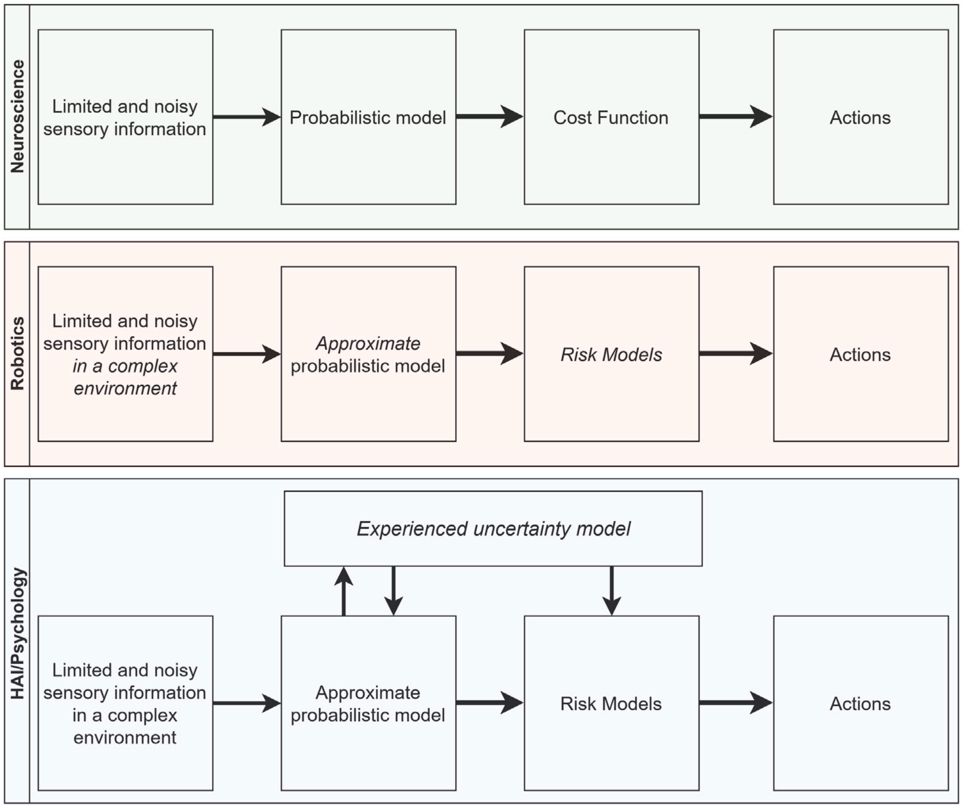

Fig. 4. The role and definition of uncertainty in guiding behavior in Neuroscience, Robotics and Human–Automation Interaction (HAI).

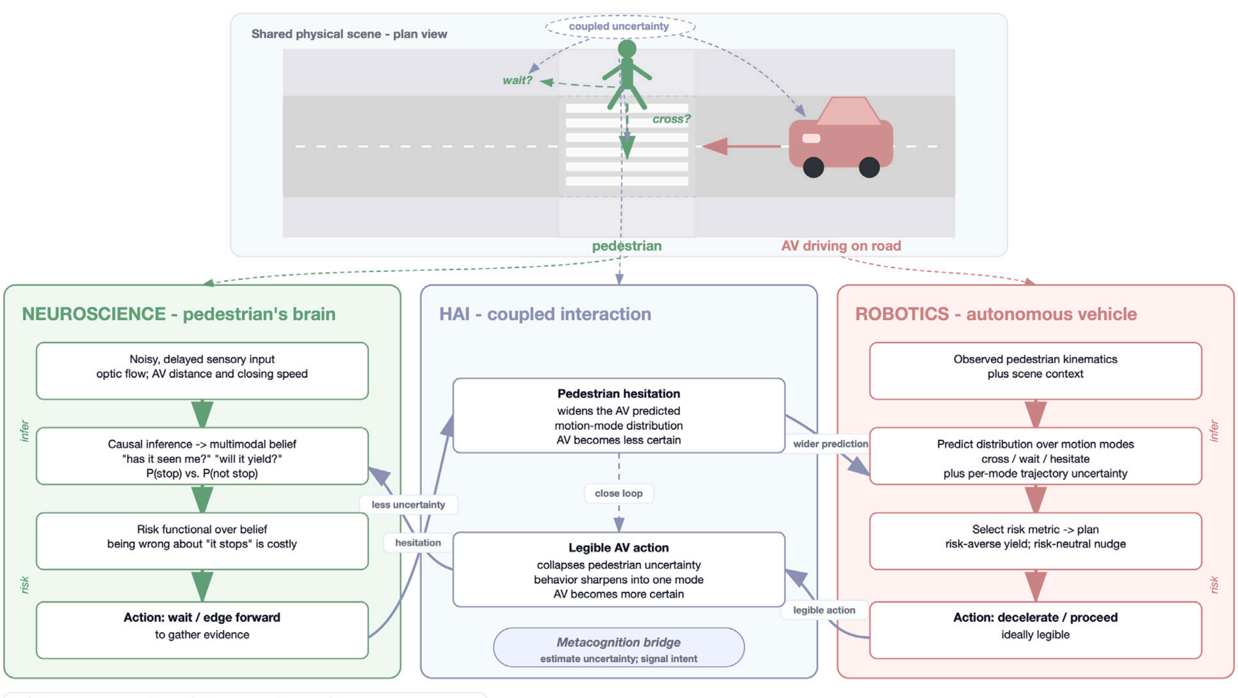

Fig. 5. A single scenario from three perspectives: a pedestrian deciding whether to cross in front of an approaching autonomous vehicle, all described as one coupled system. The pedestrian (*left*, neuroscience) infers the AV's intent from noisy, delayed sensory input and applies a risk functional to a multimodal posterior ("will it stop?"). The AV (*right*, robotics) solves the mirror-image problem: it predicts a distribution over the pedestrian's motion modes and selects a risk metric that maps this onto a plan. The two inference-then-risk loops are *coupled* (*center*, HAI): the pedestrian's psychological uncertainty produces hesitation that increases the motion-mode distribution the AV must predict, while a clearly communicated AV action reduces the pedestrian's uncertainty and sharpens their behavior into a single predictable mode. Metacognition (Section 5.2) is the suggested bridge to close this loop.

three fields, and maps cleanly onto the two-stage notion of Section 1. We outline a three-part correspondence.

Monitoring vs uncertainty quantification. In neuroscience, metacognitive monitoring produces confidence judgments and error awareness as estimates of the reliability of one's own inference. In robotics, this corresponds to the uncertainty quantification of Section 3.2: calibrated aleatoric and epistemic uncertainty, conformal-prediction bounds, and the use of epistemic uncertainty to detect that a situation lies outside the training distribution (Kendall and Gal, 2017; Neumeier et al., 2024). A metacognitively effective AV needs to have reported confidence that is well calibrated to its actual error rate.

Control vs risk-metric selection. Metacognitive control adapts behavior to monitored confidence, such as slowing down when uncertain, and gathering information when prediction error is high. This is the function served in robotics by the choice and tuning of a risk metric (Section 3.3): risk-averse planning (e.g., CVaR) is preferred when epistemic uncertainty is high, whereas efficient risk-neutral behavior is preferred when the situation is familiar (Botvinick and Toussaint, 2012). In Metacognition terms, the risk-metric then becomes a choice that an agent could modulate online from its own confidence.

Communication vs legibility. Finally, people do more than just estimate and act on their uncertainty: they also communicate calibrated confidence to others so that partners can anticipate them. This predictability objective is already pursued in AV motion planning (Section 3.4): signaling intent early (Geldenbott and Leung., 2024), shaping goal regions to reduce others' uncertainty (Bastarache et al., 2023), and producing self-predictable motion aligned with human expectations (Gil et al., 2025). Communicating calibrated confidence as measured by metacognition is then the logical solution by which an AV can reduce the *psychological* uncertainty of Section 4 in the people around it.

Together, these correspondences suggest a dynamic description of the side-by-side loops of Fig. 4: an AV estimates its own uncertainty, lets that estimate modulate how cautiously it plans, and communicates the result clearly. This would be a first stab at closing the gap between the objective uncertainty it manages internally and the subjective uncertainty it induces externally. For now, this is a hypothesis: calibrated metacognitive models that scale to full driving scenes are an active topic (Fleming and Daw, 2017). One open challenge for example is mutual adaptation, where as agents continually adjusts to each other, coupled interactions can oscillate; clear, confidence-calibrated communication is then a likely fruitful approach to keeping such mutual adjustment from locking up.

6. Discussion and conclusion

This perspective sets out to understand how the behavior of autonomous vehicles and systems can be informed by, and aligned with, human uncertainty about their actions. Our examinations from three perspectives underscores the diverse nature of uncertainty and its critical role in the interrelated fields of neuroscience, robotics, and human–automation interaction (Schwarting et al., 2019). It also highlights the challenge: to integrate autonomous vehicles into society will require a probabilistic-computation first approach, for example by scaling uncertainty-native neuroscience frameworks to the complex environment of autonomous vehicles, while also properly accounting for the likely and diverse potential effects of an autonomous vehicle's behavior on that of humans.

Neuroscience reveals the brain's sophisticated, often-times presumed Bayesian, strategies for navigating inherent sensory, motor, and environmental uncertainty, offering inspiration for robust AI (Knill and Pouget, 2004; Pezzulo et al., 2018). At the same time, the current mechanistic frameworks in neuroscience are quite far removed from being able to tackle the kind of complex problems considered in robotics. The robotics perspective advances the mathematical and computational frameworks that autonomous vehicles (AVs) employ to predict, plan, and manage risk in dynamic environments (Sadigh et al., 2016; Sunberg and Kochenderfer, 2018; Mészáros et al., 2024; Mustafa et al., 2023). These technical approaches, while powerful, often grapple with computational intractability and may not fully capture the nuances of interacting with humans. The psychological perspective highlights that human interaction with automation is strongly shaped by perceived uncertainty, which influences trust, stress, acceptance, and behavior, often independent of the objective uncertainty managed by the system itself (Hoff and Bashir, 2015; Kohn et al., 2021; Lingam et al., 2025c; Franssen et al., 2024a).

These respective areas of focus naturally follow from the problem of computational tractability: solving high-level planning and decision-making under uncertainty requires approximations and algorithms that are actively being developed (Mészáros et al., 2024). Only for low-level tasks is the problem sufficiently constrained to probe the biological solutions that the brain uses at a mechanistic level (e.g. Körding et al., 2007). At the high-level, cognitive neuroscience and psychology have a long tradition of observing and modeling behavior, yet corresponding mechanistic models are both underconstrained and computationally highly intensive. The approaches in robotics as outlined underscore this, as they illuminate both the many factors involved and the combinatorial explosions that accompany full descriptions.

Thus, a central challenge is the disconnect between the mathematical representations of uncertainty in AVs and the high-level cognitive and emotional factors influenced by notions of uncertainty that drive human responses. Bridging this gap is vital for developing AVs that are not only technically proficient in managing their own uncertainties but also behave in ways that are understandable, predictable, and inspire confidence in human users and bystanders. Note that the interaction is bidirectional: AVs must model human uncertainty and derived behavior, and humans experience uncertainty about AVs that hinders interaction, acceptance and integration.

We have argued (Section 5.2) that the most promising bridge across these perspectives is the study of *metacognition*, which studies the monitoring and control of one's own cognitive processes (Fleming and Daw, 2017). An AV that estimates and calibrates its own uncertainty (Kendall and Gal, 2017) uses those estimates to adapt its planning through risk-sensitive mode switching (Botvinick and Toussaint, 2012), and communicates the resulting calibrated confidence in a clear manner (Bartneck and Moltchanova, 2020). Such an approach would address both the objective uncertainty it manages internally and the psychological uncertainty it induces in others in a single mechanism. Mechanistically, such metacognitive models remain an active research direction towards the planning-under-uncertainty that autonomous vehicles require.

Overall, addressing uncertainty in the context of autonomous vehicles demands an integrated, interdisciplinary approach. The development of AVs that seamlessly and safely integrate into human-populated environments relies on our ability to understand, model, and manage uncertainty from technical, biological, and psychological viewpoints (Mavrogiannis et al., 2023). Future research must focus on creating unified frameworks that translate insights across these domains. This includes designing AVs that not only robustly handle their operational uncertainties using advanced prediction and planning, but also actively consider and mitigate the psychological uncertainty experienced by humans. This involves developing more predictable AV behaviors, potentially inspired by how humans naturally manage and communicate uncertainty. Ultimately, high-level neuroscientific models have to account for both: modeling the ability to mechanistically plan long-range actions while both accounting for and calculating measures of uncertainty, all in a computationally tractable manner. Together, such a holistic understanding will enable autonomous vehicles that are truly cooperative, intuitive, and trustworthy partners in our increasingly shared world, creating safer and more efficient interactions for all.

CRediT authorship contribution statement

J.R.E. Cooke: Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Conceptualization. **S.N. Lingam:** Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Conceptualization. **P. Alavian:** Writing – review & editing, Writing – original draft, Methodology, Investigation. **C. Collast:** Writing – review & editing, Investigation. **G.C.H.E. de Croon:** Writing – review & editing, Supervision, Investigation. **M. Franssen:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Conceptualization. **J. Hagenaars:** Writing – review & editing, Investigation, Conceptualization. **J. Kober:** Writing – review & editing, Writing – original draft, Supervision, Investigation, Conceptualization. **P. Morsink:** Writing – review & editing, Validation, Supervision. **K. Liu:** Writing – review & editing, Investigation. **Y-H. Liu:** Writing – review & editing, Validation, Methodology. **E. Matthews:** Writing – review & editing, Validation, Conceptualization. **J. Mejias:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Conceptualization. **A. Palffy:** Writing – review & editing, Validation, Resources. **C.M.A. Pennartz:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **W. Sanberg:** Writing – review & editing, Validation. **L. Selen:** Writing – review & editing, Supervision, Methodology. **T. Sun:** Writing – review & editing, Methodology, Investigation. **N. van Nes:** Writing – review & editing, Validation, Funding acquisition. **R. Verstegen:** Writing – review & editing, Writing – original draft, Visualization, Investigation. **M. Martens:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization. **J. Alsono-Mora:** Writing – review & editing, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization. **W.P. Medendorp:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization. **S.M. Bohte:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Sander Bohte reports financial support was provided by Dutch Research Council. Tao Sun reports financial support was provided by Dutch Research Council. Pieter Medendorp reports financial support was provided by Dutch Research Council. J. Alonso-Mora reports financial support was provided by Dutch Research Council. Marieke Martens reports financial support was provided by Dutch Research Council. R. Verstegen reports financial support was provided by Dutch Research Council. C.M.A. Pennartz reports financial support was provided by Dutch Research Council. J. Mejias reports financial support was provided by Dutch Research Council. K. Liu reports financial support was provided by Dutch Research Council. J. Kober reports financial support was provided by Dutch Research Council. J. Hagenaars reports financial support was provided by Dutch Research Council. M. Franssen reports financial support was provided by Dutch Research Council. G.C.H.E. de Croon reports financial support was provided by Dutch Research Council. P. Alavian reports financial support was provided by Dutch Research Council. S.N. Lingam reports financial support was provided by Dutch Research Council. J.R.E. Cooke reports financial support was provided by Dutch Research Council. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

SB, TS, PM, JA-M, MM, RV, CP, JM, KL, JK, JH, MF, GdC, PA, SL, JC were supported by NWO-NWA grant NWA.1292.19.298.

Data availability

Data will be made available on request.

References

Acerbi, L., Dokka, K., Angelaki, D.E., Ma, W.J., 2018. Bayesian comparison of explicit and implicit causal inference strategies in multisensory heading perception. PLoS Comput. Biol. 14 (7), e1006110. http://dx.doi.org/10.1371/journal.pcbi.1006110.

Alchian, A.A., 1950. Uncertainty, evolution, and economic theory. J. Polit. Econ. 58 (3), 211–221.

Anderson, E.C., Carleton, R.N., Diefenbach, M., Han, P.K., 2019. The relationship between uncertainty and affect. Front. Psychol. 10, 469966. http://dx.doi.org/10.3389/FPSYG.2019.02504/BIBTEX.

Bartneck, C., Moltchanova, E., 2020. Expressing uncertainty in human-robot interaction. Plos One 15 (7), e0235361.

Bastarache, J.L., Nielsen, C., Smith., S.L., 2023. On legible and predictable robot navigation in multi-agent environments. IEEE Int. Conf. Robot. Autom.

Bazilinskyy, P., Sakuma, T., de Winter, J., 2021. What driving style makes pedestrians think a passing vehicle is driving automatically? Appl. Ergon. 95, 103428.

Berniker, M., Penny, S., 2019. A normative approach to neuromotor control. Biol. Cybernet. 113 (1), 83–92. http://dx.doi.org/10.1007/s00422-018-0777-7.

Berret, B., Verdel, D., Burdet, E., Jean, F., 2024. Co-contraction embodies uncertainty: An optimal feedforward strategy for robust motor control. PLoS Comput. Biol. 20 (11), e1012598. http://dx.doi.org/10.1371/journal.pcbi.1012598, URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1012598.

Botvinick, M., Toussaint, M., 2012. Planning as inference. Trends Cogn. Sci. 16 (10), 485–488.

Braun, D.A., Aertsen, A., Wolpert, D.M., Mehring, C., 2009. Learning optimal adaptation strategies in unpredictable motor tasks. J. Neurosci. 29 (20), 6472–6478. http://dx.doi.org/10.1523/JNEUROSCI.3075-08.2009.

Brucklacher, M., Pezzulo, G., Mannella, F., Galati, G., Pennartz, C.M., 2025. Learning to segment self-generated from externally caused optic flow through sensorimotor mismatch circuits. Neural Netw. 181, 106716.

Brüdigam, T., Ahmic, K., Leibold, M., Wollherr, D., 2018. Legible model predictive control for autonomous driving on highways. IFACPapersOnLine.

Busemeyer, J.R., Townsend, J.T., 1993. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. Psychol Rev 100 (3), 432.

Chen, Y., Rosolia, U., Ubellacker, W., Csomay-Shanklin, N., Ames, A.D., 2022. Interactive multi-modal motion planning with branch model predictive control. IEEE Robot. Autom. Lett. 7 (2).

Choi, J.K., Ji, Y.G., 2015. Investigating the importance of trust on adopting an autonomous vehicle. Int. J. Hum.-Comput. Interact. 31, 692–702. http://dx.doi.org/10.1080/10447318.2015.1070549, URL https://www.tandfonline.com/doi/abs/10.1080/10447318.2015.1070549.

Cisek, P., 2007. Cortical mechanisms of action selection: the affordance competition hypothesis. Phil. Trans. R. Soc. B 362 (1485), 1585.

Cover, T.M., Thomas, J.A., 2006. Elements of Information Theory, second ed. Wiley-Interscience, Hoboken, N.J.

Crevecoeur, F., Sepulchre, R.J., Thonnard, J.L., Lefèvre, P., 2011. Improving the state estimation for optimal control of stochastic processes subject to multiplicative noise. Automatica 47 (3), 591–596. http://dx.doi.org/10.1016/j.automatica.2011.01.026.

Cross, K.P., Cook, D.J., Scott, S.H., 2024. Rapid online corrections for proprioceptive and visual perturbations recruit similar circuits in primary motor cortex. Eneuro 11 (2), http://dx.doi.org/10.1523/ENEURO.0083-23.2024, ENEURO.0083-23.2024.

Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nature Neurosci. 8 (12), 1704–1711.

de Groot, O., Brito, B., Ferranti, L., Gavrila, D., Alonso-Mora, J., 2021. Scenario-based trajectory optimization in uncertain dynamic environments. IEEE Robot. Autom. Lett. 6 (3).

de Winkel, K.N., Katliar, M., Diers, D., Bülthoff, H.H., 2018. Causal inference in the perception of verticality. Sci. Rep. 8 (1), 5483. http://dx.doi.org/10.1038/s41598-018-23838-w, URL https://www.nature.com/articles/s41598-018-23838-w.

de Winter, J., Hoogmoed, J., Stapel, J., Dodou, D., Bazilinskyy, P., 2023. Predicting perceived risk of traffic scenes using computer vision. Transp. Res. Part F: Traffic Psychol. Behav. 93, 235–247. http://dx.doi.org/10.1016/J.TRF.2023.01.014, URL https://linkinghub.elsevier.com/retrieve/pii/S1369847823000177.

Dey, D., Habibovic, A., Löcken, A., Wintersberger, P., Pfleging, B., Riener, A., Martens, M., Terken, J., 2020. Taming the eHMI jungle: A classification taxonomy to guide, compare, and assess the design principles of automated vehicles' external human-machine interfaces. Transp. Res. Interdiscip. Perspect. 7, 100174.

Dey, D., Matviienko, A., Berger, M., Pfleging, B., Martens, M., Terken, J., 2021. Communicating the intention of an automated vehicle to pedestrians: The contributions of eHMI and vehicle behavior. It-Inform. Technol. 63 (2), 123–141.
Dieterich, R., Endrass, T., Kathmann, N., 2016. Uncertainty is associated with increased selective attention and sustained stimulus processing. Cogn. Affect., & Behav. Neurosci. 16, 447–456.
Dixit, A., Lindemann, L., Wei, S.X., Cleaveland, M., Pappas, G.J., Burdick, J.W., 2023. Adaptive conformal prediction for motion planning among dynamic agents. In: Learning for Dynamics and Control Conference. PMLR, pp. 300–314.
Douglass, J.K., Wilkens, L., Pantazelou, E., Moss, F., 1993. Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance. Nature 365 (6444), 337–340. http://dx.doi.org/10.1038/365337a0.
Dragan, A.D., Lee, K.C., Srinivasa, S.S., 2013. Legibility and predictability of robot motion. In: 2013 8th ACM/IEEE International Conference on Human-Robot Interaction. HRI, IEEE, pp. 301–308.
Echeveste, R., Aitchison, L., Hennequin, G., Lengyel, M., 2020. Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. Nature Neurosci. 23 (9), 1138–1149. http://dx.doi.org/10.1038/s41593-020-0671-1.
Ernst, M.O., Banks, M.S., 2002. Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415 (6870), 429–433. http://dx.doi.org/10.1038/415429a, URL https://www.nature.com/articles/415429a.
Faas, S.M., Mathis, L.-A., Baumann, M., 2020a. External HMI for self-driving vehicles: Which information shall be displayed? Transp. Res. Part F: Traffic Psychol. Behav. 68, 171–186.
Faas, S.M., Mattes, S., Kao, A.C., Baumann, M., 2020b. Efficient paradigm to measure street-crossing onset time of pedestrians in video-based interactions with vehicles. Inf. 11, 360. http://dx.doi.org/10.3390/INFO11070360, https://www.mdpi.com/2078-2489/11/7/360/htm, https://www.mdpi.com/2078-2489/11/7/360.
Faisal, A.A., Selen, L.P.J., Wolpert, D.M., 2008. Noise in the nervous system. Nature Rev. Neurosci. 9 (4), 292–303. http://dx.doi.org/10.1038/nrn2258.
Ferro, G.M., Sornette, D., 2020. Stochastic representation decision theory: How probabilities and values are entangled dual characteristics in cognitive processes. Plos One 15 (12), e0243661.
Fetsch, C.R., DeAngelis, G.C., Angelaki, D.E., 2010. Visual–vestibular cue integration for heading perception: applications of optimal cue integration theory. Eur. J. Neurosci. 31 (10), 1721–1729. http://dx.doi.org/10.1111/j.1460-9568.2010.07207.x, URL https://onlinelibrary.wiley.com/doi/10.1111/j.1460-9568.2010.07207.x.
Fleming, S.M., Daw, N.D., 2017. Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. Psychol Rev 124 (1), 91.
Franklin, D.W., Wolpert, D.M., 2011. Computational mechanisms of sensorimotor control. Neuron 72 (3), 425–442. http://dx.doi.org/10.1016/j.neuron.2011.10.006.
Franssen, M., Dong, H., Verstegen, R., Martens, M.H., 2024a. Exploring the effects of interior design elements on the uncertainty of passengers of fully automated cars. In: Accelerating Open Access Science in Human Factors Engineering and Human-Centered Computing. pp. 213–223. http://dx.doi.org/10.54941/ahfe1005211.
Franssen, M., Verstegen, R., Bazilinskyy, P., Martens, M.H., 2024b. Exploring the correlation between emotions and uncertainty in daily travel. In: Praetorius, G., Sellberg, C., Patriarca, R. (Eds.), Accelerating Open Access Science in Human Factors Engineering and Human-Centered Computing. In: Advances in Human Factors of Transportation, AHFE International, pp. 527–536. http://dx.doi.org/10.54941/ahfe1005244, 15th International Conference op Applied Human Factors and Ergonomics and the Affiliated Conferences, AHFE 2024, AHFE 2024 ; Conference date: 24-07-2024 Through 27-07-2024.
Franssen, M., Verstegen, R., Bazilinskyy, P., van der Spek, E.D., Martens, M.H., 2026. Psychological uncertainty in human experience and behavior: a systematic review on measurements, induction, and effects. Front. Psychol. 17, 1844980. http://dx.doi.org/10.3389/FPSYG.2026.1844980, URL https://www.frontiersin.org/articles/10.3389/fpsyg.2026.1844980/full.
Friston, K., 2005. A theory of cortical responses. Phil. Trans. R. Soc. B 360 (1456), 815–836. http://dx.doi.org/10.1098/rstb.2005.1622, URL https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2005.1622.
Gal, Y., Ghahramani, Z., 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: ICML. PMLR, pp. 1050–1059.
Gandhi, M.S., Vlahov, B., Gibson, J., William, G., Theodorou1, E.A., 2021. Robust model predictive path integral control: Analysis and performance guarantees. IEEE Robot. Autom. Lett.
Gawlikowski, J., Tassi, C.R.N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., et al., 2021. A survey of uncertainty in deep neural networks. arXiv preprint arXiv:2107.03342.
Geisslinger, M., Poszler, F., Betz, J., Lütge, C., Lienkamp, M., 2021. Autonomous driving ethics: From trolley problem to ethics of risk. Philos. Technol.
Geisslinger, M., Poszler, F., Lienkamp, M., 2023. An ethical trajectory planning algorithm for autonomous vehicles. Nat. Mach. Intell.
Geldenbott, J., Leung., K., 2024. Legible and proactive robot planning for prosocial human-robot interactions. IEEE Int. Conf. Robot. Autom.
Ghanem, R., Higdon, D., Owhadi, H., et al., 2017. Handbook of Uncertainty Quantification. vol. 6, Springer New York.
Ghavamzadeh, M., Mannor, S., Pineau, J., Tamar, A., et al., 2015. Bayesian reinforcement learning: A survey. Found. Trends® Mach. Learn. 8 (5–6), 359–483.
Gil, R.C., Ornia, D.J., Mustafa, K.A., Alonso-Mora, J., 2025. Predictability awareness for efficient and robust multi-agent coordination. Int. Conf. Auton. Agents Multiagent Syst.
Griffiths, T.L., Chater, N., Tenenbaum, J.B., 2024. Bayesian Models of Cognition: Reverse Engineering the Mind. The MIT Press, Cambridge, Massachusetts.
Gros, C., Kester, L., Martens, M., Werkhoven, P., 2025. Addressing ethical challenges in automated vehicles: bridging the gap with hybrid AI and augmented utilitarianism. AI Ethics 5 (3), 2757–2770.
Grupe, D.W., Nitschke, J.B., 2013. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. Nat. Rev. Neurosci. 14 (7), 488–501. http://dx.doi.org/10.1038/nrn3524, URL https://www.nature.com/articles/nrn3524.
Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A., 2018. Social gan: Socially acceptable trajectories with generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2255–2264.
Gupta, M.M., Solo, A.M., 2010. On the morphology of uncertainty in human perception and cognition. In: Proceedings of the First Interdisciplinary CHESS Interactions Conference. World Scientific, pp. 257–271.
Haefner, R.M., Beck, J., Savin, C., Salmasi, M., Pitkow, X., 2024. How does the brain compute with probabilities?. http://dx.doi.org/10.48550/arXiv.2409.02709, arXiv preprint arXiv:2409.02709, URL http://arxiv.org/abs/2409.02709, Publisher: arXiv.
Hakobyan, A., Kim, G.C., Yang, I., 2019. Risk-aware motion planning and control using cvar-constrained optimization. IEEE Robot. Autom. Lett. 4 (4).
Harris, C.M., Wolpert, D.M., 1998. Signal-dependent noise determines motor planning. Nature 394 (6695), 780–784. http://dx.doi.org/10.1038/29528.
Hetherington, N.J., Croft, E.A., Van der Loos, H.M., 2021. Hey robot, which way are you going? Nonverbal motion legibility cues for human-robot spatial interaction. IEEE Robot. Autom. Lett. 6 (3), 5010–5015.
Hoff, K.A., Bashir, M., 2015. Trust in automation: Integrating empirical evidence on factors that influence trust. Hum. Factors 57 (3), 407–434.
Izawa, J., Rane, T., Donchin, O., Shadmehr, R., 2008. Motor adaptation as a process of reoptimization. J. Neurosci. 28 (11), 2883–2891. http://dx.doi.org/10.1523/JNEUROSCI.5359-07.2008.
Jones, K.E., Hamilton, A.F.d.C., Wolpert, D.M., 2002. Sources of signal-dependent noise during isometric force production. J. Neurophysiol. 88 (3), 1533–1544.
Kaelbling, L.P., Littman, M.L., Cassandra, A.R., 1998. Planning and acting in partially observable stochastic domains. Artificial Intelligence 101 (1–2), 99–134.
Kalantari, A.H., 2023. Modelling vehicle-pedestrian interactions at unsignalised locations employing game-theoretic models (Ph.D. thesis). University of Leeds.
Kalidindi, H.T., Cross, K.P., Lillicrap, T.P., Omrani, M., Falotico, E., Sabes, P.N., Scott, S.H., 2021. Rotational dynamics in motor cortex are consistent with a feedback controller. ELife 10, e67256. http://dx.doi.org/10.7554/eLife.67256.
Kappes, A., Nussberger, A.M., Faber, N.S., Kahane, G., Savulescu, J., Crockett, M.J., 2018. Uncertainty about the impact of social decisions increases prosocial behaviour. Nat. Hum. Behav. 2, 573–580. http://dx.doi.org/10.1038/S41562-018-0372-X, URL https://pubmed.ncbi.nlm.nih.gov/31209312/.
Kasuga, S., Crevecoeur, F., Cross, K.P., Balalaie, P., Scott, S.H., 2022. Integration of proprioceptive and visual feedback during online control of reaching. J. Neurophysiol. 127 (2), 354–372. http://dx.doi.org/10.1152/jn.00639.2020, URL https://journals.physiology.org/doi/full/10.1152/jn.00639.2020.
Kawasaki, A., Seki, A., 2021. Multimodal trajectory predictions for autonomous driving without a detailed prior map. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3723–3732.
Kendall, A., Gal, Y., 2017. What uncertainties do we need in bayesian deep learning for computer vision? Adv. Neural Inf. Process. Syst. 30.
Kirchmair, L., Paulo, N., 2023. Taking ethics seriously in AV trajectory planning algorithms. Nat. Mach. Intell.
Knight, F.H., 1921. Risk, uncertainty and profit. Hart, Schaffner Marx.
Knill, D.C., Pouget, A., 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci. 27 (12), 712–719. http://dx.doi.org/10.1016/j.tins.2004.10.007, URL https://linkinghub.elsevier.com/retrieve/pii/S0166223604003352.
Knill, D.C., Richards, W. (Eds.), 1996. Perception as Bayesian Inference. Cambridge University Press, Cambridge, URL https://www.cambridge.org/core/books/perception-as-bayesian-inference/0442F577F5E4CD874FA6819978574C8F.
Kohn, S.C., De Visser, E.J., Wiese, E., Lee, Y.-C., Shaw, T.H., 2021. Measurement of trust in automation: A narrative review and reference guide. Front. Psychol. 12, 604977.
Körding, K.P., Beierholm, U., Ma, W.J., Quartz, S., Tenenbaum, J.B., Shams, L., 2007. Causal inference in multisensory perception. PLOS ONE 2 (9), e943. http://dx.doi.org/10.1371/journal.pone.0000943, URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0000943.
Körding, K.P., Wolpert, D.M., 2004. Bayesian integration in sensorimotor learning. Nature 427 (6971), 244–247. http://dx.doi.org/10.1038/nature02169, URL https://www.nature.com/articles/nature02169.
Lakshminarayanan, B., Pritzel, A., Blundell, C., 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. NIPS 30.
Lee, T.S., Mumford, D., 2003. Hierarchical Bayesian inference in the visual cortex. J. Opt. Soc. Amer. A 20 (7), 1434–1448.

Lee, J.D., See, K.A., 2004. Trust in automation: Designing for appropriate reliance. Hum. Factors 46 (1), 50–80. http://dx.doi.org/10.1518/hfes.46.1.50_30392, PMID: 15151155.

Lew, T., Bonalli, R., Pavone, M., 2023. Risk-averse trajectory optimization via sample average approximation. IEEE Robot. Autom. Lett. 9 (2).

Lindemann, L., Cleaveland, M., Shim, G., Pappas, G.J., 2023. Safe planning in dynamic environments using conformal prediction. IEEE Robot. Autom. Lett.

Lindley, D.V., 2013. Understanding Uncertainty. John Wiley & Sons.

Lingam, S.N., Franssen, M., Petermeijer, S.M., Martens, M., 2024a. Challenges and future directions for human-drone interaction research: An expert perspective. Int. J. Hum.-Comput. Interact. 41 (12), 7905–7921. http://dx.doi.org/10.1080/10447318.2024.2400756.

Lingam, S.N., Petermeijer, S.M., Obaid, M., Martens, M., 2025a. DroDelivery: User-centered concepts for intent-based communications of a delivery drone in public spaces. In: Adjunct Proceedings of the 27th International Conference on Mobile Human-Computer Interaction. In: MobileHCI '25 Adjunct, Association for Computing Machinery, New York, NY, USA, pp. 1–9. http://dx.doi.org/10.1145/3737821.3749600.

Lingam, S.N., Petermeijer, S.M., Obaid, M., Martens, M., 2025b. Towards public-drone interactions: Communicating delivery intentions to recipients and bystanders. In: Ardito, C., Diniz Junqueira Barbosa, S., Conte, T., Freire, A., Gasparini, I., Palanque, P., Prates, R. (Eds.), Human-Computer Interaction – INTERACT 2025. Springer Nature Switzerland, Cham, pp. 317–321.

Lingam, S.N., Petermeijer, S.M., Obaid, M., Martens, M., 2026. Delivery from the sky: investigating visual cues to communicate robot intentions in simulated public spaces. Sci. Rep.

Lingam, S.N., Petermeijer, S.M., Torre, I., Bazilinskyy, P., Ljungblad, S., Martens, M., 2025c. Behavioral effects of a delivery drone on feelings of uncertainty: A virtual reality experiment. ACM Trans. Human-Robot Interact. 14 (4), http://dx.doi.org/10.1145/3729538.

Lingam, S.N., Verstegen, R., Petermeijer, S.M., Martens, M., 2024b. Design of delivery drone and the interaction: A public user perspective. In: Proceedings of the 12th International Conference on Human-Agent Interaction. pp. 356–357. http://dx.doi.org/10.1145/3687272.3690880.

Lingam, S.N., Verstegen, R., Petermeijer, S.M., Martens, M., 2025d. Human interactions with delivery drones in public spaces: design recommendations from recipient and bystander perspectives. Front. Robot. AI 12, http://dx.doi.org/10.3389/frobt.2025.1580289, URL https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2025.1580289.

Liu, J.-Y., Cooke, J.R., Selen, L.J., Medendorp, W.P., 2025. Visuoinertial and visual feedback in online steering control. PLoS Comput. Biol. 21 (8), e1012659.

Longtin, A., 1993. Stochastic resonance in neuron models. J. Stat. Phys. 70 (1), 309–327. http://dx.doi.org/10.1007/BF01053970.

Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A., 2006. Bayesian inference with probabilistic population codes. Nature Neurosci. 9 (11), 1432–1438. http://dx.doi.org/10.1038/nn1790.

Majumdar, A., Pavone, M., 2020. How should a robot assess risk? towards an axiomatic theory of risk in robotics. In: Robotics Research: The 18th International Symposium ISRR. Springer, pp. 75–84.

Mante, V., Sussillo, D., Shenoy, K.V., Newsome, W.T., 2013. Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature 503 (7474), 78–84. http://dx.doi.org/10.1038/nature12742.

Markkula, G., Madigan, R., Nathanael, D., Portouli, E., Lee, Y.M., Dietrich, A., Billington, J., Schieben, A., Merat, N., 2020. Defining interactions: A conceptual framework for understanding interactive behaviour in human and automated road traffic. Theor. Issues Ergon. Sci. 21 (6), 728–752.

Marthe, A., Bounan, S., Garivier, A., Vernade, C., 2025. Efficient risk-sensitive planning via entropic risk measures. arXiv preprint arXiv:2502.20423.

Mavrogiannis, C., Baldini, F., Wang, A., Zhao, D., Trautman, P., Steinfeld, A., Oh, J., 2023. Core challenges of social robot navigation: A survey. ACM Trans. Hum.-Robot Interact. 12 (3), 1–39.

Mayer, R.C., Davis, J.H., Schoorman, F.D., 1995. An integrative model of organizational trust. Acad. Manag. Rev. 20, 709. http://dx.doi.org/10.2307/258792.

Medendorp, W.P., Heed, T., 2019. State estimation in posterior parietal cortex: Distinct poles of environmental and bodily states. Prog. Neurobiol. 183, 101691. http://dx.doi.org/10.1016/j.pneurobio.2019.101691.

Mészáros, A., Schumann, J.F., Alonso-Mora, J., Zgonnikov, A., Kober, J., 2024. TrajFlow: Learning distributions over trajectories for human behavior prediction. In: 2024 IEEE Intelligent Vehicles Symposium. IV, IEEE, pp. 184–191.

Mitchell, V.-W., 1999. Consumer perceived risk: conceptualisations and models. Eur. J. Mark. 33, 163–195. http://dx.doi.org/10.1108/03090569910249229/FULL/XML.

Mitrovic, D., Klanke, S., Osu, R., Kawato, M., Vijayakumar, S., 2010. A computational model of limb impedance control based on principles of internal model uncertainty. PLOS ONE 5 (10), e13601. http://dx.doi.org/10.1371/journal.pone.0013601, URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0013601.

Moore, J.B., Zhou, X.Y., Lim, A.E., 1999. Discrete time LQG controls with control dependent noise. Systems Control Lett. 36 (3), 199–206. http://dx.doi.org/10.1016/S0167-6911(98)00092-9.

Morriss, J., Tupitsa, E., Dodd, H.F., Hirsch, C.R., 2022. Uncertainty makes me emotional: Uncertainty as an elicitor and modulator of emotional states. Front. Psychol. 13, 777025. http://dx.doi.org/10.3389/FPSYG.2022.777025/BIBTEX, URL https://www.frontiersin.org.

Mustafa, K.A., de Groot, O., Wang, X., Kober, J., Alonso-Mora, J., 2023. Probabilistic risk assessment for chance-constrained collision avoidance in uncertain dynamic environments. IEEE Int. Conf. Robot. Autom..

Muthugala, M.V.J., Jayasekara, A.B.P., 2016. Enhancing human-robot interaction by interpreting uncertain information in navigational commands based on experience and environment. In: 2016 IEEE International Conference on Robotics and Automation. ICRA, IEEE, pp. 2915–2921.

Nagengast, A.J., Braun, D.A., Wolpert, D.M., 2010. Risk-sensitive optimal feedback control accounts for sensorimotor behavior under uncertainty. PLoS Comput. Biol. 6 (7), e1000857. http://dx.doi.org/10.1371/journal.pcbi.1000857, URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1000857.

Nair, S.H., Govindarajan, V., Lin, T., Meissen, C., Tseng, H.E., Borrelli, F., 2022. Stochastic MPC with multi-modal predictions for traffic intersections. IEEE Int. Conf. Intell. Transp. Syst.

Narayanan, S., Moslemi, R., Pittaluga, F., Liu, B., Chandraker, M., 2021. Divide-and-conquer for lane-aware diverse trajectory prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15799–15808.

Nash, C., Cole, D., 2020. Identification and validation of a driver steering control model incorporating human sensory dynamics. Veh. Syst. Dyn. 58 (4), 495–517.

Nasr Azadani, M., Boukerche, A., 2023. A social-aware vehicle path forecasting method using graph neural networks. In: Proceedings of the Int'L ACM Symposium on Design and Analysis of Intelligent Vehicular Networks and Applications. pp. 61–68.

Neumeier, M., Dorn, S., Botsch, M., Utschick, W., 2024. Reliable trajectory prediction and uncertainty quantification with conditioned diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3461–3470.

Niedoba, M., Lavington, J., Liu, Y., Lioutas, V., Sefas, J., Liang, X., Green, D., Dabiri, S., Zwartsenberg, B., Scibior, A., et al., 2023. A diffusion-model of joint interactive navigation. Adv. Neural Inf. Process. Syst. 36, 55995–56011.

Noel, J.-P., Bill, J., Ding, H., Vastola, J., DeAngelis, G.C., Angelaki, D.E., Drugowitsch, J., 2023. Causal inference during closed-loop navigation: parsing of self-and object-motion. preprint, Neuroscience, URL http://biorxiv.org/lookup/doi/10.1101/2023.01.27.525974.

Nyberg, T., Pek, C., Col, L.D., Noren, C., Tumova, J., 2021. Risk-aware motion planning for autonomous vehicles with safety specifications. Intell. Veh. Symposium.

Orbán, G.o., Berkes, P., Fiser, J., Lengyel, M., 2016. Neural variability and sampling-based probabilistic representations in the visual cortex. Neuron 92 (2), 530–543. http://dx.doi.org/10.1016/j.neuron.2016.09.038, URL https://www.sciencedirect.com/science/article/pii/S0896627316306390.

Orhan, A.E., Ma, W.J., 2017. Efficient probabilistic inference in generic neural networks trained with non-probabilistic feedback. Nat. Commun. 8 (1), 138.

Parasuraman, R., Sheridan, T.B., Wickens, C.D., 2000. A model for types and levels of human interaction with automation. IEEE Trans. Syst. Man Cybern.-Part A: Syst. Humans 30 (3), 286–297.

Parr, T., Pezzulo, G., Friston, K.J., 2022. Active Inference: the Free Energy Principle in Mind, Brain, and Behavior. MIT Press.

Pearson, M.J., Dora, S., Struckmeier, O., Knowles, T.C., Mitchinson, B., Tiwari, K., Kyrki, V., Bohte, S., Pennartz, C.M., 2021. Multimodal representation learning for place recognition using deep hebbian predictive coding. Front. Robot. AI 8, 732023.

Pezzulo, G., Friston, K.J., 2019. The value of uncertainty: An active inference perspective. Behav. Brain Sci. 42.

Pezzulo, G., Van der Meer, M.A., Lansink, C.S., Pennartz, C.M., 2014. Internally generated sequences in learning and executing goal-directed behavior. Trends Cogn. Sci. 18 (12), 647–657.

Pezzulo, G., Rigoli, F., Friston, K.J., 2018. Hierarchical active inference: a theory of motivated control. Trends Cogn. Sci. 22 (4), 294–306.

Pruszynski, J.A., Kurtzer, I., Nashed, J.Y., Omrani, M., Brouwer, B., Scott, S.H., 2011. Primary motor cortex underlies multi-joint integration for fast feedback control. Nature 478 (7369), 387–390. http://dx.doi.org/10.1038/nature10436.

Rao, R.P.N., Ballard, D.H., 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nature Neurosci. 2 (1), 79–87. http://dx.doi.org/10.1038/4580, URL https://www.nature.com/articles/nn0199_79.

Rasouli, A., Tsotsos, J.K., 2019. Autonomous vehicles that interact with pedestrians: A survey of theory and practice. IEEE Trans. Intell. Transp. Syst. 21 (3), 900–918.

Ren, K., Chen, C., Sung, H., Ahn, H., Mitchell, I.M., Kamgarpour, M., 2024. Recursively feasible chance-constrained model predictive control under Gaussian mixture model uncertainty. IEEE Trans. Control. Syst. Technol.

Rohe, T., Ehlis, A.-C., Noppeney, U., 2019. The neural dynamics of hierarchical Bayesian causal inference in multisensory perception. Nat. Commun. 10 (1), 1907. http://dx.doi.org/10.1038/s41467-019-09664-2.

Rose, A.M., Rose, J.M., Rotaru, K., Sanderson, K.-A., Thibodeau, J.C., 2022. Effects of uncertainty visualization on attention, arousal, and judgment. Behav. Res. Account. 34 (1), 113–139.

Rowe, W.D., 1994. Understanding uncertainty. Risk Anal. 14, 743–750. http://dx.doi.org/10.1111/J.1539-6924.1994.TB00284.X, https://onlinelibrary.wiley.com/doi/full/10.1111/j.1539-6924.1994.tb00284.x, https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1539-6924.1994.tb00284.x, https://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924.1994.tb00284.x.
Rusu, S.I., Pennartz, C.M., 2020. Learning, memory and consolidation mechanisms for behavioral control in hierarchically organized cortico-basal ganglia systems. Hippocampus 30 (1), 73–98.
Sadigh, D., Sastry, S., Seshia, S.A., Dragan, A.D., 2016. Planning for autonomous cars that leverage effects on human actions. In: Robotics: Science and Systems. Vol. 2, Ann Arbor, MI, USA, pp. 1–9.
SAE International, 2021. J3016: Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. https://www.sae.org/standards/content/j3016_202104/.
Sahani, M., Dayan, P., 2003. Doubly distributional population codes: Simultaneous representation of uncertainty and multiplicity. Neural Comput. 15 (10), 2255–2279. http://dx.doi.org/10.1162/089976603322362356.
Salzmann, T., Ivanovic, B., Chakravarty, P., Pavone, M., 2020. Trajectron++: Dynamically-Feasible Trajectory Forecasting with Heterogeneous Data. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (Eds.), Computer Vision – ECCV 2020. Springer International Publishing, Cham, pp. 683–700. http://dx.doi.org/10.1007/978-3-030-58523-5_40.
Sarter, N.B., Woods, D.D., Billings, C., Salvendy, C., 1997. Automation surprises. h. In: Andbook of Human Factors and Ergonomics. Wiley.
Schaefer, K.E., Chen, J.Y., Szalma, J.L., Hancock, P.A., 2016. A meta-analysis of factors influencing the development of trust in automation: Implications for understanding autonomy in future systems. Hum. Factors 58 (3), 377–400.
Schlag, J., Schlag-Rey, M., 2002. Through the eye, slowly; Delays and localization errors in the visual system. Nature Rev. Neurosci. 3 (3), 191. http://dx.doi.org/10.1038/nrn750.
Schöller, C., Knoll, A., 2021. Flomo: Tractable motion prediction with normalizing flows. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS, IEEE, pp. 7977–7984.
Schwarting, W., Pierson, A., Alonso-Mora, J., Karaman, S., Rus, D., 2019. Social behavior for autonomous vehicles. Proc. Natl. Acad. Sci. 116 (50), 24972–24978.
Scott, S.H., 2004. Optimal feedback control and the neural basis of volitional motor control. Nature Rev. Neurosci. 5 (7), 532–545. http://dx.doi.org/10.1038/nrn1427.
Shi, S., Jiang, L., Dai, D., Schiele, B., 2022. Motion transformer with global intention localization and local movement refinement. Adv. Neural Inf. Process. Syst. 35, 6531–6543.
Shi, L., Wang, L., Zhou, S., Hua, G., 2023. Trajectory unified transformer for pedestrian trajectory prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. IEEE Computer Society, Los Alamitos, CA, USA, pp. 9675–9684. http://dx.doi.org/10.1109/ICCV51070.2023.00887.
Simon, H.A., 1955. A behavioral model of rational choice. Q. J. Econ. 99–118.
Sobel, M.J., 1982. The variance of discounted Markov decision processes. J. Appl. Probab. 19 (4), 794–802.
Song, H.F., Yang, G.R., Wang, X.-J., 2016. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: A simple and flexible framework. PLoS Comput. Biol. 12 (2), e1004792. http://dx.doi.org/10.1371/journal.pcbi.1004792.
Stavisky, S.D., Kao, J.C., Ryu, S.I., Shenoy, K.V., 2017. Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. Neuron 95 (1), 195–208.e9. http://dx.doi.org/10.1016/j.neuron.2017.05.023.
Stengel, R.F., 1994. Optimal control and estimation. In: Dover Books on Advanced Mathematics, Dover Publications, New York.
Suk, H., Kim, S., 2025. Uncertainty-aware multimodal trajectory prediction via a single inference from a single model. Sensors (Basel, Switzerland) 25 (1), 217.
Sun, T., Bohté, S., 2025. Uncertainty-aware spiking neural networks for regression. In: Submitted. pp. XX–YY.
Sun, J., Wang, Z., Li, J., Lu, C., 2021. Unified and fast human trajectory prediction via conditionally parameterized normalizing flow. IEEE Robot. Autom. Lett. 7 (2), 842–849.
Sun, T., Yin, B., Bohté, S., 2023. Efficient uncertainty estimation in spiking neural networks via MC-dropout. In: International Conference on Artificial Neural Networks. Springer, pp. 393–406.
Sunberg, Z., Kochenderfer, M., 2018. Online algorithms for POMDPs with continuous state, action, and observation spaces. In: Proceedings of the International Conference on Automated Planning and Scheduling. Vol. 28, pp. 259–263.
Tanabe, S., 2013. Population codes in the visual cortex. Neurosci. Res. 76 (3), 101–105. http://dx.doi.org/10.1016/j.neures.2013.03.010.
Tian, K., Markkula, G., Wei, C., Lee, Y.M., Madigan, R., Hirose, T., Merat, N., Romano, R., 2023. Deconstructing pedestrian crossing decisions in interactions with continuous traffic: An anthropomorphic model. IEEE Trans. Intell. Transp. Syst. 25 (3), 2466–2478.
Todorov, E., 2005. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. Neural Comput. 17 (5), 1084–1108. http://dx.doi.org/10.1162/0899766053491887.
Todorov, E., Jordan, M.I., 2002. Optimal feedback control as a theory of motor coordination. Nature Neurosci. 5 (11), 1226–1235. http://dx.doi.org/10.1038/nn963.
Tolksdorf, L., Tejada, A., van de Wouw, N., Birkner, C., 2023. Risk in stochastic and robust model predictive path-following control for vehicular motion planning. IEEE Intell. Veh. Symposium.
Torres, J.J., Marro, J., Mejias, J.F., 2011. Can intrinsic noise induce various resonant peaks? New J. Phys. 13 (5), 053014. http://dx.doi.org/10.1088/1367-2630/13/5/053014.
van Beers, R.J., Haggard, P., Wolpert, D.M., 2004. The role of execution noise in movement variability. J. Neurophysiol. 91 (2), 1050–1063. http://dx.doi.org/10.1152/jn.00652.2003.
van Beers, R.J., Sittig, A.C., Gon, J.J.D.v.d., 1999. Integration of proprioceptive and visual position-information: An experimentally supported model. J. Neurophysiol. 81 (3), 1355–1364. http://dx.doi.org/10.1152/jn.1999.81.3.1355, URL https://journals.physiology.org/doi/full/10.1152/jn.1999.81.3.1355.
van Bergen, R.S., Ji Ma, W., Pratte, M.S., Jehee, J.F.M., 2015. Sensory uncertainty decoded from visual cortex predicts behavior. Nature Neurosci. 18 (12), 1728–1730. http://dx.doi.org/10.1038/nn.4150.
Varadarajan, B., Hefny, A., Srivastava, A., Refaat, K.S., Nayakanti, N., Cornman, A., Chen, K., Douillard, B., Lam, C.P., Anguelov, D., et al., 2022. Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction. In: 2022 IEEE International Conference on Robotics and Automation (ICRA). IEEE Computer Society, Los Alamitos, CA, USA, pp. 7814–7821. http://dx.doi.org/10.1109/ICRA46639.2022.9812107.
Victor, T.W., Tivesten, E., Gustavsson, P., Johansson, J., Sangberg, F., Ljung Aust, M., 2018. Automation expectation mismatch: Incorrect prediction despite eyes on threat and hands on wheel. Hum. Factors 60 (8), 1095–1116.
von Hünerbein, B., Jordan, J., Oude Lohuis, M., Marchesi, P., Olcese, U., Pennartz, C.M., Senn, W., Petrovici, M.A., 2026. Increased perceptual reliability reduces membrane potential variability in cortical neurons. J. Neurosci. 46 (1).
Walker, E.Y., Cotton, R.J., Ma, W.J., Tolias, A.S., 2020. A neural basis of probabilistic computation in visual cortex. Nature Neurosci. 23 (1), 122–129. http://dx.doi.org/10.1038/s41593-019-0554-5.
Wang, C., Wang, Y., Xu, M., Crandall, D.J., 2022. Stepwise goal-driven networks for trajectory prediction. IEEE Robot. Autom. Lett. 7 (2), 2716–2723.
Whittle, P., 1981. Risk-sensitive linear/quadratic/gaussian control. Adv. in Appl. Probab. 13 (4), 764–777. http://dx.doi.org/10.2307/1426972, URL https://www.cambridge.org/core/journals/advances-in-applied-probability/article/abs/risksensitive-linearquadraticgaussian-control/9D29B8E7D13589A3823E75525181117F.
Wickens, C.D., Helton, W.S., Hollands, J.G., Banbury, S., 2021. Engineering psychology and human performance. Routledge.
Wierda, T.S., Dora, S., Pennartz, C.M., Mejias, J.F., 2025. Diverse and flexible behavioral strategies arise in recurrent neural networks trained on multisensory decision making. PLoS Comput. Biol. 21 (10), e1013559.
Windschitl, P.D., Wells, G.L., 1996. Measuring psychological uncertainty: Verbal versus numeric methods. J. Exp. Psychol.: Appl. 2, 343–364. http://dx.doi.org/10.1037/1076-898X.2.4.343, URL /doiLanding?doi=10.1037%2F1076-898X.2.4.343.
Wojciechowska, A., Frey, J., Sass, S., Shafir, R., Cauchard, J.R., 2019. Collocated human-drone interaction: Methodology and approach strategy. In: 2019 14th ACM/IEEE International Conference on Human-Robot Interaction. HRI, IEEE, pp. 172–181.
Wolpert, D.M., Landy, M.S., 2012. Motor control is decision-making. Curr. Opin. Neurobiol. 22 (6), 996–1003. http://dx.doi.org/10.1016/j.conb.2012.05.003, URL https://linkinghub.elsevier.com/retrieve/pii/S0959438812000827.
Yang, S.C.-H., Lengyel, M., Wolpert, D.M., 2016. Active sensing in the categorization of visual patterns. In: Tsao, D.Y. (Ed.), ELife 5, e12215. http://dx.doi.org/10.7554/eLife.12215, URL https://doi.org/10.7554/eLife.12215.
Yeh, A., Ratsamee, P., Kiyokawa, K., Uranishi, Y., Mashita, T., Takemura, H., Fjeld, M., Obaid, M., 2017. Exploring proxemics for human-drone interaction. In: Proceedings of the 5th International Conference on Human Agent Interaction. pp. 81–88.
Yin, J., Zhang, Z., Tsiotras, P., 2022. Risk-aware model predictive path integral control using conditional value-at-risk. IEEE Int. Conf. Robot. Autom.
Zhi, Y., Bao, Z., Zhang, S., He, R., 2021. Bigru based online multi-modal driving maneuvers and trajectory prediction. Proc. Inst. Mech. Eng. D: J. Automob. Eng. 235 (14), 3431–3441.
Zhou, J., Gao, Y., Johansson, O., Olofsson, B., Frisk, E., 2025. Robust predictive motion planning by learning obstacle uncertainty. IEEE Trans. Control. Syst. Technol.
Zhou, Z., Ye, L., Wang, J., Wu, K., Lu, K., 2022. Hivt: Hierarchical vector transformer for multi-agent motion prediction. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE Computer Society, Los Alamitos, CA, USA, pp. 8823–8833. http://dx.doi.org/10.1109/CVPR52688.2022.00862.