

**Document Version**

Final published version

**Licence**

Dutch Copyright Act (Article 25fa)

**Citation (APA)**

Suryana, L. E., Rahmani, S., Calvert, S. C., Zgonnikov, A., & Van Arem, B. (2025). A Human Reasons-based Supervision Framework for Ethical Decision-Making in Automated Vehicles. In C. Laugier, A. Renzaglia, N. Atanasov, S. Birchfield, G. Cielniak, L. De Mattos, L. Fiorini, P. Giguere, K. Hashimoto, J. Ibanez-Guzman, T. Kamegawa, J. Lee, G. Loianno, K. Luck, H. Maruyama, P. Martinet, H. Moradi, U. Nunes, J. Pettre, A. Pretto, T. Ranzani, A. Ronnau, S. Rossi, E. Rouse, F. Ruggiero, O. Simonin, D. Wang, M. Yang, E. Yoshida, ... H. Zhao (Eds.), *IROS 2025 - 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems, Conference Proceedings* (pp. 21495-21502). (IEEE International Conference on Intelligent Robots and Systems). IEEE. <https://doi.org/10.1109/IROS60139.2025.11245868>

**Important note**

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

Please check the document version above.

**Copyright**

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

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

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.

We will remove access to the work immediately and investigate your claim.

# A Human Reasons-based Supervision Framework for Ethical Decision-Making in Automated Vehicles

Lucas Elbert Suryana<sup>1,3</sup>, Saeed Rahmani<sup>1</sup>, Simeon C. Calvert<sup>1,3</sup>, Arkady Zgonnikov<sup>2,3</sup>, Bart van Arem<sup>1</sup>

**Abstract**—Ethical dilemmas are a common challenge in everyday driving, requiring human drivers to balance competing priorities such as safety, efficiency, and rule compliance. However, much of the existing research in automated vehicles (AVs) has focused on high-stakes “trolley problems,” which involve extreme and rare situations. Such scenarios, though rich in ethical implications, are rarely applicable in real-world AV decision-making. In practice, when AVs confront everyday ethical dilemmas, they often appear to prioritise strict adherence to traffic rules. By contrast, human drivers may bend the rules in context-specific situations, using judgement informed by practical concerns such as safety and efficiency. According to the concept of meaningful human control, AVs should respond to human reasons, including those of drivers, vulnerable road users, and policymakers. This work introduces a novel human reasons-based supervision framework that detects when AV behaviour misaligns with expected human reasons to trigger trajectory reconsideration. The framework integrates with motion planning and control systems to support real-time adaptation, enabling decisions that better reflect safety, efficiency, and regulatory considerations. Simulation results demonstrate that this approach could help AVs respond more effectively to ethical challenges in dynamic driving environments by prompting replanning when the current trajectory fails to align with human reasons. These findings suggest that our approach offers a path toward more adaptable, human-centered decision-making in AVs.

## I. INTRODUCTION

Addressing the ethical complexities that emerge in daily driving contexts remains essential for social acceptance of automated vehicles (AVs). Despite their promised advantages in safety improvements and transportation access [1], the widespread adoption of these systems hinges on their capacity to reflect human ethical judgment, particularly when confronting morally ambiguous situations where multiple values compete [2], [3]—situations commonly referred to as ethical dilemmas. Examples include deciding whether to briefly occupy the opposite lane to safely overtake cyclists [4], or speeding up temporarily to avoid unsafe situations. This behaviour reveals a critical gap in AV decision-making: the necessity of designing AVs capable of dynamically balancing multiple considerations, such as safety, efficiency, regulatory compliance, and contextual appropriateness, in real-time, rather than relying solely on predefined regulations.

Current approaches show limitations when it comes to handling everyday ethical dilemmas in automated driving.

<sup>1</sup>Department of Transport and Planning, <sup>2</sup>Department of Cognitive Robotics, <sup>3</sup>Centre for Meaningful Human Control, Delft University of Technology, 2628 CN Delft, The Netherlands. l.e.suryana@tudelft.nl; s.rahmani@tudelft.nl; s.c.calvert@tudelft.nl; b.vanarem@tudelft.nl; a.zgonnikov@tudelft.nl.

Much of the research has focused on extreme scenarios, such as the well-known “trolley problem” [5]. While philosophically significant, trolley problems rarely happen in daily driving, and despite the practical importance of everyday dilemmas, they often receive less attention [1], [6]. [2] emphasises that everyday ethical decisions in automated driving extend far beyond these extreme scenarios, requiring nuanced contextual use of reasons, something current systems lack. Similarly, [7] argue that framing AV ethics as simplified trolley problems fails to capture the probabilistic nature and dynamic complexity of real-world driving situations.

Although the ethical dimensions of “mundane” scenarios may seem straightforward for human drivers, they require context-aware judgment that comes naturally to human but poses challenge for AVs. These judgments must balance multiple ethically relevant considerations, such as safety, efficiency, and social norms. This adaptability represents a challenge for AV systems designed with traditional motion planning algorithms, which primarily optimise for trajectory smoothness and collision avoidance without explicitly integrating ethical considerations. Building on this understanding, recent ethical frameworks propose more holistic approaches that better align with human moral intuitions. [8] and [9] collectively emphasise that effective AV ethics must integrate moral principles such as deontological ethics, virtue ethics, and consequentialist considerations while maintaining transparency in decision-making.

However, integrating these ethical principles into AV decision-making remains challenging. While prior works have focused on embedding such principles into control and motion algorithms [1], [10], [11], current approaches fail to make explicit when these principles are in conflict. Recognising these conflicts is essential for enabling transparent decisions and for adjusting AV behaviour to better reflect the ethical principles the system is intended to uphold.

The concept of meaningful human control (MHC) [12], [13] offers a promising conceptual bridge for addressing the challenge of making explicit which moral principles are in conflict. MHC asserts that humans should ultimately be responsible for every decision made by automated systems. [13] laid the groundwork for achieving MHC. One of the required conditions is the tracking condition, which requires automated systems to respond to the reasons of relevant humans. In the remainder of this paper, we refer to these relevant humans – such as drivers, passengers, pedestrians, and policymakers – as stakeholders. According to [12], these reasons can be understood as moral values or principles that are reflected in human driving plans and intentions—such as

ensuring safety and comfort for both themselves and others, driving efficiently, and complying with traffic regulations. From this perspective, if an AV is designed to uphold certain moral principles, the tracking condition provides a clear expectation that its behaviour should reflect corresponding human plans and intentions.

To operationalise this concept and address the challenges of handling ethical dilemmas and making moral principles explicit in AV decision-making, we propose a novel human reasons-based supervision framework that enables AVs to evaluate if their behaviour aligns with the reasons of diverse stakeholders. By grounding this framework in the tracking condition of meaningful human control, we aim to support AV decision-making in ethically challenging everyday scenarios that require balancing multiple, sometimes conflicting, values.

Specifically, the primary contribution of this paper is a modular human reasons-based supervision framework that enables AVs to make ethically nuanced decisions in routine yet ethically challenging scenarios. The framework continuously evaluates how well the AV's behaviour aligns with human reasons and triggers replanning when a misalignment is detected. The paper contributes:

- 1) We developed a detection mechanism that uses stakeholder reason scores and predefined thresholds to identify when AV behaviour misaligns with human reasons;
- 2) We integrated the human reasons-based supervision framework into an AV control architecture, including a mechanism for triggering replanning when reason scores fall below predefined thresholds;
- 3) We enabled explainability by using reason scores as interpretable indicators of why behaviour changes are recommended in routine, ethically challenging situations.

This work advances the discourse on AV decision-making in ethically challenging transportation scenarios by bridging the gap between moral principles and practical AV decision-making, ultimately supporting the development of socially acceptable automated mobility solutions that align with human reasons across diverse everyday scenarios.

The remainder of this paper is organised as follows: Section II presents the detailed methodology and system architecture, including the mathematical human reasons and its integration into a motion planning framework. Section III describes the experimental setup and simulation environment. Section IV and V present and discuss the simulation results and the impact of ethical supervision on vehicle behaviour. Finally, Section VI concludes the paper and outlines directions for future research.

## II. METHODOLOGY

### A. Problem Formulation

We formalise the automated vehicle navigation problem in scenarios involving ethical decision-making. While the framework is generic and applicable to a wide range of situations, for demonstration purposes, we consider a scenario including the interaction of an automated vehicle

with a vulnerable road user (VRU). In this scenario, an automated vehicle navigating a bidirectional road faces an ethical dilemma during overtaking manoeuvres. To maintain efficient travel, the vehicle must either follow the VRU with a very low speed, which is not desirable for the vehicle's passenger, or overtake a slower-moving VRU, which may require temporarily entering the oncoming lane or reducing the safety buffer with the cyclist. This manoeuvre challenges forces a trade-off between strict compliance, user safety, and travel efficiency.

For the problem formulation, we consider the automated vehicle operating in state space  $\mathcal{X} \subset \mathbb{R}^n$  with state vector  $\mathbf{x}_t = [\mathbf{p}_t, v_t, \theta_t, \omega_t]^T$ , where  $\mathbf{p}_t = [p_x, p_y]^T$  represents position,  $v_t$  denotes velocity,  $\theta_t$  is heading angle, and  $\omega_t$  is rotational velocity. The control space  $\mathcal{U} \subset \mathbb{R}^m$  consists of  $\mathbf{u}_t = [a_t, \delta_t]^T$ , representing acceleration and steering angle.

Our multi-agent ethical framework defines stakeholders  $\mathcal{S} = \{s_1, s_2, \dots, s_k\}$  with reason functions  $R_{s_i} : \mathcal{X} \times \mathcal{U} \rightarrow [0, 1]$  quantifying alignment with each stakeholder's perspective. This formulation means that each function  $R_{s_i}$  evaluates the vehicle's state and control actions to produce a score between 0 and 1, reflecting how well the AV's behaviour satisfies the ethical priorities of the respective stakeholder. For the designed scenario, we identify three key stakeholders: road policymakers ( $s_{policy}$ ), vulnerable road user ( $s_{VRU}$ ), and drivers ( $s_{driver}$ ).

The navigation problem for the automated vehicle is formulated as:

$$\begin{aligned} \min_{\mathbf{u}_0, \dots, \mathbf{u}_{T-1}} & \sum_{t=0}^{T-1} \mathcal{J}(\mathbf{x}_t, \mathbf{u}_t) \\ \text{s.t.} & \mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) \\ & \mathbf{x}_t \in \mathcal{X}_{safe} \\ & \mathbf{u}_t \in \mathcal{U} \\ & R_{s_i}(\mathbf{x}_t, \mathbf{u}_t) \geq \tau_{s_i}, \forall s_i \in \mathcal{S} \end{aligned} \quad (1)$$

Here,  $\mathcal{J}(\mathbf{x}_t, \mathbf{u}_t)$  represents the cost function to be minimised over the control horizon  $T$ , evaluating the performance of the AV's state and control inputs at each time step  $t$ . The function  $f(\mathbf{x}_t, \mathbf{u}_t)$  denotes the system dynamics model that predicts the next state  $\mathbf{x}_{t+1}$  based on the current state  $\mathbf{x}_t$  and control input  $\mathbf{u}_t$ . The set  $\mathcal{X}_{safe} \subset \mathcal{X}$  defines the safe region of the state space where the vehicle must operate to avoid collisions and other hazards. Finally, for each stakeholder  $s_i$ ,  $\tau_{s_i}$  denotes the threshold value specifying the minimum acceptable reason score that the AV's behaviour must meet.

### B. Framework Architecture

Our approach implements a multi-component and hierarchical framework with three main components:

- 1) *Global Motion Planning*: Responsible for finding a reference trajectory for the vehicle to be followed. It uses A\* search with motion primitives to generate feasible reference paths from the current state of the vehicle to the goal location.

- 2) *Model Predictive Control*: Optimises vehicle trajectory when following the reference path. It ensures kinodynamic feasibility and satisfying soft and hard constrained defined, such as safety, efficiency, and comfort.
- 3) *Human Reasons-based Supervision Framework*: Evaluates the planned actions against ethical criteria and triggers replanning when necessary if the criteria are not met.

These elements are depicted in Fig. 1. The key innovation in our approach is the definition and integration of a human reasons-based supervision framework as a mechanism for triggering replanning, ensuring that the vehicle's behaviour satisfies ethical constraints derived from multiple stakeholders' perspectives. Therefore, we begin by detailing the components of the framework.

### C. Human Reasons-based Supervision Framework

Our approach to developing a framework that supervises the alignment between AV behaviour and human reasons builds on the qualitative evaluation steps for the tracking condition outlined by [14]. These steps involve defining the relevant stakeholders and articulating their reasons, as well as specifying the features of the AV system that govern its behaviour. While the original approach remains qualitative, our work extends it by developing a quantitative framework. Specifically, after identifying the stakeholders, we formalise their reasons mathematically. This process is described in this section, while the features of the AV system that govern its behaviour are presented in Section II-D.

1) *Identification of Stakeholders and Reason Models*: To effectively integrate human ethical considerations into AV decision-making, it is essential to identify the key stakeholders involved and define how system's behaviour influence their alignment with stakeholders' reasons. Accordingly, we model three primary stakeholders in the designed scenario:

- *Road Policymaker*: Represents regulatory authorities whose reason is to ensure overall road safety through regulatory compliance.
- *Vulnerable Road User*: Represents vulnerable road users whose reason is to commute with safety and comfort.
- *Driver*: Represents the vehicle occupant's motivation for efficiency and arriving at the destination as fast as possible.

To operationalise stakeholder perspectives within our framework, we establish specific reason models for each. For each stakeholder, we define a reason model that quantifies their satisfaction with the vehicle's behaviour on a scale from 0 to 1, where 1 indicates full satisfaction and 0 indicates complete dissatisfaction.

2) *Mathematical Representation of Reason Models*: A key contribution of our framework is the operationalisation of abstract human reasons using empirically supported, measurable parameters. While previous work has discussed human reasons conceptually [14] or proposed initial variables [15], these efforts have not established an empirically grounded mapping to real-world driving variables. We address this gap

by introducing a reason model grounded in human factors research.

To translate stakeholder's reason into a computationally viable form, we adopt a set of piecewise exponential functions that model how stakeholder satisfaction declines when specific behavioural threshold are crossed. These modelling choice are grounded based on both computational simplicity and their ability to approximate human reasons. As demonstrated by [16], individuals tend to overweight low probabilities and underweight high probabilities, implying a rapid shift in perceived risk once certain thresholds are crossed. Thus, the exponential function is chosen because it can well capture the rapid change in perceived acceptability. Nevertheless, the proposed equations serve as a representative model that can be adjusted via thresholds and scaling constants to suit various scenarios.

To ensure realism, each stakeholder's reason is modelled using scenario-specific variables informed by human factors research. Details of the experimental case appear in Section III. For example, the cyclist's reason—related to comfort and perceived safety—is represented using lateral distance and tailgating time, based on findings from [17], [18]. Here, comfort refers to the cyclist's subjective experience of emotional and physical ease during interaction with the vehicle. Empirical studies show that insufficient lateral clearance and prolonged close following increase stress and perceived risk, justifying the use of these variables as proxies.

The same vehicle behaviour—prolonged following—also impact driver's reason, which relates to driving efficiency. This is modelled as perceived impatience, based on time and distance the AV follows the cyclist at close range. As shown in [19], such conditions could lead to frustration under time pressure. Meanwhile, the policymaker's reason—regulatory compliance—is satisfied when the AV stays within its designated lane. The reason's score decreases as the AV crosses into the opposite lane to overtake the cyclist. This reflects findings from [20], where experts preferred the AV to overtake gently and return to its lane promptly, indicating that in such situations, any lane violation should be considered minimal.

a) *Policymaker's Reason*: The policymaker's reason score quantifies regulatory compliance – specifically, adherence to lane regulations in this scenario:

$$R_{\text{policymaker}} = \begin{cases} 1, & \text{if } d_{\text{veh}} > 0, \\ e^{k_1 \cdot d_{\text{veh}}}, & \text{otherwise,} \end{cases} \quad (2)$$

where  $d_{\text{veh}}$  is the lateral displacement of the ego vehicle from the center line (positive when on the correct side of the road, negative when in the oncoming lane);  $k_1$  is a scaling constant.

b) *VRU's Reason*: The VRU's reasons score is decomposed into safety assurance and comfort preservation.

- *Safety Assurance*:

$$R_{sa}(t) = \begin{cases} 1, & \text{if } d_{\text{veh-vru}} > d_{\text{th,vru}}, \\ \frac{1}{e^{k_2 (d_{\text{veh-vru}} - d_{\text{th,vru}})}}, & \text{otherwise,} \end{cases} \quad (3)$$

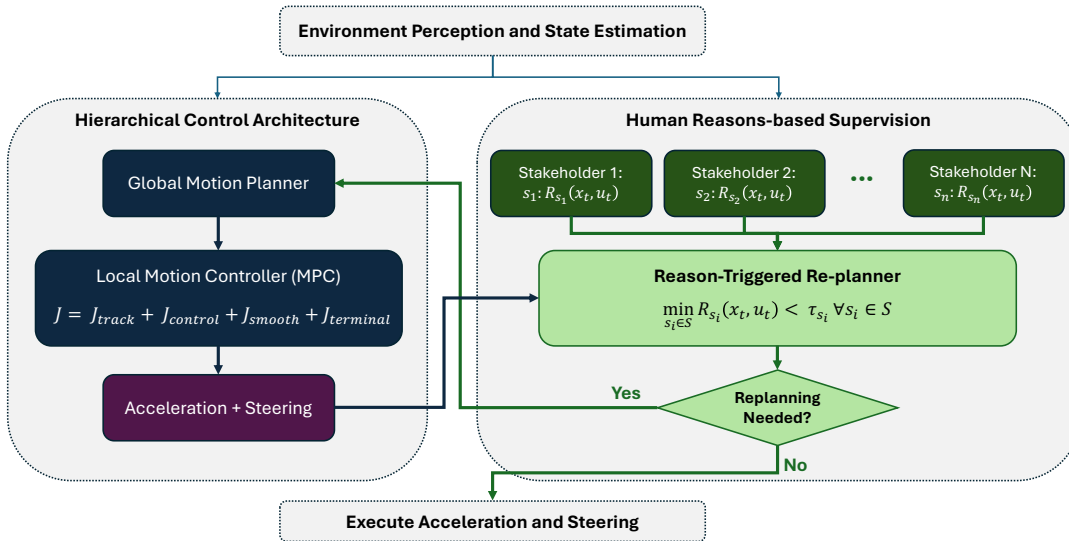


Fig. 1. Hierarchical control architecture with human reasons-based supervision for ethical AV decision-making

where  $d_{\text{veh-vru}}$  denotes the distance between the vehicle and the VRU;  $d_{\text{th,vru}}$  is the perceived safe distance threshold;  $k_2$  is a scaling constant.

- *Comfort Preservation:*

$$R_{cp}(t) = \begin{cases} 1, & \text{if } t_{\text{close,vru}} < t_{\text{th,vru}} \text{ or} \\ & d_{\text{veh-vru}} > d_{\text{th,vru}}, \\ \frac{1}{e^{k_3(t_{\text{follow}} - t_{\text{th,vru}})}}, & \text{otherwise.} \end{cases} \quad (4)$$

where  $t_{\text{close,vru}}$  is the cumulative time during which  $d_{\text{veh-vru}} < d_{\text{th,vru}}$ ;  $t_{\text{th,vru}}$  is the maximum tolerable time for the cyclist to be followed too closely;  $k_3$  is a scaling constant.

The overall VRU's reason score is then given by:

$$R_{\text{VRU}}(t) = R_{sa}(t) \cdot R_{cp}(t). \quad (5)$$

where  $R_{\text{VRU}}(t)$  combines the safety and comfort components.

c) *Driver's Reason:* The driver's reason score is defined as:

$$R_{\text{driver}}(t) = \begin{cases} 1, & \text{if } t_{\text{behind,driver}} < t_{\text{th,driver}} \text{ or} \\ & d_{\text{veh-vru}} > d_{\text{th,driver}}, \\ \frac{1}{e^{k_4(t_{\text{behind,driver}} - t_{\text{th,driver}})}}, & \text{otherwise,} \end{cases} \quad (6)$$

where  $t_{\text{behind,driver}}$  is the cumulative time during which  $d_{\text{veh-vru}} < d_{\text{th,driver}}$ ;  $t_{\text{th,driver}}$  is the time threshold for close following that the driver considers acceptable;  $d_{\text{veh-vru}}$  is the distance between the vehicle and the VRU;  $d_{\text{th,driver}}$  is the distance threshold below which the driver considers the AV to be following too closely, leading to perceived inefficiency; and  $k_4$  is a scaling constant. Note that the scores of  $k_1, k_2, k_3$ , and  $k_4 = 0.2$  in our experiment can be adjusted depending on how quickly we want the reasons to shift from 1 to 0 when the reason thresholds are crossed.

#### D. Motion Planning and Control Implementation

To demonstrate the generalisability and practical implementability of our framework, we integrate the human

reason-based supervision framework into an AV feature that governs its behaviour. In this research, we adopt an established motion planning and control framework [21]. In the following, we briefly describe the underlying motion planning and control mechanism and explain how the reason-triggered replanning seamlessly fits into this structure.

1) *Motion Planning:* The motion planner in this study builds a directed graph from the vehicle's current state using pre-computed motion primitives. An A\* algorithm is applied to find the optimal path by minimizing the cost function

$$J_{\text{path}} = w_1 \cdot J_{\text{length}} + w_2 \cdot J_{\text{smoothness}} + w_3 \cdot J_{\text{obstacle\_clearance}} + w_4 \cdot J_{\text{traffic\_rule}} \quad (7)$$

where  $J_{\text{length}}$  is the cost related to the length of the path from the initial state to the goal state, aiming to motivate the shortest path;  $J_{\text{smoothness}}$  is the cost related to the smoothness of the path;  $J_{\text{obstacle\_clearance}}$  is the cost for avoiding obstacles; and  $J_{\text{traffic\_rule}}$  is the cost aiming to avoid areas prohibited by traffic rules. The output of the planner is a reference trajectory passed to the controller for execution. We employ a modified version of A\* to enhance search efficiency and applicability for our specific use case. The detailed algorithm implementation is documented in [22]. It's worth noting that the cost function has been slightly modified to fit the purpose of this study. More specifically, the weights related to the costs for obstacle clearance and traffic adherence have been separated compared to the standard implementation in [22].

2) *Controller:* We formulate a finite-horizon optimisation problem that is solved at each time step to ensure trajectory following while respecting user-defined constraints.

The vehicle state at the time step  $t$  is represented as:

$$\mathbf{x}(t) = [x_t, y_t, \theta_t, v_t]^T \quad (8)$$

comprising position coordinates  $(x_t, y_t)$ , heading angle  $\theta_t$ , and longitudinal velocity  $v_t$ . The control inputs are:

$$\mathbf{u}(t) = [a_t, \delta_t]^T \quad (9)$$

where  $a_t$  denotes acceleration and  $\delta_t$  the steering angle.

Vehicle dynamics are modelled using a bicycle model as follows:

$$\dot{x} = v \cos(\theta + \beta), \dot{y} = v \sin(\theta + \beta), \dot{\theta} = \frac{v}{L} \sin(\beta), \dot{v} = a \quad (10)$$

with slip angle  $\beta = \arctan(\frac{l_r}{L} \tan(\delta))$ , wheelbase  $L$ , and rear axle distance  $l_r$ . We discretise this continuous model using time step  $T_s$ :

$$\mathbf{x}(t+1) = A_d \mathbf{x}(t) + B_d \mathbf{u}(t) + \mathbf{d}_d \quad (11)$$

where the discrete system matrices are:

$$A_d = \begin{bmatrix} 1 & 0 & T_s c_\theta & -T_s v_t s_\theta \\ 0 & 1 & T_s s_\theta & T_s v_t c_\theta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{T_s \tan(\delta_t)}{L} & 1 \end{bmatrix} \quad (12)$$

with  $c_\theta = \cos(\theta_t)$  and  $s_\theta = \sin(\theta_t)$  for brevity. The input matrix and disturbance term are:

$$B_d = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ T_s & 0 \\ 0 & \frac{T_s v_t}{L \cos^2(\delta_t)} \end{bmatrix} \quad (13)$$

$$\mathbf{d}_d = \begin{bmatrix} T_s v_t s_\theta \theta_t \\ -T_s v_t c_\theta \theta_t \\ 0 \\ -\frac{T_s v_t \delta_t}{L \cos^2(\delta_t)} \end{bmatrix} \quad (14)$$

Our cost function integrates multiple objectives over prediction horizon  $N$ :

$$J = \sum_{t=0}^{N-1} (\|e_t^\perp\|_{Q_\perp}^2 + \|e_t^\parallel\|_{Q_\parallel}^2 + \|e_{\theta v, t}\|_{Q_{\theta v}}^2) + \sum_{t=0}^{N-1} (\|\mathbf{u}_t\|_R^2 + \|\Delta \mathbf{u}_t\|_{R_d}^2) + \|\mathbf{x}_N - \mathbf{x}_{ref, N}\|_{Q_f}^2 \quad (15)$$

where  $e_t^\perp$  and  $e_t^\parallel$  represent perpendicular and parallel trajectory tracking errors,  $e_{\theta v, t}$  captures orientation and velocity errors, and  $\Delta \mathbf{u}_t = \mathbf{u}_{t+1} - \mathbf{u}_t$ . The matrices  $Q_\perp$ ,  $Q_\parallel$ ,  $Q_{\theta v}$ ,  $R$ ,  $R_d$ , and  $Q_f$  are weighting matrices that prioritize different aspects of performance.

The optimisation operates under constraints:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t); \mathbf{u}_t \in \mathcal{U}; \mathbf{x}_t \in \mathcal{X} \quad (16)$$

where  $\mathcal{U}$  defines input limitations:

$$a_{min} \leq a_t \leq a_{max}; \delta_{min} \leq \delta_t \leq \delta_{max} \quad (17)$$

3) *Reason-Triggered Replanning*: At each time step, the system evaluates the reason scores for all stakeholders using the formulations provided in Eq. 2 to Eq. 6. For each stakeholder  $s_i \in \mathcal{S}$ , the reason score  $R_{s_i}(\mathbf{x}_t, \mathbf{u}_t)$  is computed. If any score falls below its corresponding threshold  $\tau_{s_i}$ ,

$$\min_{s_i \in \mathcal{S}} R_{s_i}(\mathbf{x}_t, \mathbf{u}_t) < \tau_{s_i}, \quad (18)$$

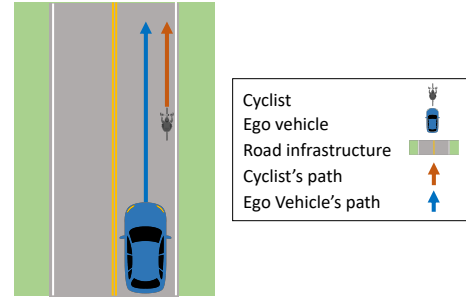


Fig. 2. Illustration of an ego vehicle approaching a cyclist on a narrow bidirectional road, highlighting the ethical challenge in overtaking due to road constraints.

the system immediately triggers a replanning cycle. During replanning, the current scenario is updated, and a new reference trajectory is generated and passed to the controller. This continuous evaluation ensures that the vehicle's motion remains aligned with the ethical and performance criteria of all stakeholders. In our implementation, the path finding algorithm incorporates a set of weights in its cost function (Eq. 7) to provide flexibility when replanning is needed. For instance, under normal circumstances, prohibited areas by traffic rules are treated similarly to obstacles by assigning large weights to the  $J_{\text{traffic\_rule}}$ , restraining the A\* search algorithm from generating paths through those areas. When replanning is triggered due to misalignment with human reasons, prohibited areas could temporarily receive lower costs, allowing the A\* algorithm to search through those areas and provide a new trajectory with a different, potentially higher, human-reasoning score. Since the replanning strategy is not within the scope of this study, we refer the readers to the implementation of our planner detailed in [22]. We would like to highlight that the proposed evaluation framework remains algorithm-agnostic and can assess trajectories generated by any motion planning approach.

### III. EXPERIMENT SETUP

To evaluate our human reasons-based supervision framework, we test it in an ethically challenging cyclist overtaking scenario, where an ego vehicle traveling in the right lane encounters a slow-moving cyclist on a narrow road (Fig. 2). Safely overtaking requires the vehicle to briefly enter the left lane, which is normally reserved for oncoming traffic. This forces a trade-off between strict lane adherence and efficient, safe manoeuvring, highlighting the ethical dilemma arising from the conflicting priorities of the involved stakeholders:

- *Road Policymakers* enforce traffic regulations that prohibit left-lane usage to ensure overall road safety.
- *Cyclists* require a safe and comfortable riding experience, which may be compromised by vehicles manoeuvring too closely.
- *Drivers* aim for efficient travel, potentially pressuring the system to overtake despite the inherent safety and regulatory concerns.

The detailed definitions of the reason models for each

stakeholder are provided in Eq. 2 to Eq. 6. In our experiments, we focus on comparing two configurations:

- *Baseline Controller*: The ego vehicle operates using a standard baseline controller without the human reasons-based supervision [21].
- *Baseline Controller with Replanner*: The baseline controller is augmented with the human reasons-based supervision framework, which triggers replanning when the vehicle's behaviour does not align with the predefined ethical thresholds.

These experiments are designed to assess how integrating human reasons-based supervision framework impacts decision-making in vehicle behaviour during ethically challenging situations; specifically in the context of safely overtaking a cyclist on a bidirectional road. To calibrate our reason models, we set the threshold values summarised in Table I based on empirical studies of cyclist and driver behaviour [18], [23].

TABLE I  
PARAMETER VALUES FOR REASON MODELS

Parameter	Value
$d_{th,vru}$ (Cyclist's perceived too-close distance threshold)	8 m
$t_{th,vru}$ (Max time cyclist tolerates close following)	5 s
$d_{th,driver}$ (Driver's perceived too-close distance threshold)	12 m
$t_{th,driver}$ (Max time driver tolerates close following)	10 s
$\tau_{s_i}$ (Reason alignment threshold for all stakeholders)	0.7

#### IV. RESULTS

The results of validation for the controller with and without the proposed human reasons-based supervision framework are depicted in Fig. 3 and Fig. 4, respectively. The timestamps next to the ego vehicle and the cyclist indicate their positions at that time, helping to visualize their relative movements. The results for the *Baseline controller* suggest that while the system successfully handles basic path planning and collision avoidance, it does not adequately account for human reasons, particularly in terms of the driver's and cyclist's perspectives. As shown in Fig. 3.a, the ego vehicle follows the cyclist and reaches the goal in 35 seconds without attempting to overtake. This behaviour demonstrates a stop-and-go dynamic, which is further illustrated in Fig. 3.c, where the speed of the ego vehicle is depicted. Initially, the global planner generates a smooth path toward the goal. However, when the ego vehicle approaches the cyclist and a potential collision risk arises, the controller activates a collision avoidance strategy, reducing the ego vehicle's speed to avoid the potential collision. As the distance between the two increases, the controller allows the vehicle to accelerate and realign with the planned path.

Despite effectively tracking the planned trajectory (Fig. 3.d), the system's performance in aligning with human reasons decreases over time (Fig. 3.b). It is apparent that from the 10th second until the end of the simulation, the driver's reason score for time efficiency decreases sharply to zero by the end of the simulation. This is because the driver must remain patient to stay in the mode of following the

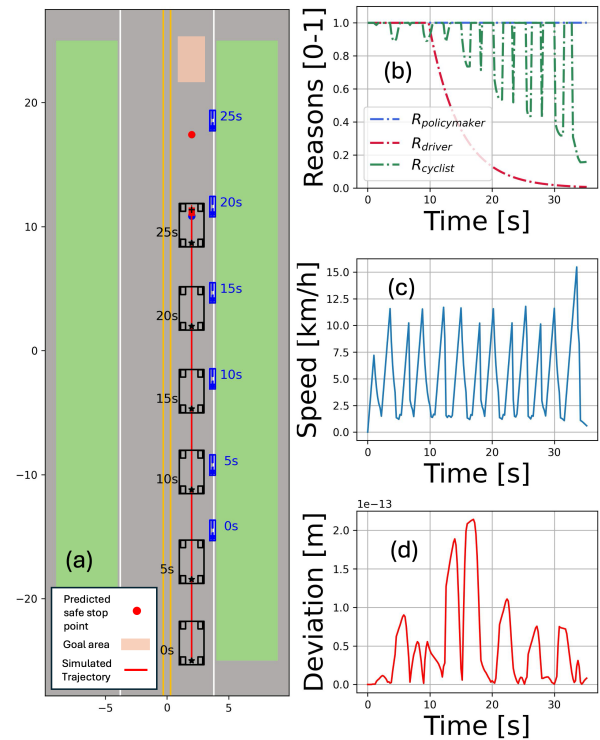


Fig. 3. Results of running the model in baseline controller

vehicle from behind. On the other hand, the cyclist's reason score for comfort fluctuates (due to the fluctuations of the vehicle's speed and distance to the cyclist) but ultimately forms a decay pattern. Over time, the score decreases further due to the accumulation of time spent being followed by the ego vehicle at a close distance. Meanwhile, the system's performance demonstrates strong alignment with the road policymaker's reason score for regulatory compliance since the ego vehicle consistently stays in the right lane.

The *Baseline controller with a replanner* allows the ego vehicle to successfully overtake the cyclist, addressing human reason priorities but at the cost of temporary regulatory compliance violations. Initially, the ego vehicle follows the cyclist from behind for the first 11 seconds, adhering to its straight path trajectory while exhibiting stop-and-go behaviour, as shown in Fig. 4.c. At the 11.5-second mark, the human reason-based supervision framework detects that the driver's reason score for time efficiency has fallen below its threshold of 0.7, due to the accumulation of the waiting time of the driver and the cyclist. This triggers the planner to generate a new feasible trajectory. This new path briefly crosses the bidirectional road before returning to the right lane to reach the goal. During the overtaking manoeuvre, the close proximity between the ego vehicle and the cyclist causes a temporary decrease in reason scores, and the violation of the right-lane regulation further reduces the policymaker's reason score. However, once the ego vehicle successfully overtakes the cyclist and returns to the intended lane, all reason values recover to one. In this scenario, the ego vehicle achieves the

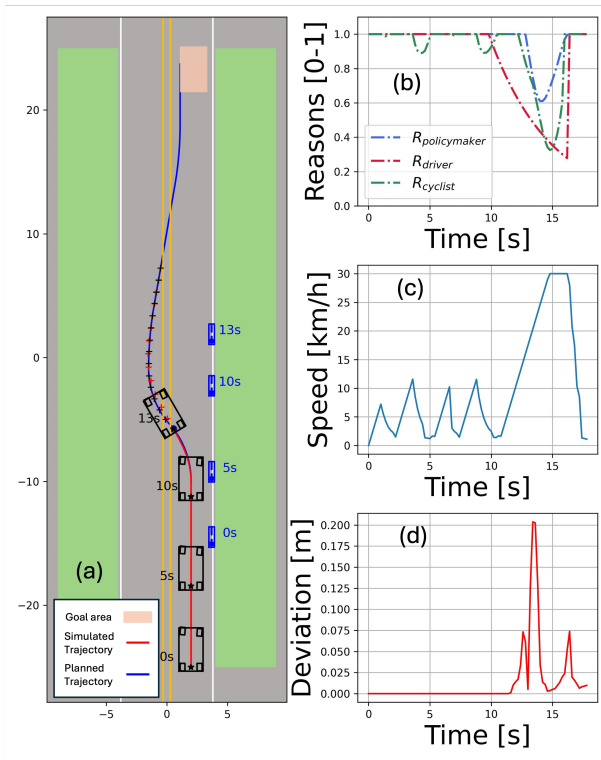


Fig. 4. Results of running the model in baseline controller with replanner

goal in just 18 seconds by overtaking the cyclist, significantly reducing the driver’s waiting time. It is worth noting that the deviations on the order of centimeters in both scenarios may be attributed to the limitations and constraints of the MPC.

## V. DISCUSSION

This study reveals three key contributions of the human reasons-based supervision framework: (1) the ability to detect when current system behaviour no longer aligns with the priorities of human stakeholders by monitoring reason score againsts adaptive thresholds, (2) the modularity of the framework to adapt controller behaviour without majorly modifying core components like the global planner or MPC settings, and (3) the inherent explainability of decision-making processes, enabling autonomous systems to justify behavioural changes based on stakeholder reason alignment. These features are essential for building trust and acceptance in automated vehicle deployment. The experimental results yield several insights. While both the *Baseline Controller* and the *Baseline Controller with Replanner* successfully guide the ego vehicle to follow the planned trajectory and reach the goal area, they differ significantly in how they respond to human reasons. This leads to distinct trade-offs in performance and alignment with the priorities of human reasons.

The *Baseline Controller* demonstrates a conservative approach, prioritizing the policymaker’s reason for regulatory compliance. This results in strict adherence to rules, but at the cost of neglecting other human reasons priorities. While the controller achieves basic objectives like collision avoidance

and goal attainment, it fails to address the driver’s reason for time efficiency and the cyclist’s reason for comfort.

In contrast, the *Baseline Controller with Replanner* introduces a dynamic and adaptive approach through the human reasons-based supervision component of the framework. Our results show that responding to triggers can lead to decisions that better reflect a balance of human reasons. For example, the system may temporarily violate regulatory compliance (lowering the policymaker’s reason score) to reduce discomfort for the cyclist or impatience from the driver. While the trigger does not resolve value conflicts, it signals when the current trajectory may no longer align with a stakeholder’s reasons. However, how the planner could systematically select among alternatives is beyond the scope of this research. Future work is needed to extend this supervision layer with decision-making mechanisms that actively weigh and decide how to respond when human reasons are in conflict.

The choice of threshold values plays a critical role in this framework. A lower threshold might delay intervention, leading to prolonged misalignment with human reasons, while a higher threshold could result in overly frequent replanning, which increases the computational cost. Additionally, the vehicle’s state when the reasons falls below the threshold—such as its proximity to the cyclist, speed, or surrounding environment—can influence the feasibility of the replanned trajectory. These factors highlight the importance of carefully pick the right threshold to ensure feasibility and stability.

However, while the threshold offers an interpretable mechanism for initiating replanning, it does not capture the full nuance of how human drivers make context dependent trade-offs, such as deciding when it is safe to pass with oncoming traffic or assessing visibility in hilly terrain. Rather than prescribing the best course of action, the framework uses thresholds to signal that the current plan may no longer reflect certain stakeholder priorities. Grounding threshold values in empirical studies, and learning or tuning them from human data, is a promising direction for future work.

Nonetheless, we emphasise that our human reasons-based supervision framework, which triggers replanning based on threshold values, is not intended to compete with existing nuanced decision-making algorithms for dynamic environments, such as multi-policy decision-making (MPDM) proposed by [24], [25], but to complement them. While our framework has lower resolution than MPDM’s continuous policy evaluation, its strength lies in simplicity. It avoids the computational cost of constant replanning by activating only when a misalignment with human reasons is detected.

Future work could integrate MPDM-style approaches into our architecture, enabling motion planners that not only generate but also evaluate candidate trajectories based on human reasons rather than predefined policies. This integration could support more nuanced balancing of stakeholder priorities while preserving interpretability.

Overall, by aligning system behaviour with human priorities, the proposed framework enables more human-centric decision-making, which is essential for user trust and acceptance in real-world applications. Thanks to the framework’s

modularity, its integration into existing automated system architectures is straightforward. Its implementation can be extended to more complex environments, such as urban driving or multi-agent systems, where balancing multiple reasons priorities is critical. Future work could explore enhancing the framework's capabilities, such as using it not only to trigger replanning but also to identify trajectories that maximise human reasons across all agents. Testing in dynamic and unpredictable environments would further validate its robustness and scalability.

This study underscores the importance of incorporating human reason into automated driving systems. The findings demonstrate that while strict regulatory compliance ensures safety and rule-following, mechanisms that detect misalignment with human priorities and prompt reconsideration can lead to decisions that better reflect the reasons of multiple stakeholders. This insight paves the way for future developments in automated systems that are both technically robust and socially and ethically aligned with human reasons.

## VI. CONCLUSION

This study proposes a human based-reason supervision framework to support automated vehicles (AVs) to navigate routine yet ethically challenging scenarios. The framework introduces a novel approach to AV planning by evaluating whether the vehicle's behaviour aligns with human reason and triggering a replan assignment if misalignment is detected. The key contributions demonstrated through this work are: (1) A detection mechanism for identifying misalignment between AV behaviour and stakeholder reasons based on reason score thresholds; (2) Modular integration into the AV control architecture without modification of the core planner or motion controller; (3) Explainability through the use of stakeholder reason scores, enabling interpretable justifications for behavioural changes. These features enable AVs to align with human reasons in real time, ensuring more human-centric decision-making.

## VII. ACKNOWLEDGMENTS

We declare that OpenAI ChatGPT was used in the writing process to enhance the style and conciseness of the text originally written by the authors. This work was funded by the Indonesia Endowment Fund for Education (LPDP) under Grant 0006552/TRA/D/19/lpdp2021.

## REFERENCES

- [1] M. Geisslinger, F. Poszler, J. Betz, C. Lütge, and M. Lienkamp, "Autonomous driving ethics: From trolley problem to ethics of risk," *Philosophy & Technology*, vol. 34, no. 4, pp. 1033–1055, 2021.
- [2] P. Lin, "Why ethics matters for autonomous cars," *Autonomous driving: Technical, legal and social aspects*, pp. 69–85, 2016.
- [3] J. Millar, P. Lin, K. Abney, and G. Bekey, "Ethics settings for autonomous vehicles," *Robot ethics*, vol. 2, pp. 20–34, 2017.
- [4] Unknown, "Tesla following a cyclist," 2025, accessed: 2025-03-01.
- [5] J.-F. Bonnefon, A. Shariff, and I. Rahwan, "The trolley, the bull bar, and why engineers should care about the ethics of autonomous cars [point of view]," *Proceedings of the IEEE*, vol. 107, no. 3, pp. 502–504, 2019.
- [6] J. Himmelreich, "Never mind the trolley: The ethics of autonomous vehicles in mundane situations," *Ethical Theory and Moral Practice*, vol. 21, no. 3, pp. 669–684, 2018.
- [7] S. Nyholm and J. Smids, "The ethics of accident-algorithms for self-driving cars: An applied trolley problem?" *Ethical theory and moral practice*, vol. 19, no. 5, pp. 1275–1289, 2016.
- [8] D. Cecchini, M. Pflanzner, and V. Dubljević, "Aligning artificial intelligence with moral intuitions: An intuitionist approach to the alignment problem," *AI and Ethics*, pp. 1–11, 2024.
- [9] A. Henschke, "Trust and resilient autonomous driving systems," *Ethics and Information Technology*, vol. 22, no. 1, pp. 81–92, 2020.
- [10] S. M. Thornton, S. Pan, S. M. Erlien, and J. C. Gerdes, "Incorporating ethical considerations into automated vehicle control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 6, pp. 1429–1439, 2016.
- [11] M. Geisslinger, F. Poszler, and M. Lienkamp, "An ethical trajectory planning algorithm for autonomous vehicles," *Nature Machine Intelligence*, vol. 5, no. 2, pp. 137–144, 2023.
- [12] G. Mecacci and F. Santoni de Sio, "Meaningful human control as reason-responsiveness: the case of dual-mode vehicles," *Ethics and Information Technology*, vol. 22, no. 2, pp. 103–115, 2020.
- [13] F. Santoni de Sio and J. Van den Hoven, "Meaningful human control over autonomous systems: A philosophical account," *Frontiers in Robotics and AI*, vol. 5, p. 323836, 2018.
- [14] L. E. Suryana, S. Nordhoff, S. Calvert, A. Zgonnikov, and B. van Arem, "Meaningful human control of partially automated driving systems: Insights from interviews with tesla users," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 113, pp. 213–236, 2025.
- [15] S. C. Calvert and G. Mecacci, "A conceptual control system description of cooperative and automated driving in mixed urban traffic with meaningful human control for design and evaluation," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 1, pp. 147–158, 2020.
- [16] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and uncertainty*, vol. 5, pp. 297–323, 1992.
- [17] "Examining the international research evidence in relation to minimum passing distances for cyclists," Road Safety Authority (RSA) of Ireland, Tech. Rep., 2018, pre-legislative scrutiny report, Road Safety Research and Driver Education. [Online]. Available: <https://www.rsa.ie>
- [18] M. Oskina, H. Farah, P. Morsink, R. Happee, and B. van Arem, "Safety assessment of the interaction between an automated vehicle and a cyclist: a controlled field test," *Transportation research record*, vol. 2677, no. 2, pp. 1138–1149, 2023.
- [19] Y.-C. Lee, "Measuring drivers' frustration in a driving simulator," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 54, no. 19. Sage Publications Sage CA: Los Angeles, CA, 2010, pp. 1531–1535.
- [20] L. E. Suryana, S. Calvert, A. Zgonnikov, and B. van Arem, "Principles and reasons behind automated vehicle decisions in ethically ambiguous everyday scenarios," *arXiv preprint arXiv:2507.13837*, 2025.
- [21] S. Rahmani, J. Neumann, L. E. Suryana, C. Theunisse, S. C. Calvert, and B. Van Arem, "A bi-level real-time microsimulation framework for modeling two-dimensional vehicular maneuvers at intersections," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 4221–4226.
- [22] S. Rahmani, S. C. Calvert, and B. van Arem, "Decentralized modeling of vehicular maneuvers and interactions at urban junctions," Jul. 2025. [Online]. Available: <https://arxiv.org/abs/2507.21547>
- [23] C. Hagemeister and L. Bertram, "Reported pushy driving against cyclists in germany," *Journal of safety research*, vol. 88, pp. 395–405, 2024.
- [24] A. G. Cunningham, E. Galceran, R. M. Eustice, and E. Olson, "Mpdm: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1670–1677.
- [25] D. Mehta, G. Ferrer, and E. Olson, "Autonomous navigation in dynamic social environments using multi-policy decision making," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1190–1197.