

From Human Negotiation to Agent Negotiation: Personal Mobility Agents in Automated Traffic

PASCAL JANSEN, Institute of Media Informatics, Ulm University, Germany

Conflicts between user preferences and automated system behavior already shape the experience of automated mobility. For example, a passenger may prefer assertive driving, yet the vehicle slows down early to follow a conservative policy or yield to other actors. Similar conflicts arise at merges, crossings, or right-of-way situations, where users must accept opaque decisions or attempt to negotiate through interfaces not designed for continuous, multi-actor relationships. This position paper argues that such approaches do not scale as mobility becomes more heterogeneous and automated. Instead, it proposes a conceptual perspective for *personal mobility agents* that act as proxies for users, encode preferences such as comfort and safety margins, and negotiate traffic behavior with other agents under shared safety rules. The central idea is a shift from moment-to-moment user negotiation interfaces to delegation and oversight interfaces, in which proxy agents manage real-time conflicts while users can shape high-level policies and preferences.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; *Interaction paradigms*; Ubiquitous and mobile computing; • **Computing methodologies** → Intelligent agents.

Additional Key Words and Phrases: mobility, human–automation interaction, agentic interactions

ACM Reference Format:

Pascal Jansen. 2026. From Human Negotiation to Agent Negotiation: Personal Mobility Agents in Automated Traffic. In *Proceedings of AutomationXP26 Workshop of the 2026 CHI Conference on Human Factors in Computing Systems, April 14, 2026, Barcelona, Spain*. ACM, New York, NY, USA, 6 pages.

1 Introduction

Mobility is a foundational socio-technical infrastructure because it enables access to work, education, healthcare, and social participation, while failures of mobility systems directly translate into safety risks [12]. At its core, everyday traffic is a continuous negotiation among multiple actors (e.g., drivers and cyclists) over shared space, including who proceeds first, what constitutes an acceptable gap, and how a driving maneuver is executed.

Automated mobility changes this negotiation by transferring parts of decision-making from humans to automated systems. Consequently, prior work has focused on how automated systems communicate with surrounding road users, for example, via external human–machine interfaces (eHMIs) that signal intent to pedestrians [3]. Across these strands, *trust calibration* (see Lee and See [8]) remains a central design objective as users must neither over-rely on automation nor avoid it, and adoption at scale depends on perceived safety, reliability, and preference alignment [3, 18]. Beyond signaling intent, recent work has also explored negotiation interfaces in automated driving, for example, in passenger–vehicle conversational negotiation scenarios [15]. These approaches move beyond static communication by allowing users to influence or discuss momentary decisions.

However, many interactions assume a one-human–one-system negotiation relationship, for example, in pedestrian–vehicle interaction at crossings [3] or in passenger–vehicle personalization such as driving-style adaptation [17]. In

Author’s Contact Information: [Pascal Jansen](mailto:pascal.jansen@uni-ulm.de), pascal.jansen@uni-ulm.de, Institute of Media Informatics, Ulm University, Ulm, Germany.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2026 Copyright held by the owner/author(s).

Manuscript submitted to ACM

contrast, prior work in multi-agent negotiation explores the design of protocols and negotiation strategies among automated actors in traffic settings, providing a formal basis for shared decision-making and maneuver coordination [1]. However, these approaches do not model persistent, user-representing preference proxies or delegation interfaces across situations. As the number of interacting actors increases, such direct negotiation approaches do not scale, as users cannot meaningfully track, evaluate, or participate in multiple concurrent negotiations.

In practice, automated traffic is a multi-actor environment involving drivers, passengers, pedestrians, automated vehicles, micromobility users (e.g., e-scooter riders), and intelligent infrastructure. Actors differ in abilities, preferences, and goals, such as age-related differences in gap acceptance [10], cultural differences in risk tolerance [6], and preferences for automated driving style [17]. Hence, traffic is a multi-objective negotiation under time pressure and partial observability, where each actor optimizes different objectives with incomplete information about others [3]. These conditions make conflicts between user intent and system constraints likely and can reduce automation engagement, trust, and acceptance [16]. To exemplify such conditions, consider an urban traffic scenario in the near future:

A passenger sits in an automated ride-sharing vehicle approaching a crossing. A cyclist approaches from the side, while a delivery robot moves along the sidewalk, and a pedestrian prepares to cross. The vehicle slows down earlier than the passenger expects, apparently yielding to the pedestrian. The cyclist, however, accelerates and crosses first, while the delivery robot pauses due to its own internal policy. From the passenger’s perspective, the situation appears inconsistent and difficult to interpret: it is unclear why the vehicle behaved conservatively, whether this reflects their own preferences, the cyclist’s behavior, or hidden infrastructure data. The user has no practical way to express a preference, such as “be more assertive in light traffic but always keep large safety gaps near pedestrians,” nor to understand how this preference would be negotiated with other actors.

In such scenarios, existing interaction approaches (e.g., [15]) would require either additional human-facing explanations, explicit negotiation dialogues, or post-hoc justifications. However, these typically assume that users have appropriate *situation awareness* (see Endsley [4]) so that they can meaningfully track and evaluate each interaction. In increasingly heterogeneous environments, such assumptions may not hold at scale. This could lead to disengagement, more overrides, generalized distrust, and ultimately reduced perceived control, trust calibration, and acceptance of automated mobility.

This position paper argues that scalable mobility interaction requires shifting from direct user-to-user negotiation to proxy-mediated negotiation. It proposes an interaction model in which each road user is represented by a **personal mobility agent** that acts as a proxy between the user and the traffic situation. The agent encodes the user’s preferences, constraints, and delegation policies, and negotiates with other actors on the user’s behalf. For example, a pedestrian’s agent may request larger safety gaps, while a vehicle’s agent may negotiate a smoother but slightly slower merge to preserve passenger comfort. Instead of requiring the user to participate in each momentary negotiation, the agents resolve these interactions among themselves and communicate only high-level outcomes to their users, if desired, such as “yielding early to maintain your preferred safety margin.”

In this view, interaction shifts from direct control of individual maneuvers toward the specification of high-level preferences and delegation policies. The design problem then becomes ensuring that proxy-based negotiation mechanisms remain safe, understandable, and acceptable to all involved actors.

While grounded in automated mobility, the proposed proxy-based negotiation model generalizes to other agentic automation contexts in which users delegate decisions to persistent agents that interact under shared constraints, such as domestic robotics, smart infrastructure, or collaborative AI systems.

2 Personal Mobility Agents in Automated Traffic Negotiations

This section outlines the possible technical foundations of personal mobility agents, describes potential mechanisms for their personalization, and exemplifies how they could operate in future automated traffic scenarios. Personal mobility agents are proposed here as a conceptual interaction model that combines existing components, including preference learning, human-in-the-loop optimization, and agentic AI.

2.1 Technological Foundations

From a technological perspective, personal mobility agents could build on several current developments. Automated vehicles and driver-assistance systems can already learn individual driving preferences, such as comfort levels or acceleration styles, from user input and observed behavior [9]. Human-in-the-loop optimization approaches further show that systems can learn such preferences efficiently through short interactive sessions. For example, Bayesian optimization has been used to adapt eHMI and in-vehicle interfaces to individual users [2, 7]. In parallel, recent progress in agentic AI, large language models, and vision–language models suggests that software agents can act on behalf of users, maintain persistent preference models, and interact with other agents or services [13]. Many of these systems can be personalized using reinforcement learning from human feedback, where the system learns preferred behavior from human ratings rather than fixed rules [11]. A personal mobility agent can be understood as a specialization of such personal agents for traffic contexts.

2.2 Core Functional Components

A personal mobility agent may perform three coupled functions:

Preference modeling The agent maintains a *preference model*, akin to a digital twin of the user, that captures aspects such as acceptable safety gaps, preferred acceleration levels, comfort limits, willingness to yield, and tolerance for delays or detours. This model, however, might remain constrained by safety and legal requirements (e.g., [5]).

Traffic negotiation During traffic events such as approaching a crossing, merging into a lane, following a slower vehicle, or entering a roundabout, the agent evaluates feasible maneuvers under safety constraints and negotiates a coordination outcome with other automated actors. For example, when two vehicles approach a merge, each agent may propose a maneuver that reflects its user’s preferences, such as “yield early for comfort” or “merge efficiently if the gap is safe.” The resulting maneuver is negotiated under shared safety rules, while the user remains at the high-level preference and policy control level.

Interface management The agent manages the *human-facing interface* by communicating outcomes. For example, after a merge, it may report: “merged conservatively due to comfort policy,” or after approaching a crossing: “yielded early to maintain your preferred safety margin.” This follows the principle of calibrated trust, where users receive intelligible feedback about system behavior [8, 18]. In addition, the interface would support *selective disclosure and contestability*, as it summarizes relevant outcomes, reveals trade-offs when behavior deviates from expectations, and allows users to inspect and revise decision policies.

2.3 Lifecycle of a Personal Mobility Agent

From the user’s perspective, the personal mobility agent might be created and refined in three phases:

- (1) During setup, the agent initializes the *preference model* using simple onboarding questions (e.g., rating comfort versus efficiency preferences in traffic), demographics, personality tests, and/or a short calibration drive or simulation.

- (2) The agent then adapts online during regular use. It observes how the user reacts to different maneuvers, such as whether they accept or override certain behaviors, and occasionally asks for quick feedback.
- (3) The agent is updated over longer time scales. It detects changes in user behavior, allows explicit policy edits (e.g., “be more assertive in light traffic”), and periodically recalibrates its preference model. The agent is therefore not a fixed model trained once, but an evolving representation that is refined both offline and online while serving as the user’s proxy.

3 Design Tensions and Research Directions

Introducing personal mobility agents that act as proxies for users shifts interaction design away from isolated control interfaces toward the design of *negotiation representations* and *delegation policies*. In this model, the user is no longer directly involved in each traffic interaction. Instead, their proxy negotiates on their behalf. This creates several design tensions that may become relevant in future mobility systems and automation experiences.

3.1 Designing for Preference Alignment Under Legal and Safety Constraints

Mobility systems are constrained by strict safety and legal requirements, which may force the proxy to deviate from users’ preferences. For example, a user may prefer small gaps when merging, but the proxy must still enforce safety thresholds.

Prior negotiation-focused HMI approaches often attempt to resolve such conflicts through explanation, persuasion, or adaptive dialogue [15, 16]. In a proxy-interaction model, however, the design problem shifts from resolving each conflict interactively in the moment to constraining policy space. This raises the question of how the agent should communicate differences between the user’s stated preferences and the behavior it executes, and whether such communication would undermine the intended benefits of relieving users from direct interaction in momentary situations.

3.2 Preventing Strategic Manipulation While Preserving Personalization

Research on LLM-mediated negotiation suggests that users appreciate adaptive communication, but may also attempt to probe or outsmart the system [15]. This issue may become more critical with persistent agents that accumulate personalization over time and repeatedly act on the user’s behalf. Research on delegation in human–automation decision contexts suggests that users’ willingness to delegate depends on perceived competence, controllability, and accountability of the delegated system [14]. Persistent proxy agents, therefore, introduce long-term governance challenges beyond momentary adaptation.

A possible direction is to separate *interaction adaptivity* (how the agent communicates) from *policy adaptivity* (what the agent is allowed to do). This suggests design requirements for manipulation-resistant preference capture, non-negotiable safety rules, and interaction patterns that discourage adversarial bargaining.

3.3 Scaling Agent Negotiation Without Breaking Mixed-Traffic Interaction

Personal mobility agents could enable efficient and effective negotiation among automated actors by exchanging preferences and constraints as data points. However, traffic will likely remain a mixed environment in which many actors lack such an agent and behave unpredictably [3]. A proxy must therefore interact not only with other agents but also with human drivers, cyclists, and pedestrians, whose intentions must be inferred from their behavior. Prior work on eHMIs has highlighted that not all actors share symmetric communication channels or shared representations [3].

Proxy-based coordination, therefore, cannot assume universal agent participation and must degrade gracefully to behavior-based inference.

This creates a tension between optimizing agent-agent negotiations and maintaining robust interaction with human actors. In addition, proxy-based negotiation may obscure responsibility when outcomes emerge from multiple interacting agents. A possible research direction is the design of *attribution mechanisms* that explain whether a behavior was driven by safety constraints, the user’s delegation policy, inferred human behavior, or a negotiated compromise with other proxies. This aligns with concerns in multi-agent automation about distributed responsibility and attribution when outcomes emerge from interacting autonomous systems rather than a single decision-maker (e.g., see [1]).

3.4 Making Delegation Legible: Transparency, Agency, and Accountability

Proxy-based interactions change not only system behavior but also the user’s experience with automation. When a proxy negotiates most situations autonomously, users may become less aware of how decisions are made, which trade-offs are applied, or which constraints shape outcomes. In some cases, this reduced involvement may lower cognitive load and might even increase acceptance because the system “just works.” In other cases, however, it may lead to loss of agency, miscalibrated trust, or difficulties in understanding why certain outcomes occurred.

This creates an ethical and experiential tension between *convenient delegation* and *meaningful awareness*. Designers must decide how visible proxy decisions should be. One option is minimal disclosure, where the proxy operates largely in the background and only surfaces important deviations or safety-critical events. Another option is reflective transparency, where the system periodically exposes negotiation summaries or policy trade-offs to support user understanding and long-term trust calibration.

This tension raises several open questions: Do users prefer to remain largely unaware of negotiation details as long as outcomes match their expectations? At what point does reduced awareness undermine perceived agency or accountability? And how should proxy systems balance cognitive load reduction with ethical requirements for transparency, contestability, and informed consent?

4 Conclusion

This position paper argues that traffic in the near future will become too heterogeneous and automated for users to meaningfully negotiate each traffic situation themselves. As automated vehicles, micromobility, robots, and connected infrastructure interact in shared spaces, conflicts between actors with different goals and preferences become inevitable, while traditional moment-to-moment human negotiation strategies struggle to scale. Instead, *personal mobility agents* may act as proxies for users. These agents encode user preferences and constraints and negotiate concrete maneuvers with other actors under shared safety rules.

This shift raises relevant research questions for HCI. Proxy-based interactions may reduce cognitive load and improve acceptance, but they may also change how users perceive agency, responsibility, and trust when most negotiations occur without their direct involvement. Understanding how to design delegation interfaces, how much awareness users need, and how proxy-mediated experiences affect long-term trust and accountability remains an open research agenda for automated mobility. The proposed proxy-mediated negotiation model extends beyond mobility to agentic automation more broadly, including domains such as domestic robotics, smart infrastructure, and collaborative AI systems, in which users delegate decision-making to interacting agents.

References

- [1] Aline Belloni, Alain Berger, Olivier Boissier, Grégory Bonnet, Gauvain Bourgne, Pierre-Antoine Chardel, Jean-Pierre Cotton, Nicolas Evreux, Jean-Gabriel Ganascia, Philippe Jaillon, et al. 2015. Dealing with Ethical Conflicts in Autonomous Agents and Multi-Agent Systems.. In *AAAI Workshop: AI and Ethics*.
- [2] Mark Colley, Pascal Jansen, Mugdha Keskar, and Enrico Rukzio. 2025. Improving external communication of automated vehicles using bayesian optimization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [3] Joost de Winter et al. 2022. External human-machine interfaces: Gimmick or necessity? *Current Opinion in Psychology* (2022). doi:10.1016/j.copsyc.2022.101051
- [4] Mica R. Endsley. 1995. Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors* 37, 1 (1995), 32–64. doi:10.1518/001872095779049543
- [5] Federal Republic of Germany. 2021. Gesetz zur Änderung des Straßenverkehrsgesetzes und des Pflichtversicherungsgesetzes – Gesetz zum autonomen Fahren (BGBl. I S. 3108). https://dejure.org/BGBl/2021/BGBl_I_S_3108. Amendment to the Road Traffic Act establishing the legal framework for autonomous driving functions (Level 4) that entered into force on 28 July 2021.
- [6] Lorena Hell, Janis Sprenger, Matthias Klusch, Yoshiyuki Kobayashi, and Christian Müller. 2021. Pedestrian behavior in japan and germany: A review. In *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 1529–1536.
- [7] Pascal Jansen, Mark Colley, Svenja Krauß, Daniel Hirsche, and Enrico Rukzio. 2025. Opticarvis: Improving automated vehicle functionality visualizations using bayesian optimization to enhance user experience. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–23.
- [8] John D. Lee and Katrina A. See. 2004. Trust in Automation: Designing for Appropriate Reliance. *Human Factors* 46, 1 (2004), 50–80. doi:10.1518/hfes.46.1.50.30392
- [9] Xishun Liao, Zhouqiao Zhao, Matthew J. Barth, Amr Abdelraouf, Rohit Gupta, Kyungtae Han, Jiaqi Ma, and Guoyuan Wu. 2025. A Review of Personalization in Driving Behavior: Dataset, Modeling, and Validation. *IEEE Transactions on Intelligent Vehicles* 10, 2 (2025), 1241–1262. doi:10.1109/TIV.2024.3425647
- [10] Régis Lobjois and Viola Cavallo. 2007. Age-related differences in street-crossing decisions: The effects of vehicle speed and time constraints on gap selection in an estimation task. *Accident analysis & prevention* 39, 5 (2007), 934–943.
- [11] OpenAI. 2025. Introducing Operator. <https://openai.com/index/introducing-operator/>. Accessed 2026-02-15.
- [12] World Health Organization. 2024. *Global status report on road safety 2023: country and territory profiles*. World Health Organization.
- [13] Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. doi:10.1145/3586183.3606763
- [14] Philipp Spitzer, Joshua Holstein, Patrick Hemmer, Michael Vössing, Niklas Kühn, Dominik Martin, and Gerhard Satzger. 2025. Human delegation behavior in human-AI collaboration: The effect of contextual information. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–28.
- [15] Annika Stampf, Mark Colley, Bettina Girst, and Enrico Rukzio. 2024. Exploring Passenger-Automated Vehicle Negotiation Utilizing Large Language Models for Natural Interaction. In *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 350–362.
- [16] Annika Stampf, Mark Colley, Ann-Kathrin Knuth, Cagla Tasci, and Enrico Rukzio. 2024. Examining Psychological Conflict-Handling Strategies for Highly Automated Vehicles to Resolve Legal User-Vehicle Conflicts. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 3 (2024), 1–25.
- [17] Laurin Vasile, Barbara Seitz, Verena Staab, Magnus Liebherr, Christoph Däsch, and Dieter Schramm. 2023. Influences of personal driving styles and experienced system characteristics on driving style preferences in automated driving. *Applied Sciences* 13, 15 (2023), 8855.
- [18] Manfred Wischnewski et al. 2023. Measuring and Understanding Trust Calibrations for Automated Systems: A Survey of Empirical Human-Subject Studies. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*.