ACM GREEKCHI-CYCHI SYMPOSIUM 2026

SPACER: NEGOTIATING WITHOUT REVEALING PREFERENCES

Panteleimon Stanimeros[1,2], Chairi Kiourt[1], Alexandros Karakasidis[1,2]

[1]Athena Research Center [2]University of Macedonia

p.stanimeros@athenarc.gr, chairiq@athenarc.gr, a.karakasidis@uom.edu.gr

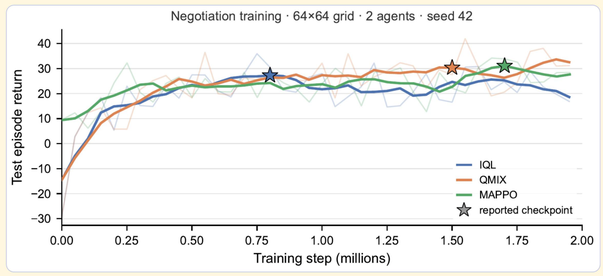
The problem

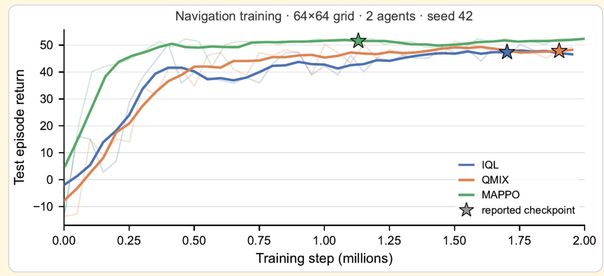


Computer-Human Interaction is not limited to graphical interfaces.

- A human and an autonomous taxi must agree on a meeting point, and neither wants to say why one spot suits it better.
- Each one sets a preference privately and negotiates directly, exchanging only points, never scores.
- But hiding scores doesn't hide preferences: which points, in what order, and the route taken all leak them away.

Training





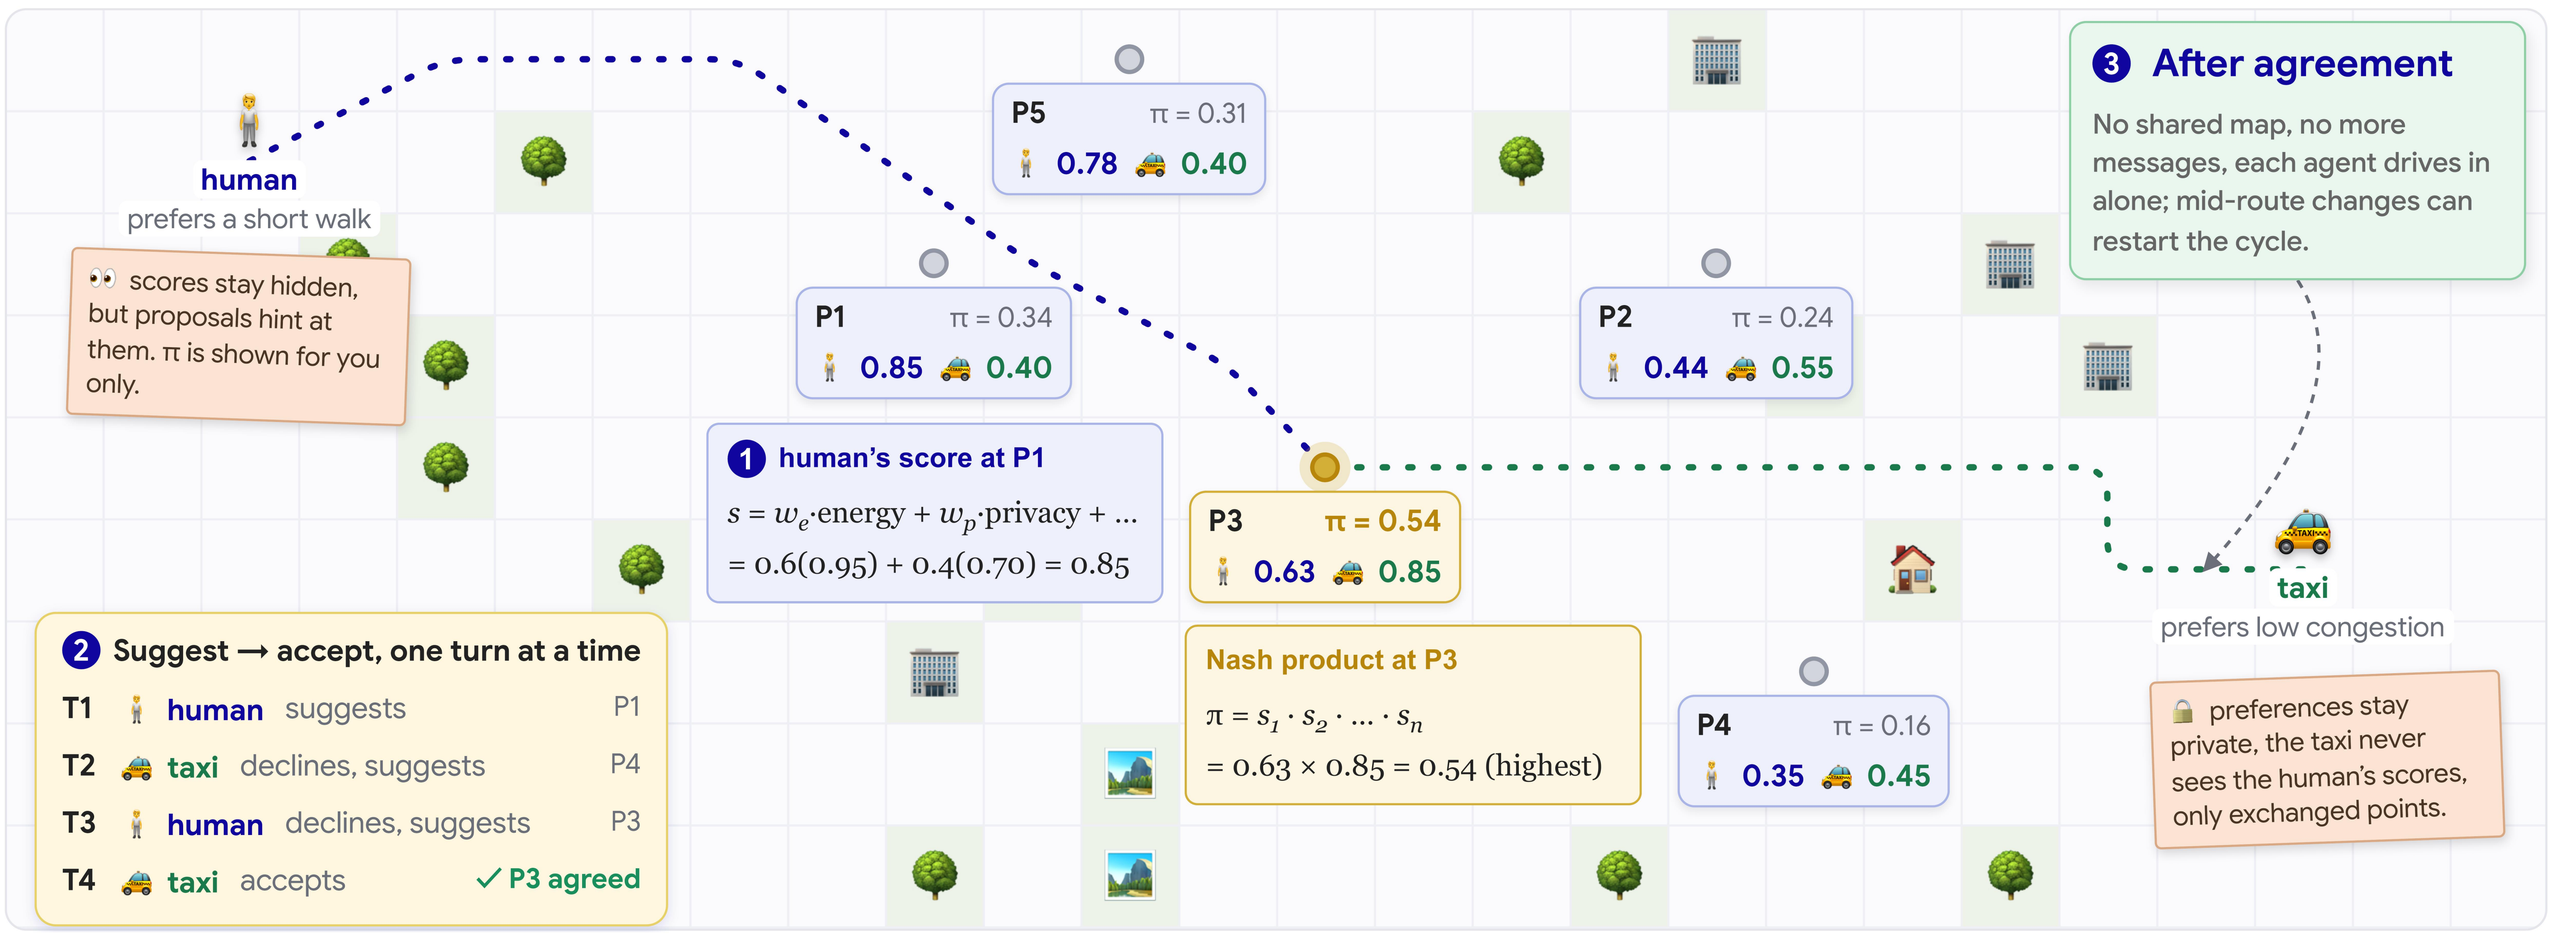
- 64×64 grid, 10% obstacles, 5 candidate points. Fresh random preferences every episode.
- Negotiation: 6-dim observation (5 private scores + peer's last suggestion), 6 actions, 24-turn cap.
- Navigation: 7-dim observation (direction, distance, 4-way lidar), the same input the baselines get.
- IQL, QMIX and MAPPO in EPyMARL, each trained alone for 2M steps, seeds 42/43/44, 100/300 eval episodes.

1 Evaluation — Each participant privately scores every candidate point using its own hidden goals and constraints.

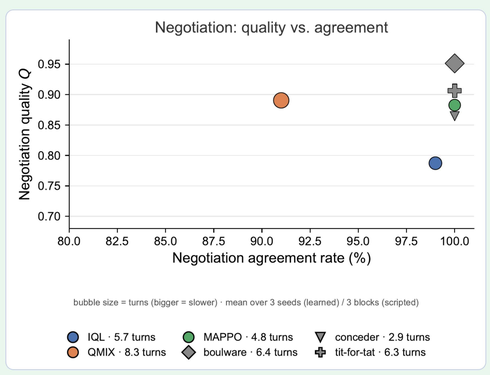
→ **2 Negotiation** — They exchange only point indices, converging on the point that maximises their scores.

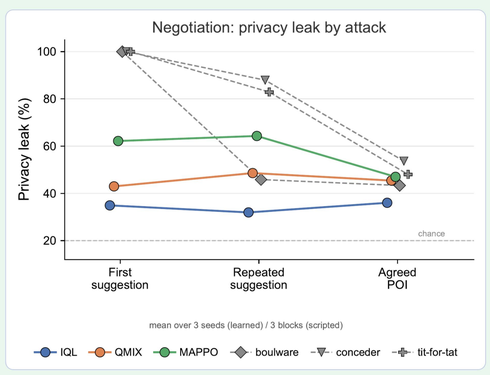
→ **3 Navigation** — Once agreed, each one navigates independently to the shared point, using only its own local view.

environment or preferences change



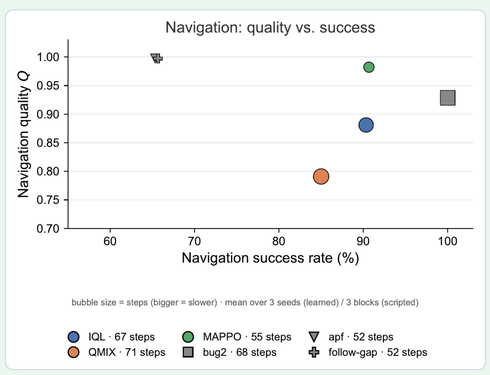
Early results: Negotiation (Stage 2)

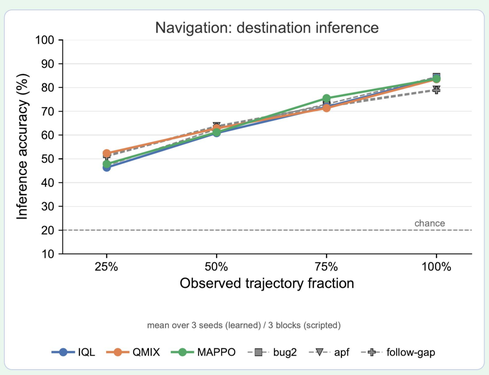




- Learned agreement and quality land close to scripted: 91-100% agreement, 0.79-0.89 quality vs. scripted's 0.86-0.95.
- But learned doesn't leak nearly as much: opening-offer leak is 35-62% against scripted's 100% (chance is 20%).

Early results: Navigation (Stage 3)





- Learned matches or beats scripted on success and straightness (MAPPO: 90% success, 0.98 quality).
- Route leak is similar either way: 46-52% a quarter-way in, rising to 79-85% on the full route.

Holding the leak down

Can an agent be trained to push the leak back toward chance without losing agreement quality?

Leakage over more episodes

Is this leak a one-shot exposure, or does it compound across repeated negotiations with the same partner?

A malicious agent

What if an agent negotiates not to agree, but to actively extract the other's private preferences?

Selected References

P. Stanimeros, C. Kiourt, and A. Karakasidis. Toward Decentralized Spatial Coordination. *ACM SIGSPATIAL, Vision Paper Track* (to appear).

P. Stanimeros and A. Karakasidis. 2025. Near for All: Privacy-Preserving Meetings for Human-Autonomous Vehicles. *SN Computer Science*.

B. Rani. 2026. Behavioral Privacy Leakage in Agentic Negotiation: Formalizing and Mitigating Inference Attacks. *AI4TCI Workshop, ARES*.

S. Farhadi and N. R. Jennings. 2021. A Faithful Mechanism for Incremental Multi-Agent Agreement Problems. *SN Computer Science*.