Early Prediction of Future Behavioral Strategy from Process Traces

Robert Kasumba, Dennis L. Barbour, Chien-Ju Ho

Washington University in St. Louis

NeurIPS 2026

Adaptive AI must anticipate how people will act



Tutors

How will this learner approach a new problem?

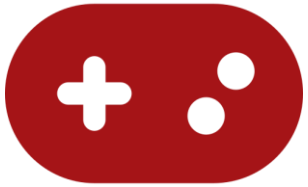

Games

How will this player tackle the next level?

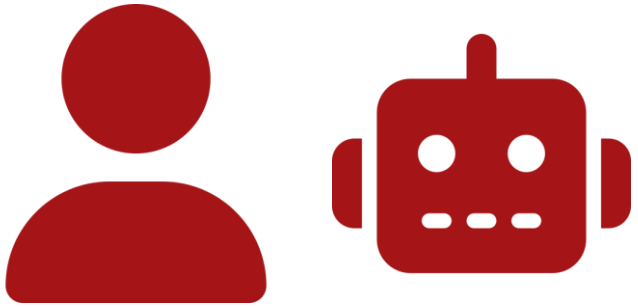

Human-AI teams

Will my partner stick with the plan or switch?

They have to adapt **before** seeing the person in the new task.

VanLehn, 2011; Paraschos & Koulouriotis, 2023; Hoffman et al., 2024; Marticorena et al., 2025

Problem: predict strategy in a task not yet seen

Can how people behaved in earlier tasks predict the strategy they will use in a new one?

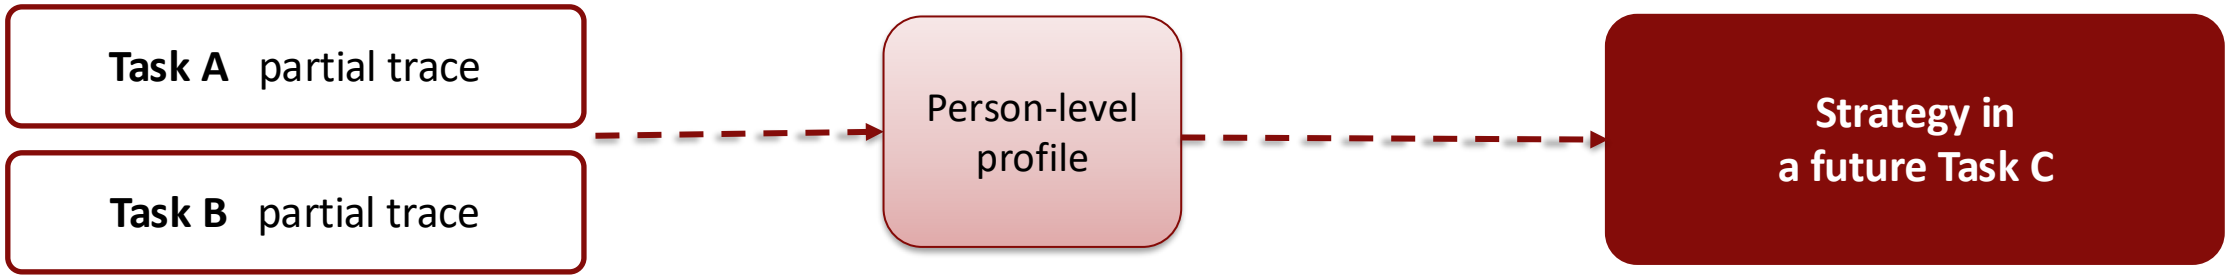

Outcomes hide process

Similar scores can come from very different strategies.

One task is not enough

A single trace mixes the person with that task's layout.

Waiting is costly

Watching the new task first delays any adaptation.

PLVM: Process-Level Latent Variable Model

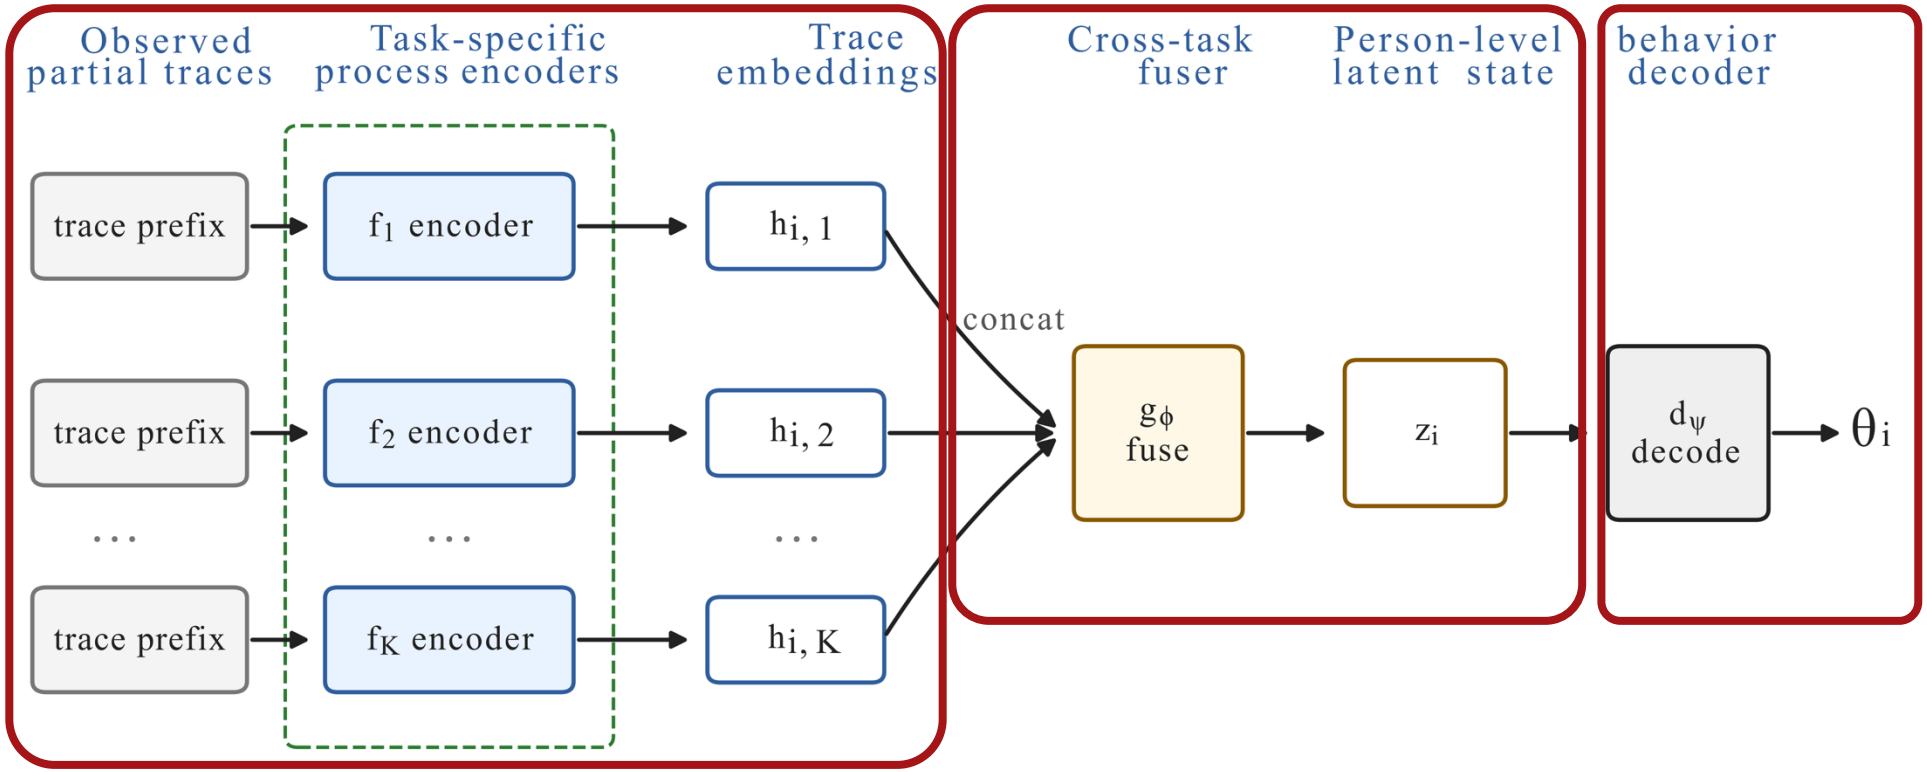

1 Encode

One sequence encoder per source task, trained without the target label

2 Fuse

Combine trace embeddings into one person-level profile *z*

3 Predict

Decode *z* into the strategy in the held-out task

Predict strategy in an unseen PowerWash level



What players see

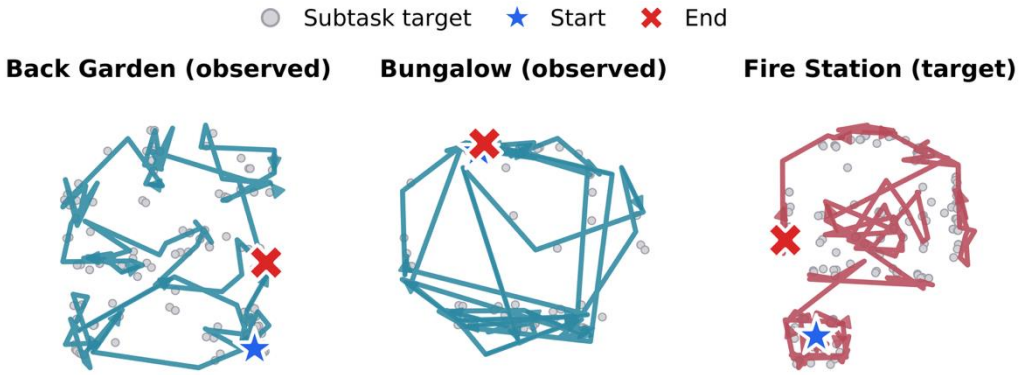

- **11,080 players** [1] clean every part of each level, in any order
- Observe up to **2 h** of two **early levels**
- Predict behavior strategy (**Planner vs. Hopper)** in a future unseen **target** level

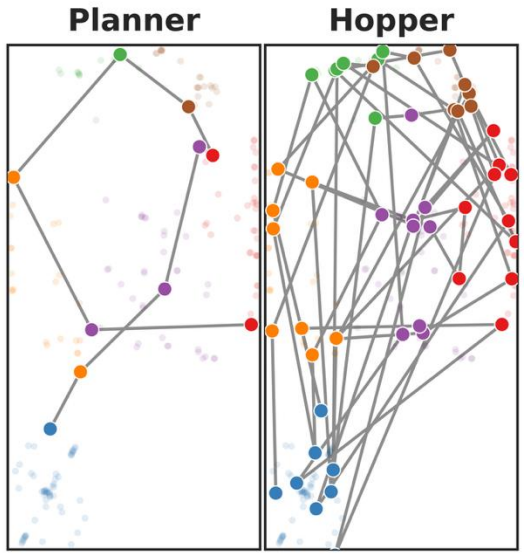

Planner: one zone at a time
Hopper: keeps switching zones

[1] Vuorre et al., 2023, Scientific Data

Process traces predict strategy; outcomes don't

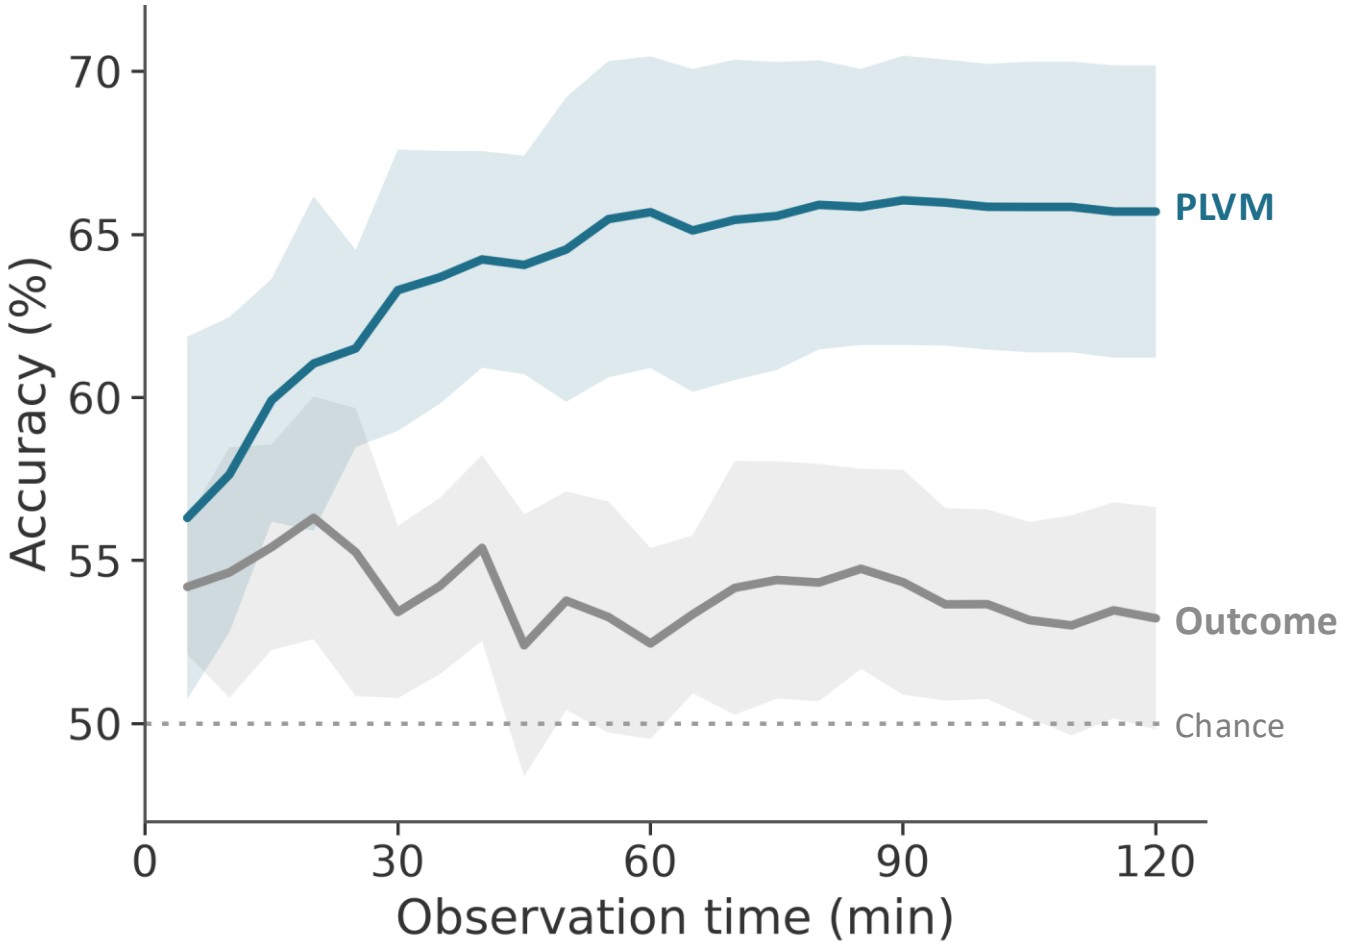

+12.5 pp

PLVM over outcome-only models (65.7% vs. 53.2% at 2 h)

Outcome-only models see scores and progress from the same 2 h, and stay near chance

Cross-task fusion helps most when observed tasks are complementary

Setup: simulated agents with a known hidden type, each observed in two tasks

Complementary: each task reveals a different aspect of the type

Redundant: both tasks reveal the same aspect

We predict the hidden type from partial behavior in both tasks

68% → 87%

Hidden-type recovery with complementary tasks: either task alone → **PLVM** fusing both

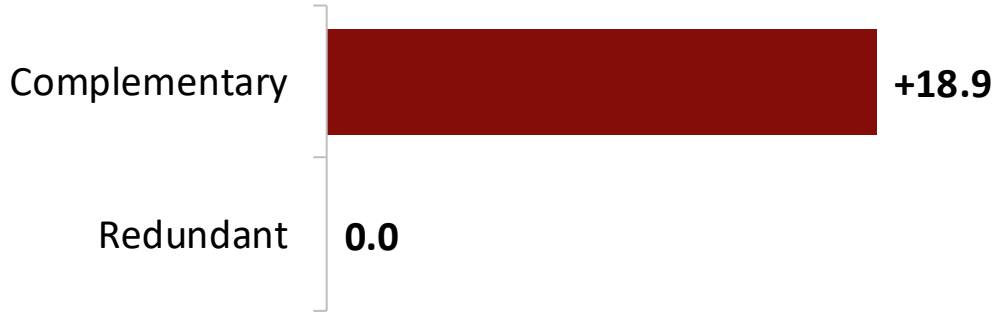

Fusion is most powerful when **no single task tells the whole story**.

Fusing levels beats either level alone with real-human data

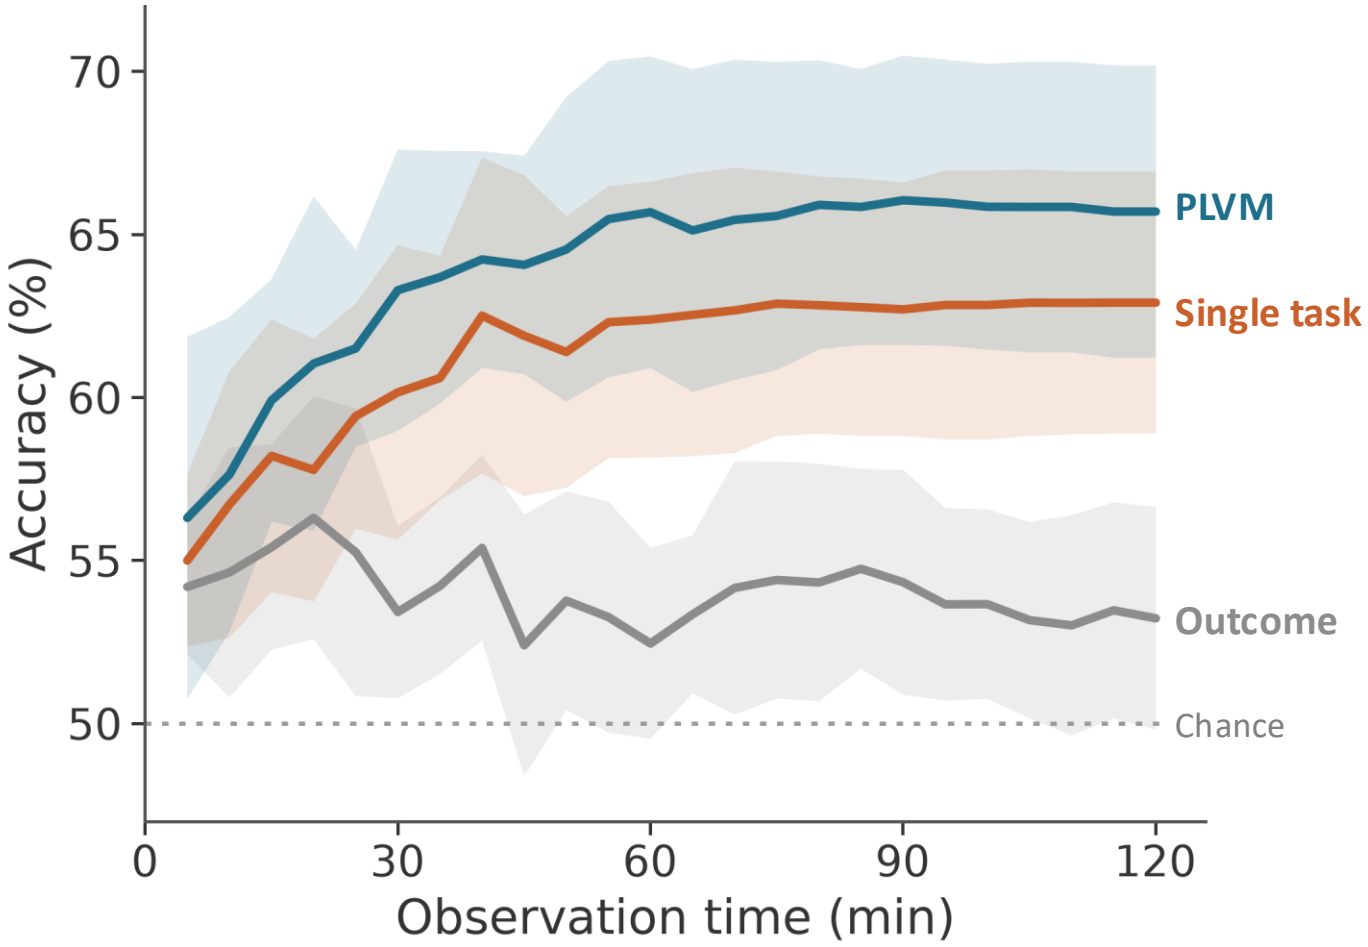

+2.8 pp

PLVM over the best single level (65.7% vs. 62.9% at 2 h)

PLVM is ahead of **both** single-level models at **every** observation time

Takeaways

How people play predicts their future strategy; **what** they score barely does.

Fusing process traces across tasks beats any single task, in simulation **and** in real gameplay.

Fusion is most powerful when tasks reveal **complementary** parts of the latent profile.

arxiv.org/abs/2605.30550

References

1. M. Vuorre, K. Magnusson, N. Johannes, J. Butlin, and A. K. Przybylski. An intensive longitudinal dataset of in-game player behaviour and well-being in PowerWash Simulator. *Scientific Data*, 10:622, 2023.
2. K. VanLehn. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46:197–221, 2011.
3. P. D. Paraschos and D. E. Koulouriotis. Game difficulty adaptation and experience personalization: A literature review. *International Journal of Human–Computer Interaction*, 39:1–22, 2023.
4. G. Hoffman, T. Bhattacharjee, and S. Nikolaidis. Inferring human intent and predicting human action in human–robot collaboration. *Annual Review of Control, Robotics, and Autonomous Systems*, 7:73–95, 2024.
5. D. C. P. Marticorena, Z. Lu, C. Wissmann, Y. Agarwal, D. Garrison, J. M. Zempel, and D. L. Barbour. Immersive virtual games: winners for deep cognitive assessment. *arXiv:2502.10290*, 2025.
6. R. Kasumba, D. C. P. Marticorena, A. Pahor, G. Ramani, I. Goffney, S. M. Jaeggi, A. R. Seitz, J. R. Gardner, and D. L. Barbour. Distributional latent variable models with an application in active cognitive testing. *IEEE Transactions on Cognitive and Developmental Systems*, 17(5):1212–1222, 2025.
7. R. Kasumba, Z. Lu, D. C. P. Marticorena, M. Zhong, P. Beggs, A. Pahor, G. Ramani, I. Goffney, S. M. Jaeggi, A. R. Seitz, J. R. Gardner, and D. L. Barbour. Bayesian distributional models of executive functioning. *Behavior Research Methods*, 58(11):310, 2026.
8. R. Kasumba, D. L. Barbour, and C.-J. Ho. Early prediction of future behavioral strategy from process traces. *NeurIPS*, 2026. arXiv:2605.30550.

Acknowledgments

Co-authors

Dennis L. Barbour

Chien-Ju Ho

Funding

McDonnell International Scholars Academy

Biswas Family Foundation

Data

PowerWash Simulator open dataset

Vuorre et al., 2023

arxiv.org/abs/2605.30550