Understanding Human Limits in Pattern Recognition: A Computational Model of Sequential Reasoning in Rock, Paper, Scissors

Logan Cross*, Erik Brockbank*, Tobias Gerstenberg, Judith E. Fan, Daniel L. K. Yamins, & Nick Haber

What are the bottlenecks in human sequential reasoning?

We apply a LLM-based cognitive model of social reasoning to the problem of predicting an opponent's next move in rock, paper, scissors.

The model captures ***human performance*** against algorithmic opponents & suggests that ***hypothesis generation*** is the primary computational bottleneck in this sequential pattern recognition setting.

METHODS

Capturing human pattern recognition with Rock, Paper, Scissors

300 rounds against an algorithmic opponent[1]

- Humans readily exploit bots with simple transition strategies ●●●●
- Humans are near chance against bots with more complex policies ●●●

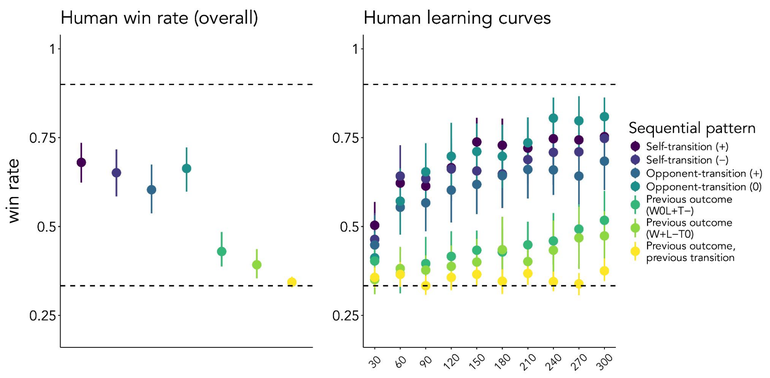

Bot Strategies Were Based On:

- 1. **Previous Move** (●●●●): Four bots reacted to the last move of *self* or *opponent*
  - *Examples: Always countering the player's last move, or copying the player's last move*
- 2. **Previous Outcome** (●●): Two bots changed their strategy based on whether they won, lost, or tied the last round
- 3. **Previous Outcome & Transition** (●): The most complex bot's strategy depended on both the previous outcome and its own last move type

Transitions

Move n \ Move $n+1$	R	P	S
R	o	+	–
P	–	o	+
S	+	–	o

A computational model of human social reasoning in adversarial settings

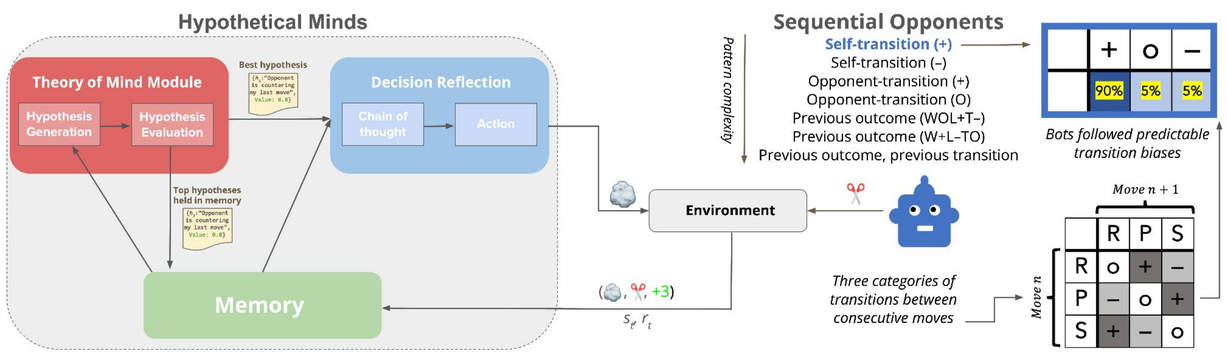

Hypothetical Minds LLM agent[2]

- **Theory of Mind module**
  - **Generates hypotheses** about opponent strategies in natural language
  - **Hypothesis Evaluation**: Scores hypotheses based on predicting behavior
- **Decision Reflection component**: uses best hypothesis to determine the best counter-move

RESULTS

Hypothetical Minds captures human sequential reasoning: similar performance against algorithmic opponents

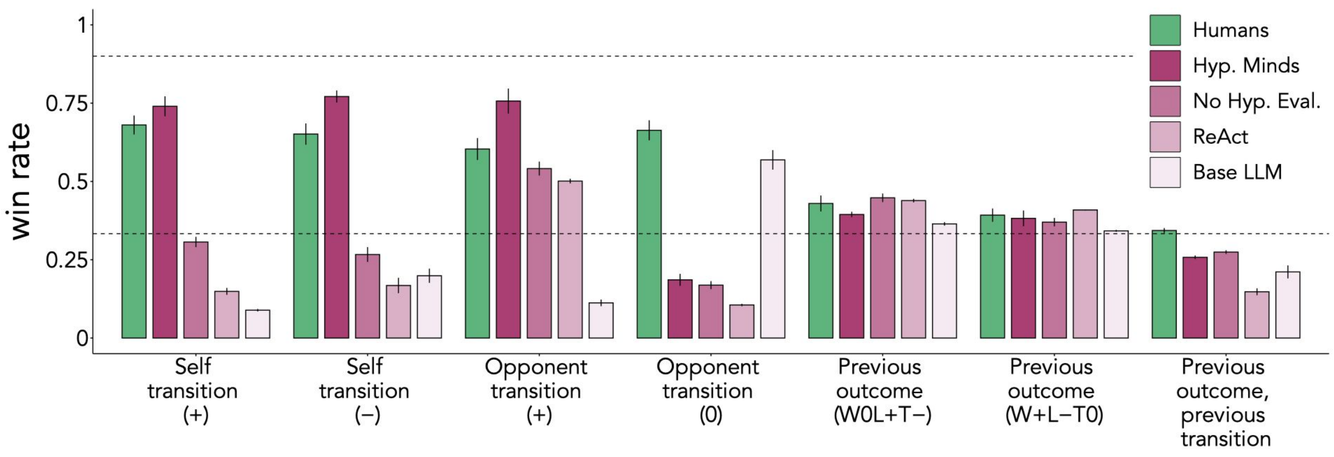

Hypothetical Minds model ■ closely mirrors human win rates ■ against 6 / 7 algorithmic opponents

Ablated models ■ ■ ■ that lack Theory of Mind reasoning fail to capture human behavior

Hypothesis generation is key bottleneck: removing hypothesis generation leads to near-ceiling performance

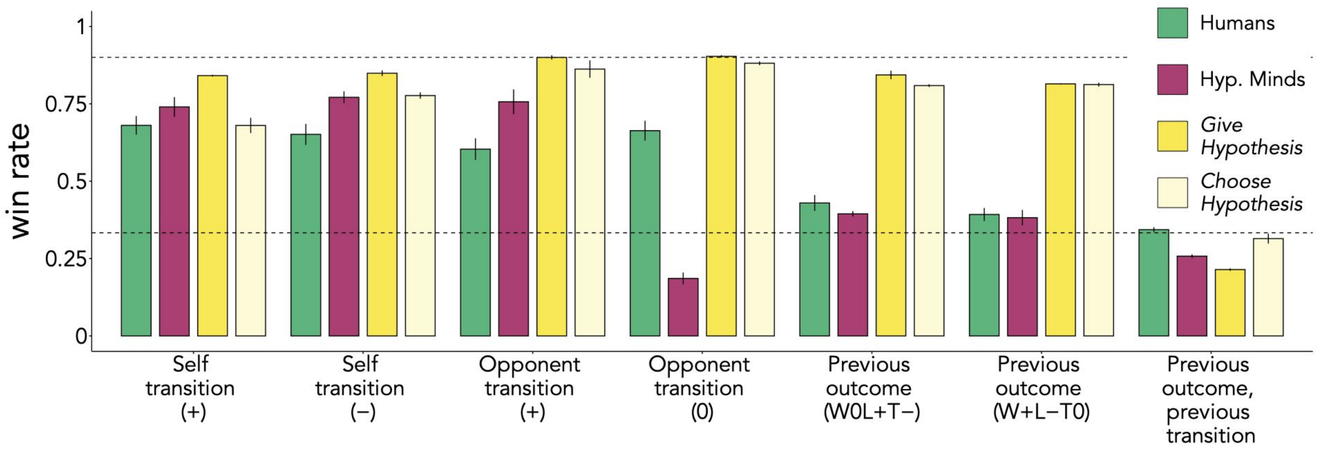

Model performs near ceiling when given a verbal description of the opponent's strategy ■ or a description of candidate opponent strategies ■

Overcoming the bottleneck: pedagogically-inspired interventions help the model discover opponent strategies

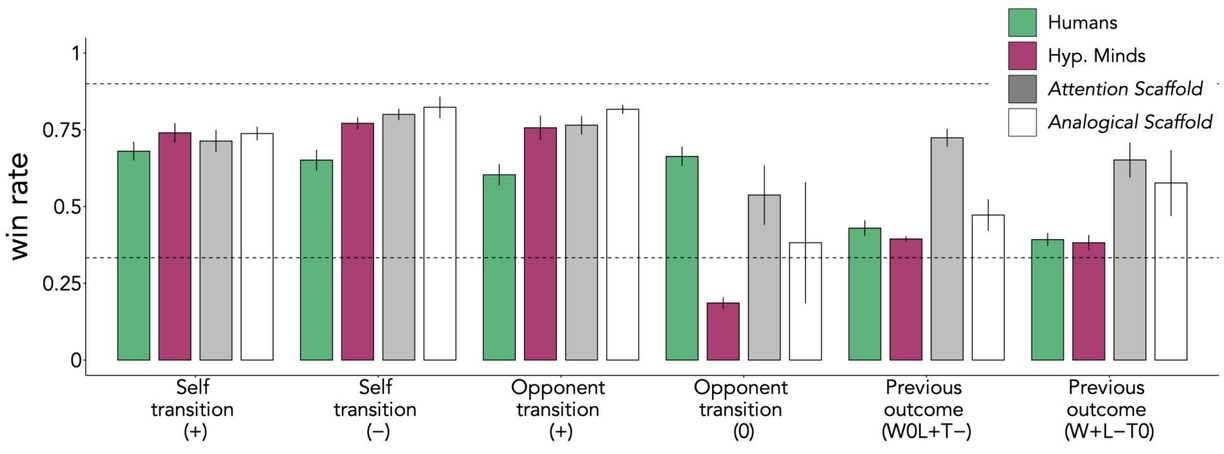

Model shows significant improvement against hyp. gen. bottlenecked strategies when given instruction to attend to features of opponent strategy ■ or similar example strategies ■

References

1. Brockbank, E., & Vul, E. (2024). Repeated rock, paper, scissors play reveals limits in adaptive sequential behavior. Cognitive Psychology, 151, 101654.
2. Cross, L., Xiang, V., Bhatia, A., Yamins, D. L. K., & Haber, N. (2025, April). Hypothetical minds: Scaffolding theory of mind for multi-agent tasks with large language models. In Proceedings of the International Conference on Learning Representations (ICLR 2025).