Lecture 14:

Problem Solving Agents (Part 2)

Visual Computing Systems
Stanford CS348K, Spring 2025

What could we do with humanlike agents?

Irregardless of whether they are LLM-based or generated using other ML technologies like reinforcement learning or imitation learning

“Bots”: virtual characters in interactive experiences

- Virtual teammates in team-based games, or competitors/enemies
- Narrative elements

Learning to race in Gran Turismo Sport
[Wurman et al. 2022]

Aside… embracing the fallibility of AI in game design

A clever example…

[💥OG] Be NPC or DIE!💢

By arashiyama

Maturity: Mild

Favorite 419K+ 31K+

Predicting behavior of real human players

- We've seen examples of predicting behavior to enable testing of game designs
  - Flappy bird

- How did the Flappy bird authors "model" human behavior?

- See next slide for another example…

Given enough training, an agent can learn to successfully navigate the skinny path, even though that path affords a low margin for movement error.

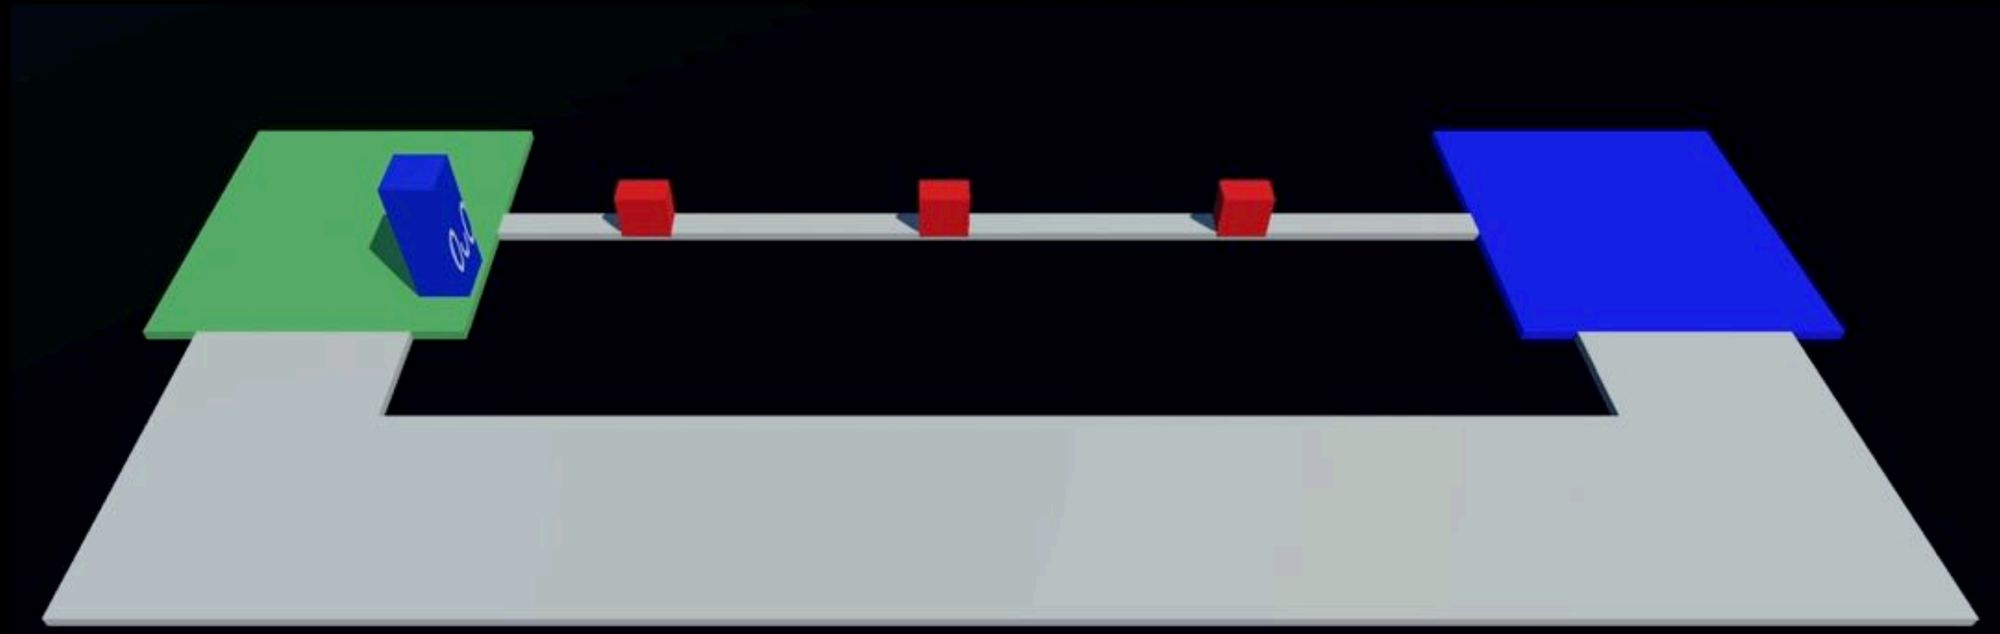

Randomly changing the agent's agents (with probability *p*) reduces success rate to near 0.

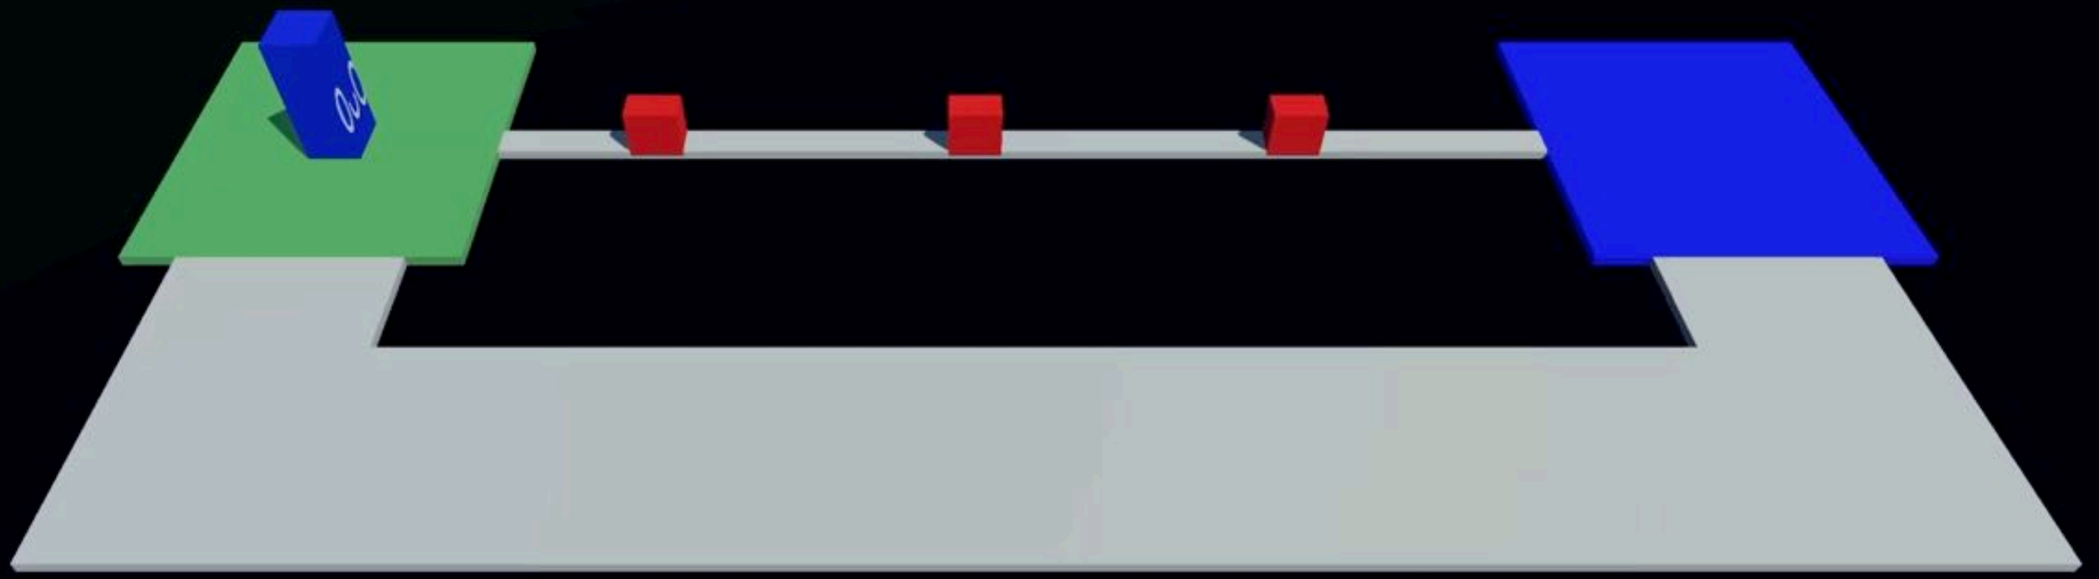

Training an agent that is “aware of” its own limitations (potential for occasional action perturbation) yields a policy that wisely takes a more conservative approach.

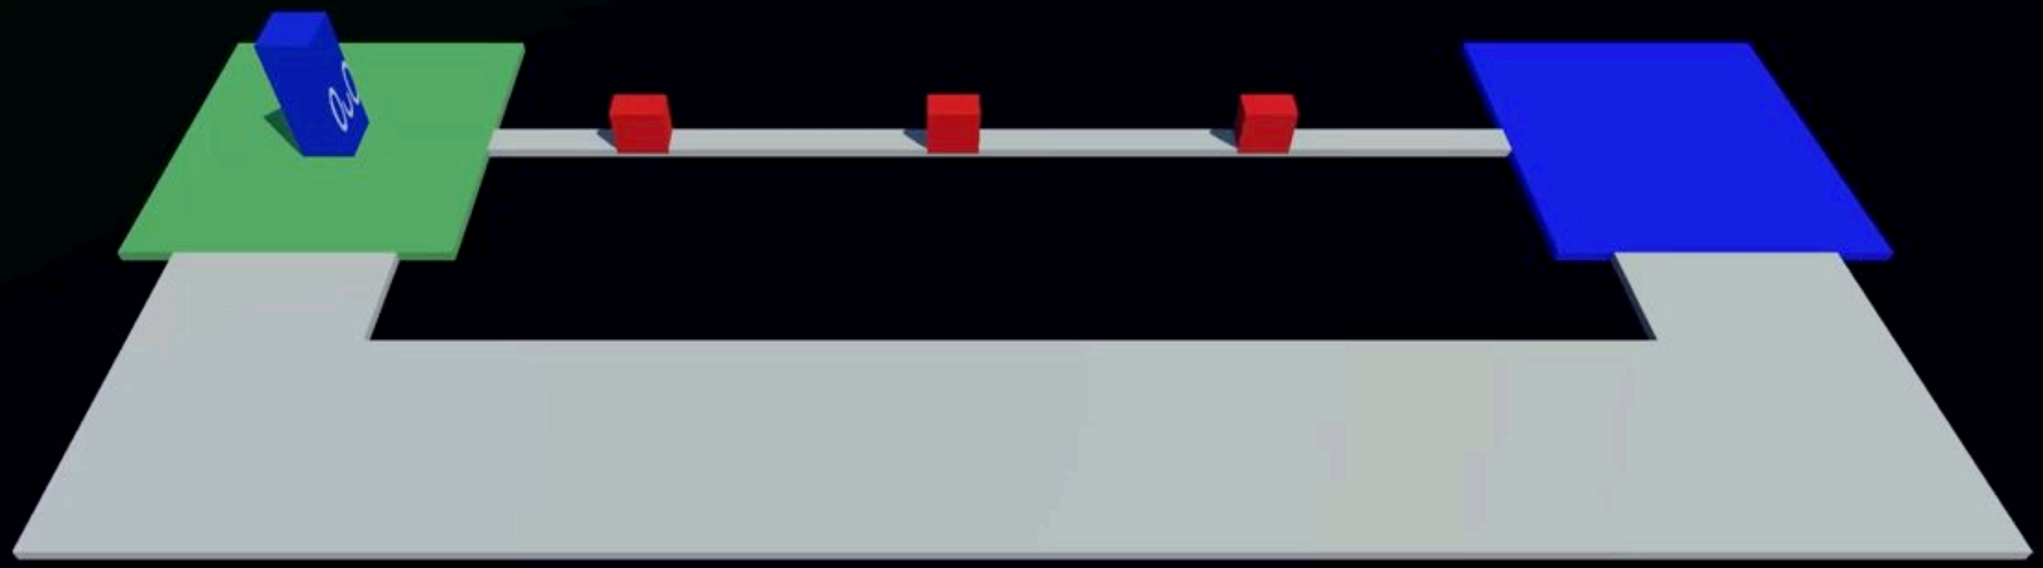

Another example: Counterstrike agents

Dataset of human play

- 123 hours
- 16 Hz data point frequency
- 2292 unique players
- 513K shots
- 29K eliminations

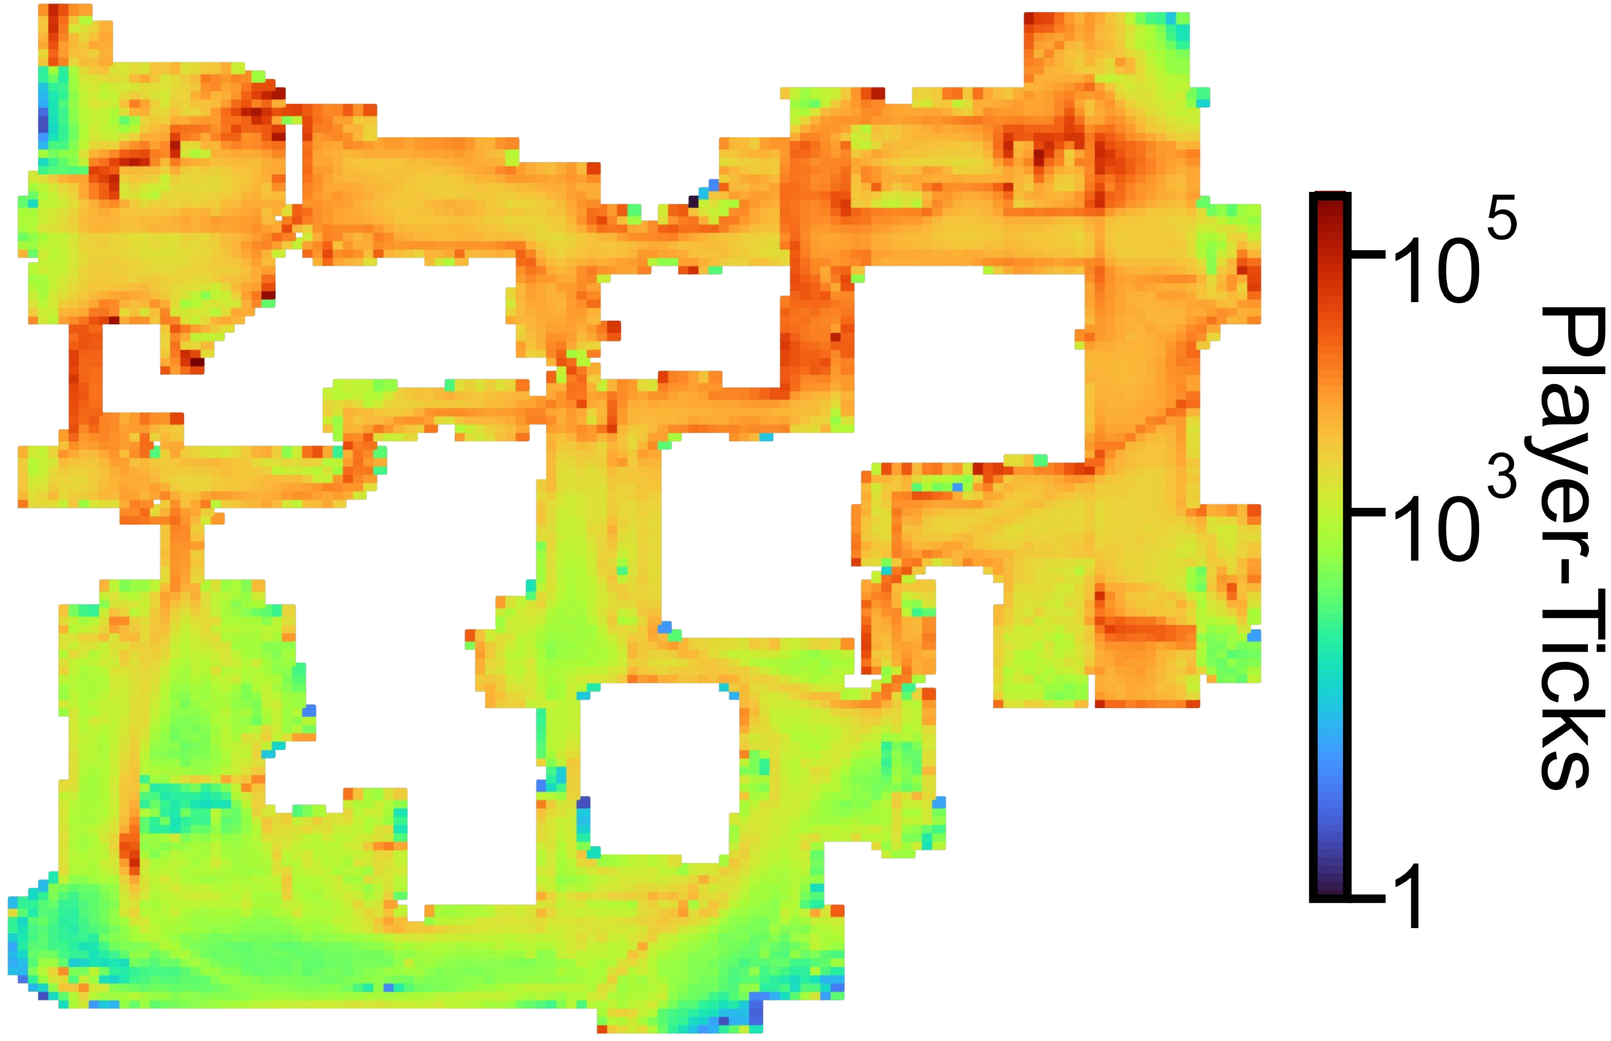

CGCO experts watched videos of play (sometimes human, sometimes bots), and asked to identify which video depicts how "they'd expect humans to move in that situation".

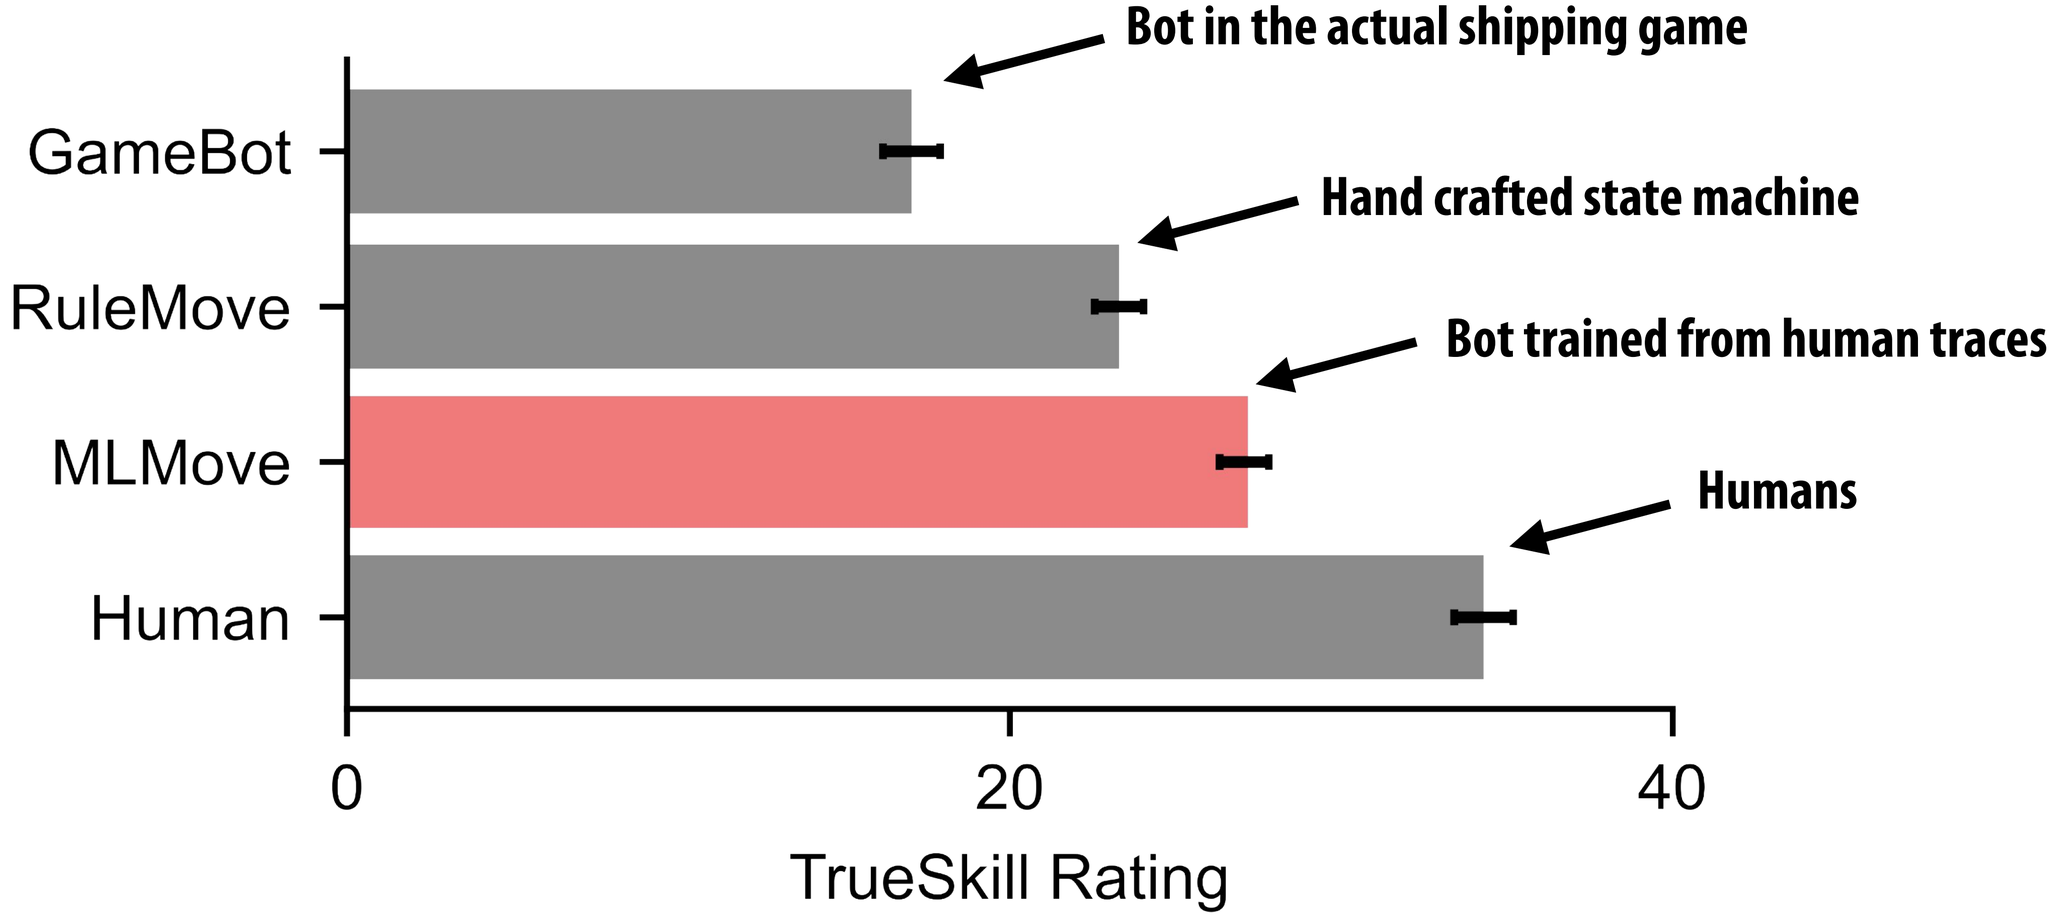

MLMove and Human Positioning Similar

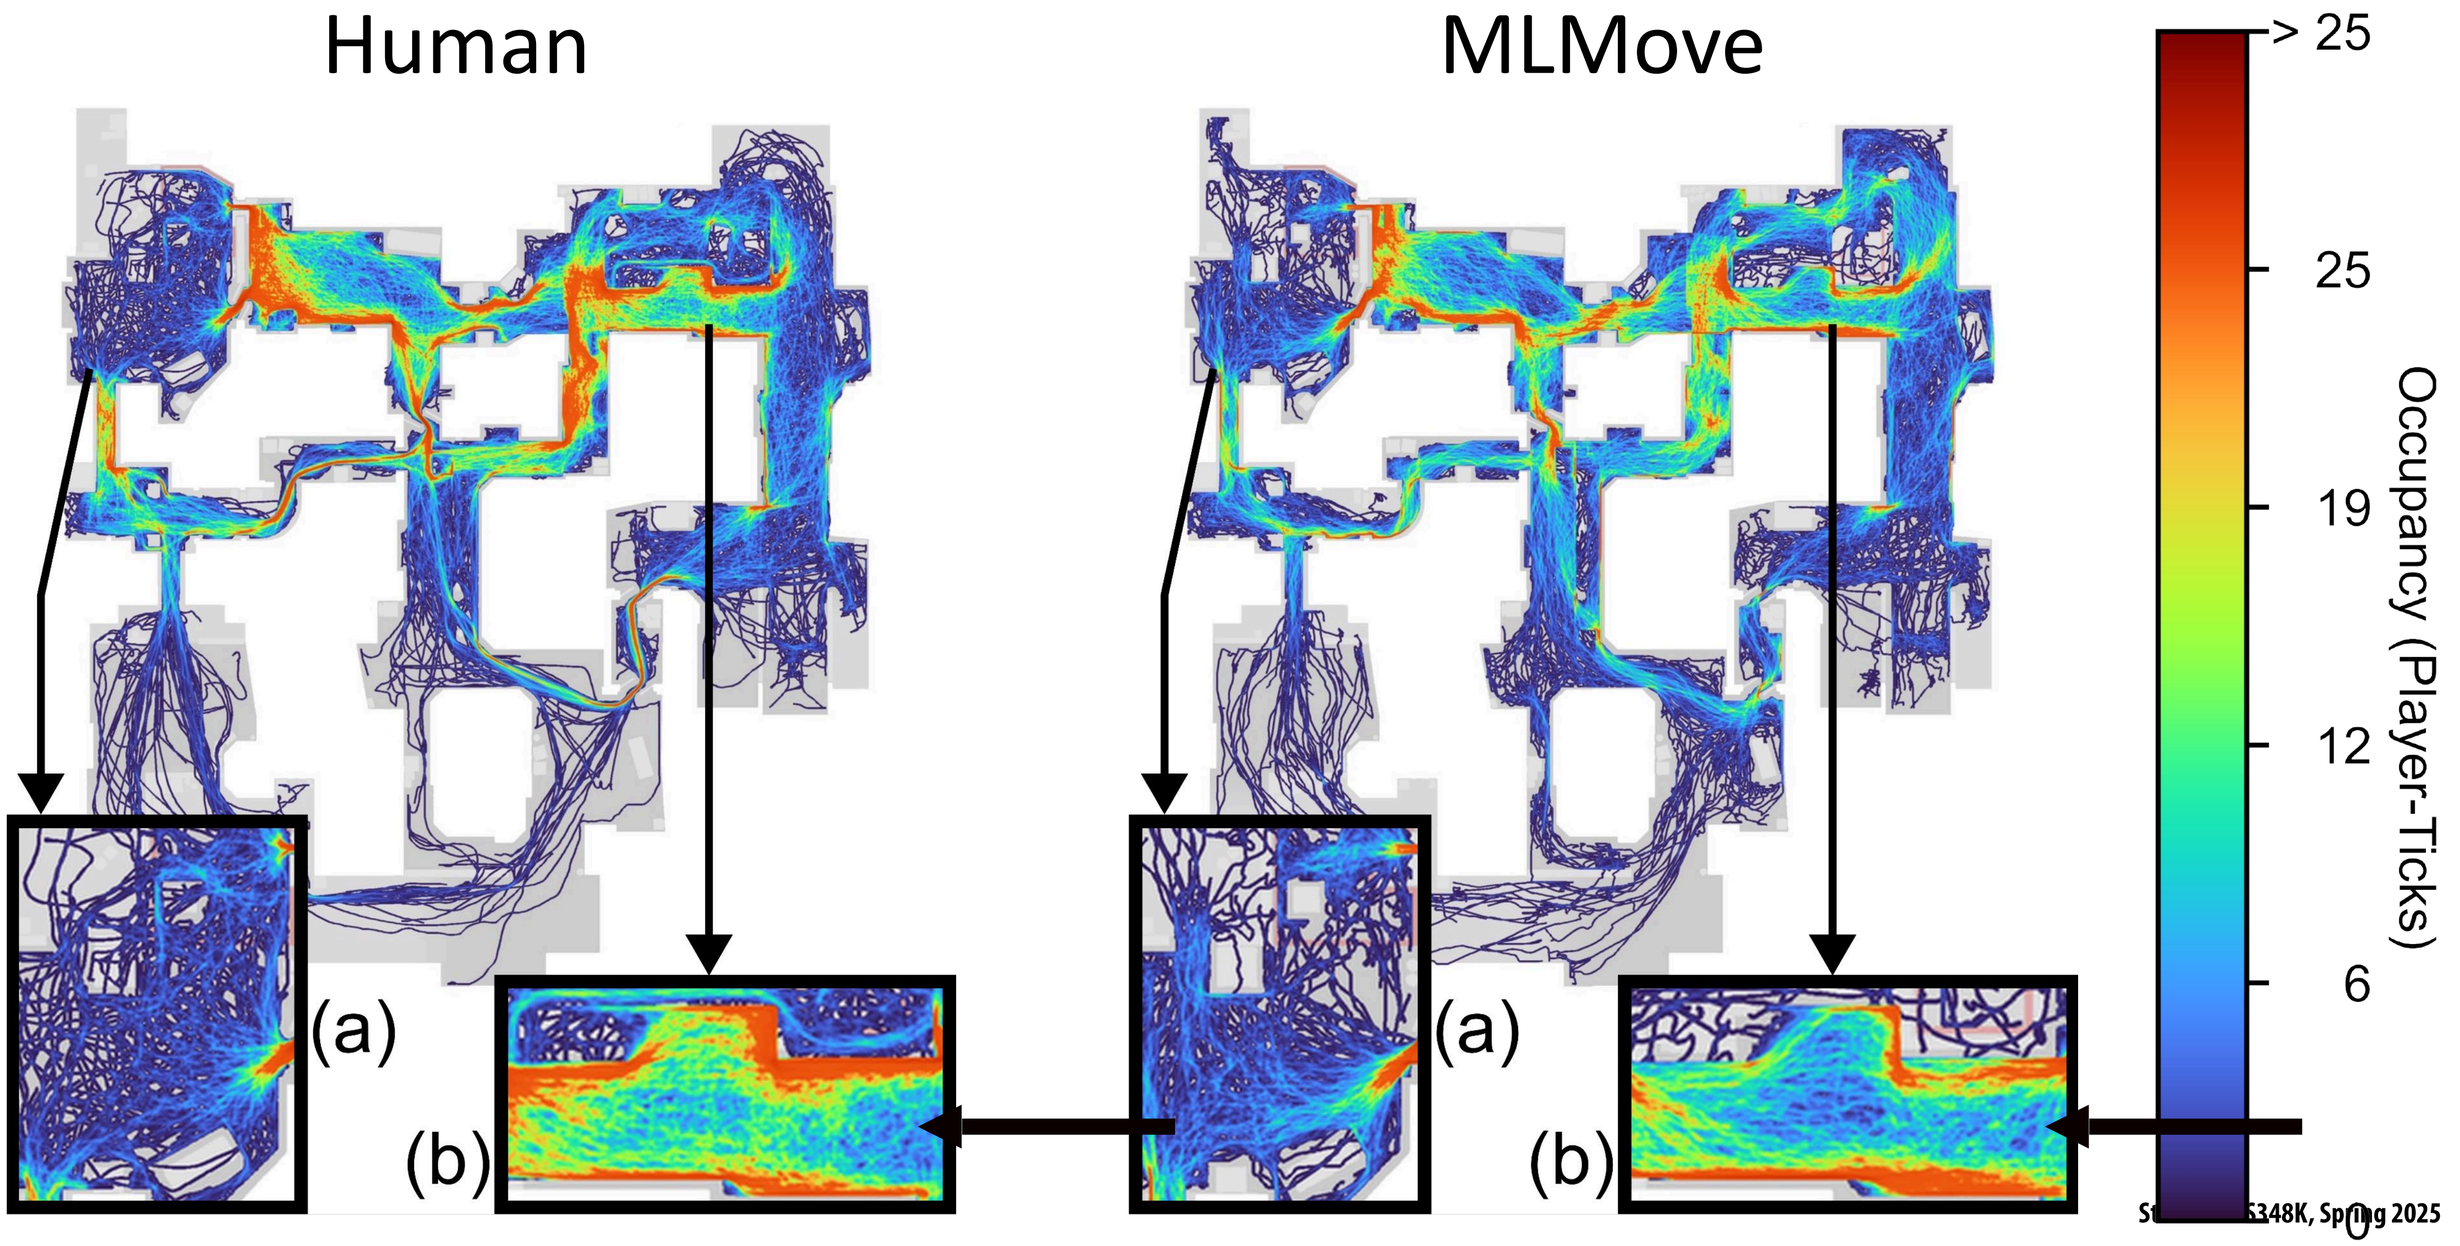

MLMove and Human Positioning Similar

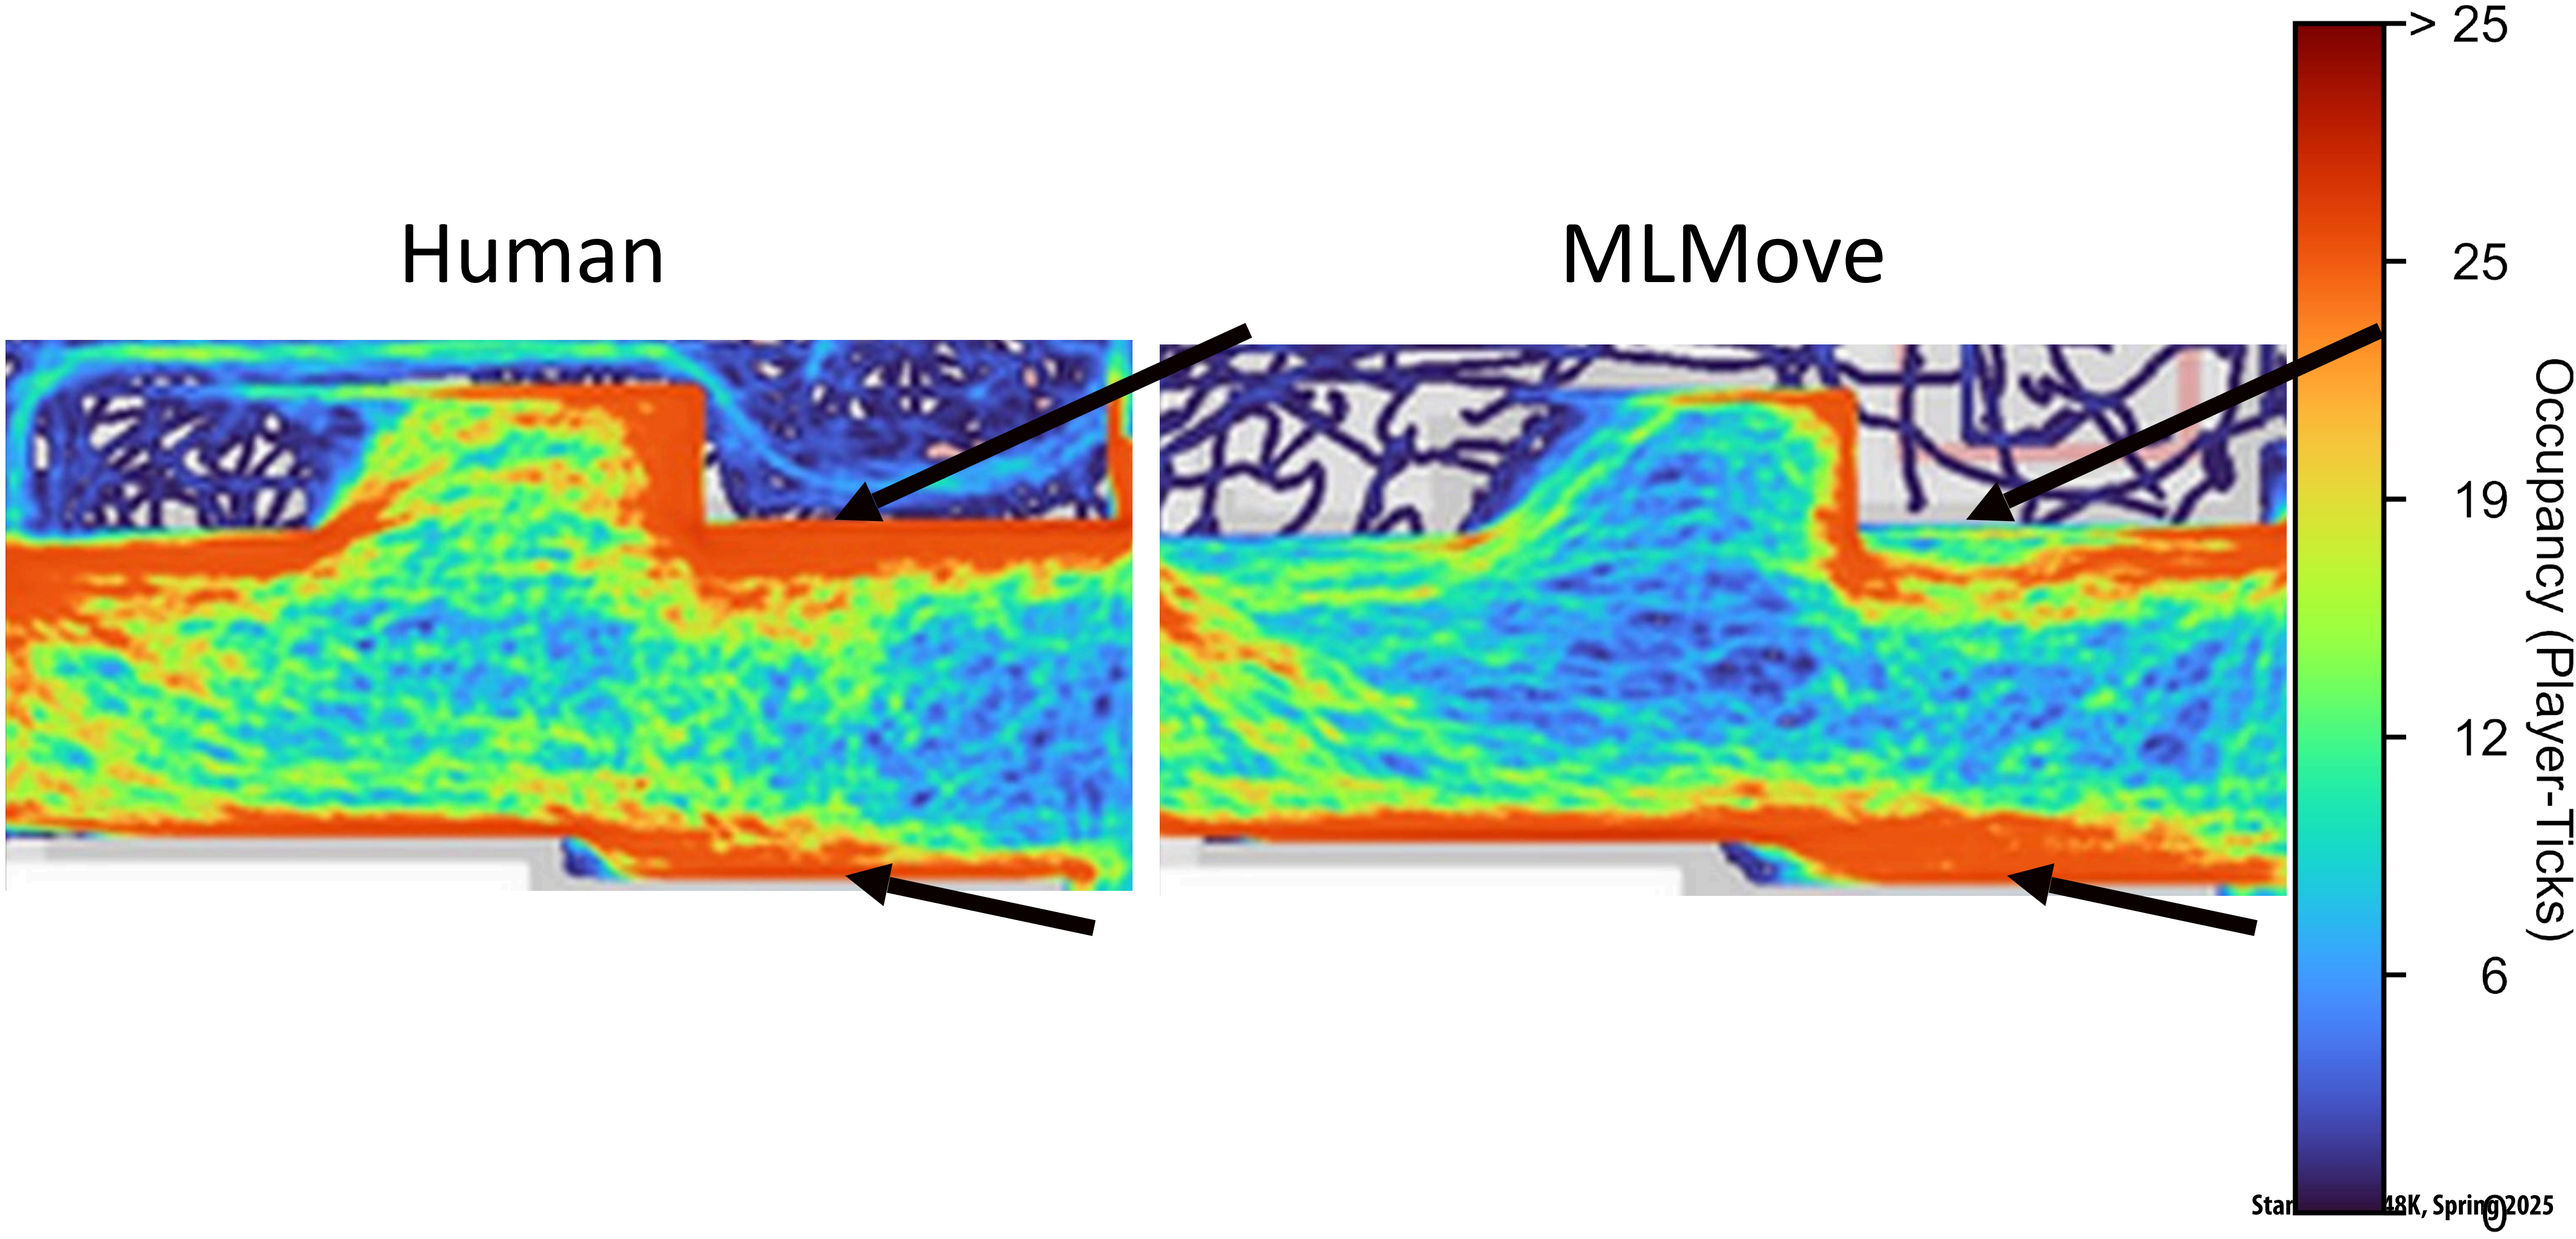

MLMove and Human Positioning Similar

Human

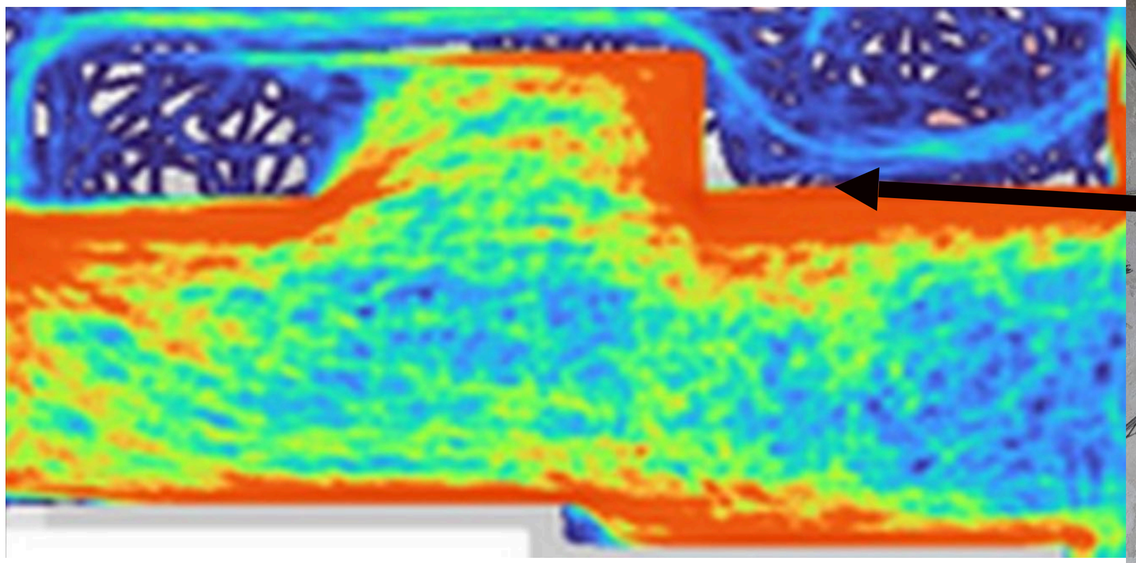

MLMove and Human Positioning Similar

Human

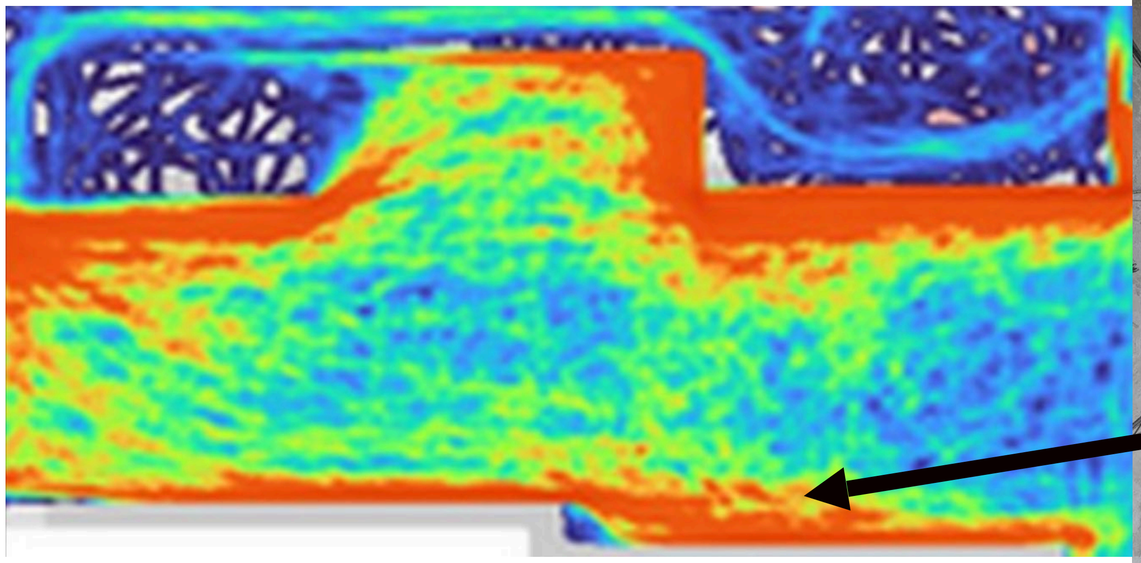

Human players typically score enemy kills from positions of “cover”

Human

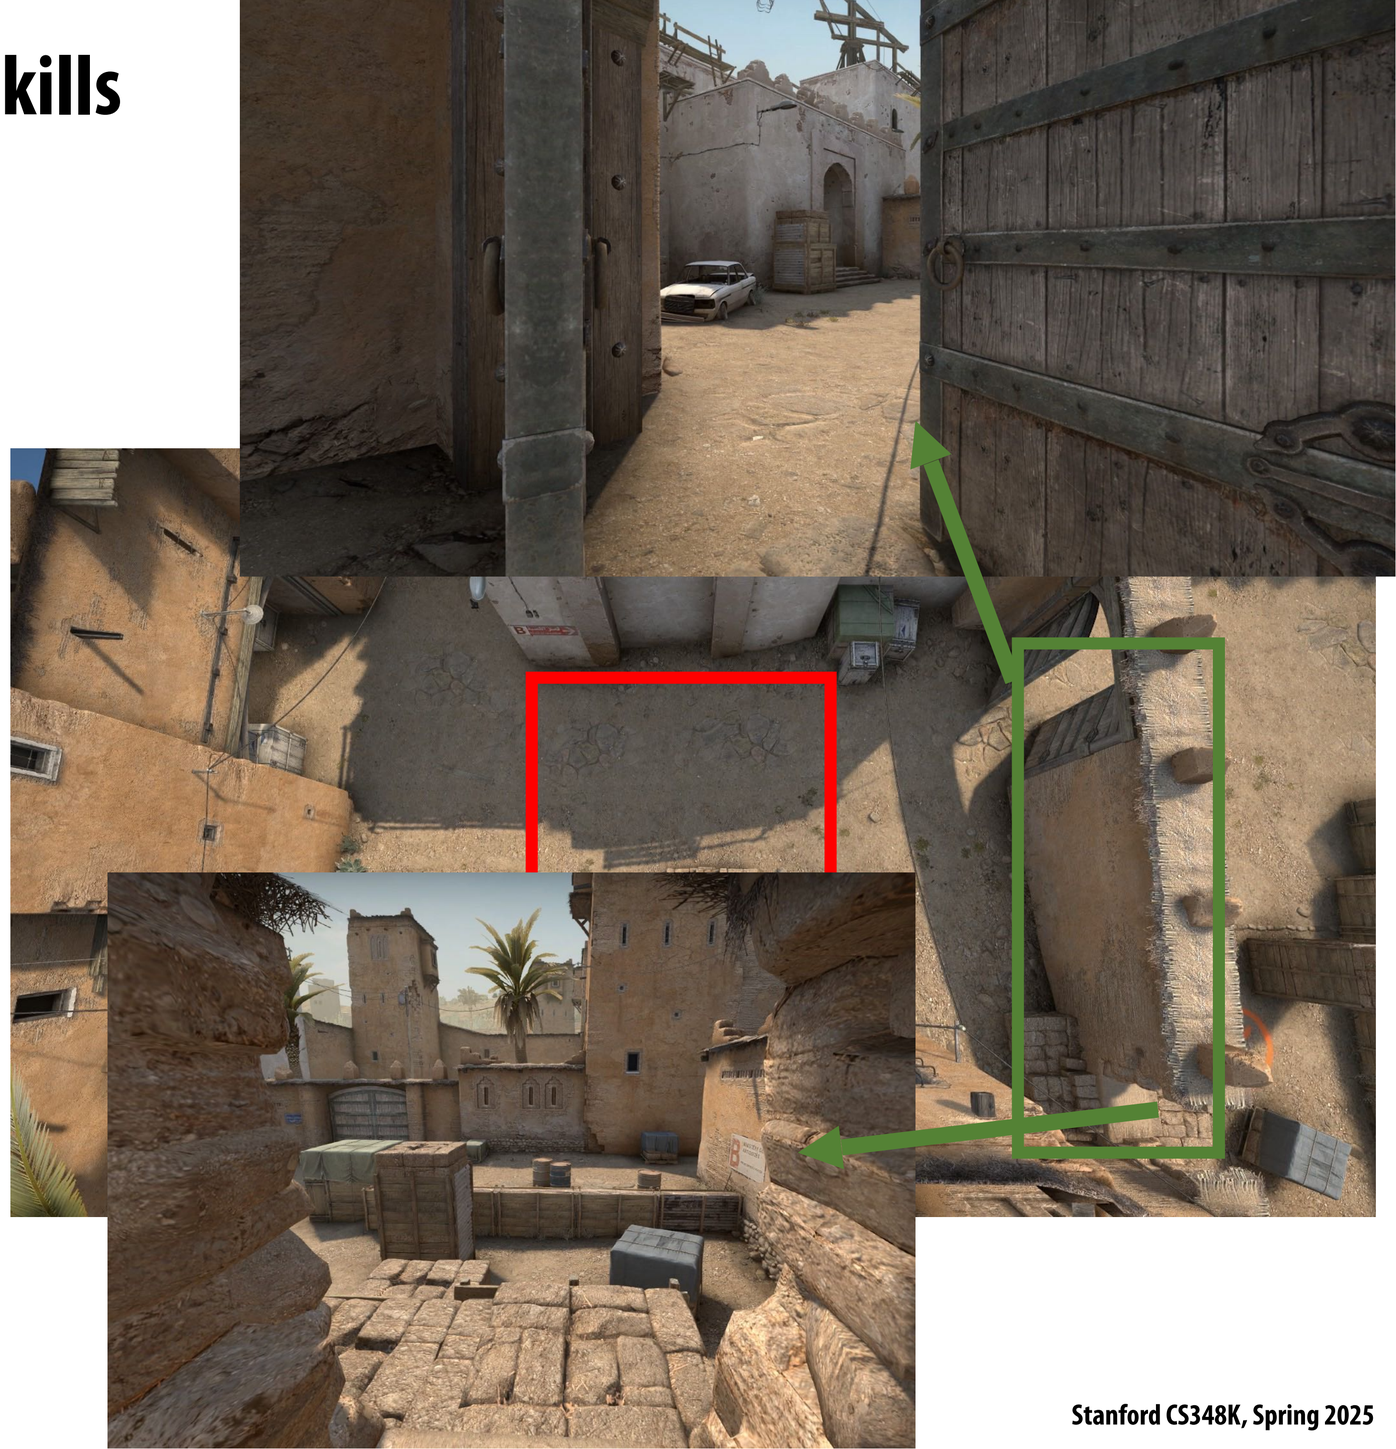

MLMove replicates human kill positions

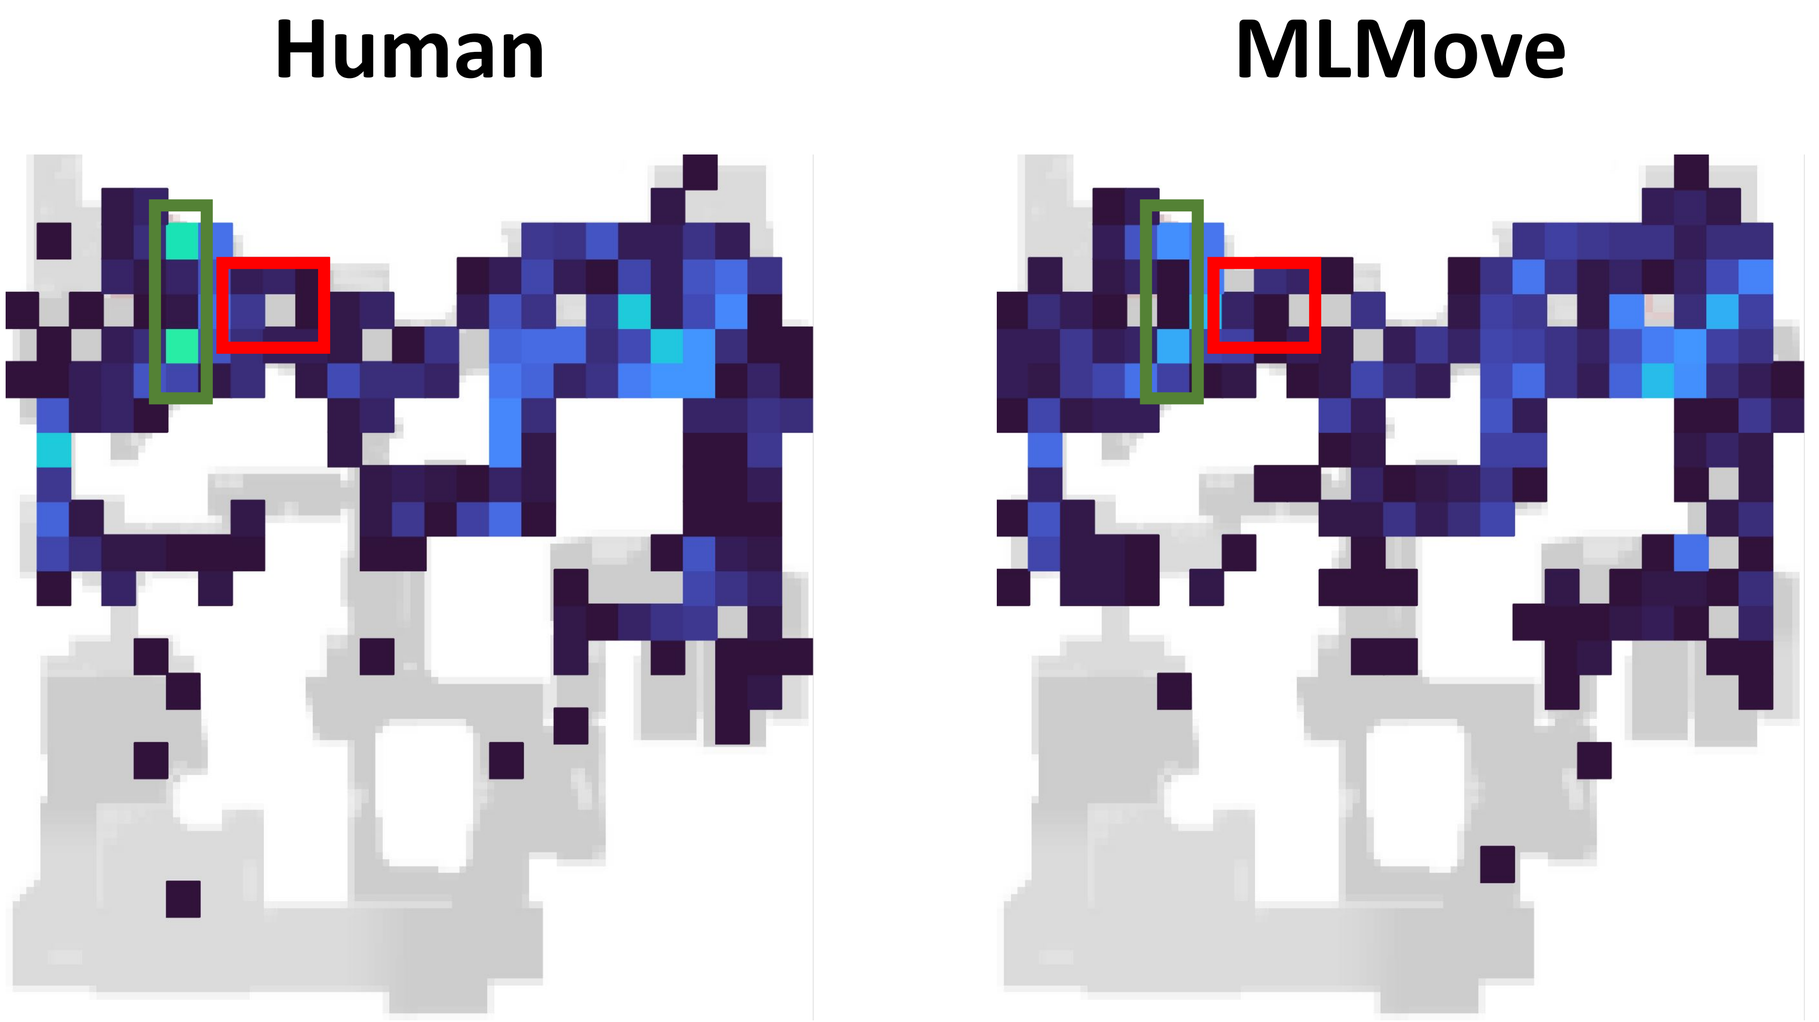

Predicting behavior of humans for many other types of “design”

- Software design — how will users on a social network respond if constantly reminded of certain anti-bullying rules ("Sim Reddit”) [Park et al. 2022]
- Public policy — how would changing the price of an item effect consumer habits
- Political science — how would a population respond to a particular campaign strategy

Modeling human behavior

- Recent study (by the same authors as Generative agents) [Park et al. 2024]
- Interview 1000 people + ask those people a detailed set of questions about their personality, ask them to perform simple tasks, etc.
- Give transcript of interview to LLM agent… see how the result agent does on the same interview

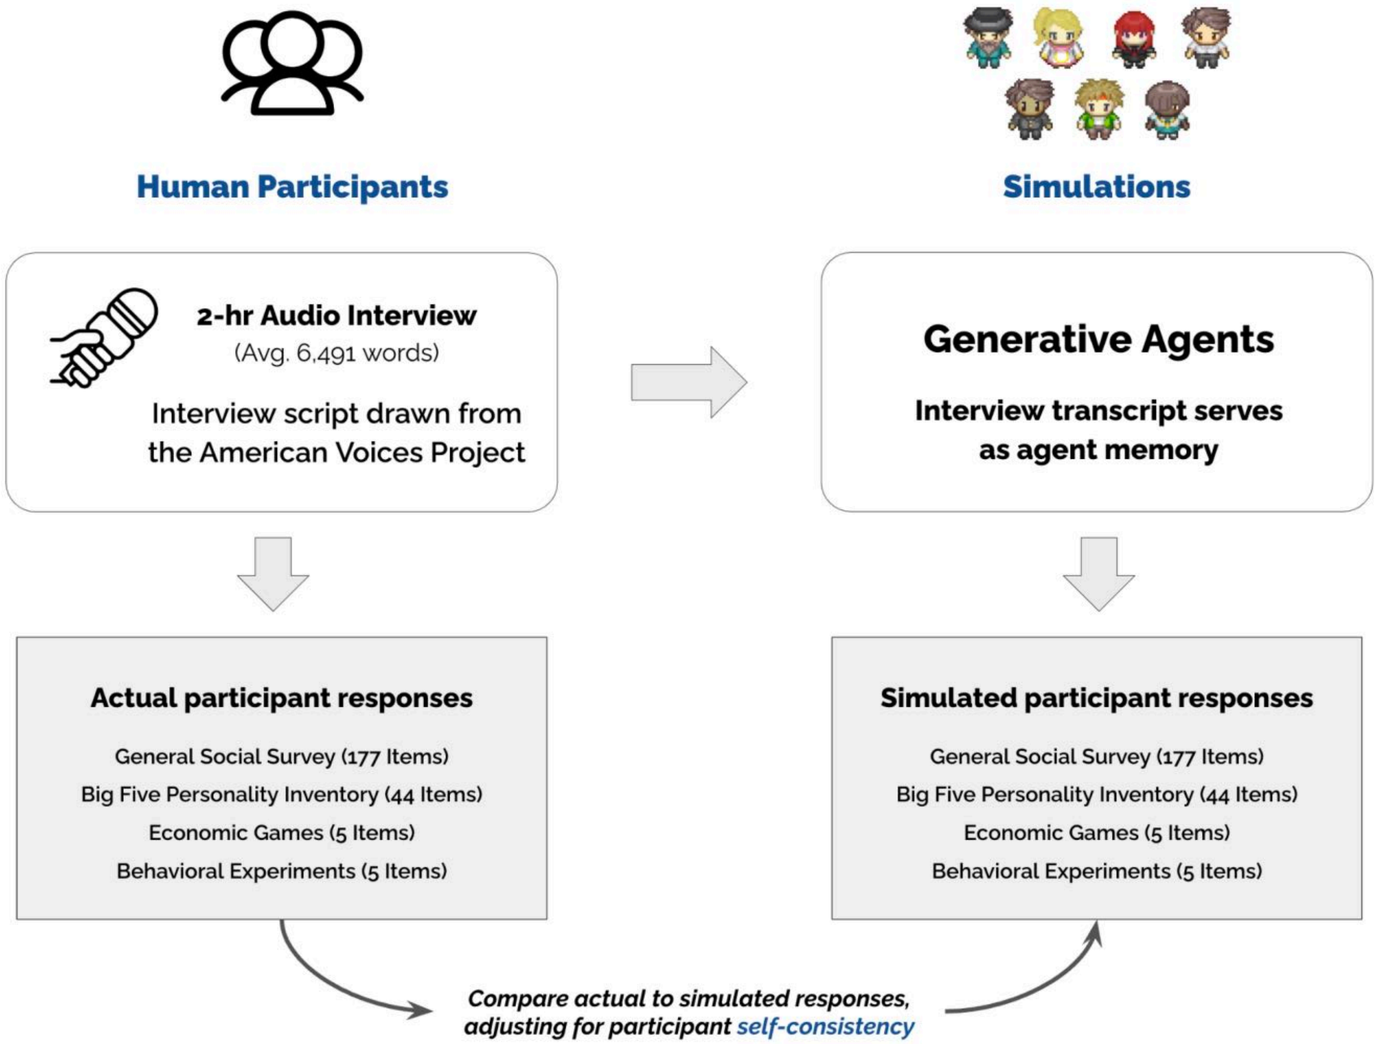

Similarity to human responses

- Generative agent built off of interview context responds to questions more like the interview than alternative ways to model a persona

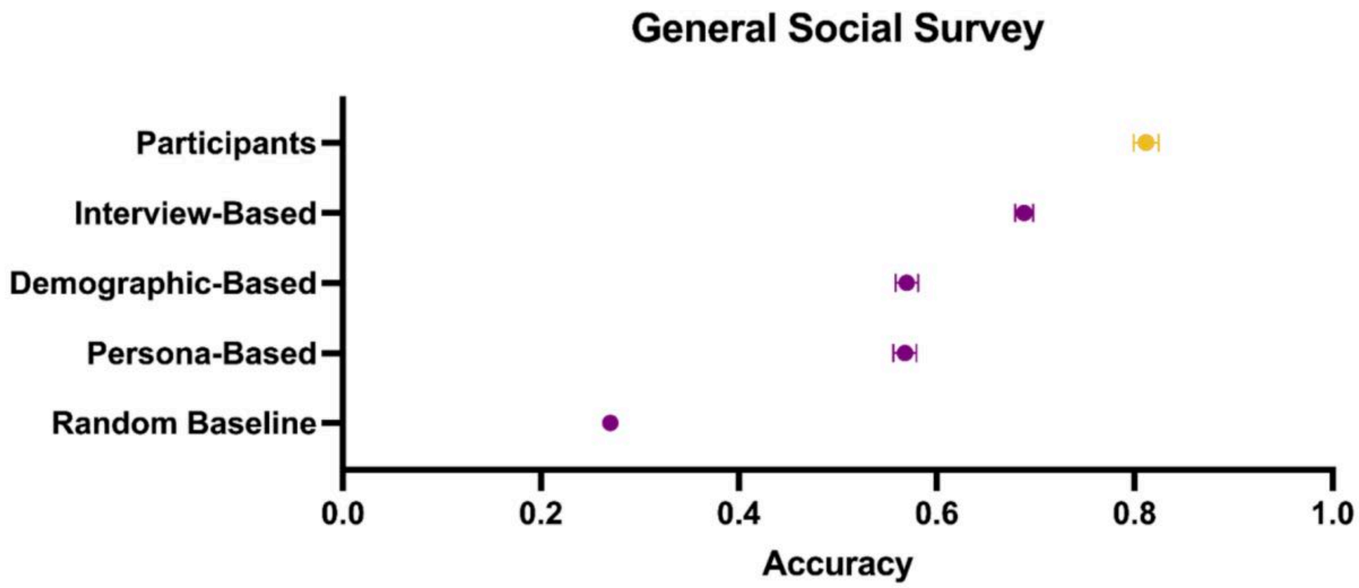

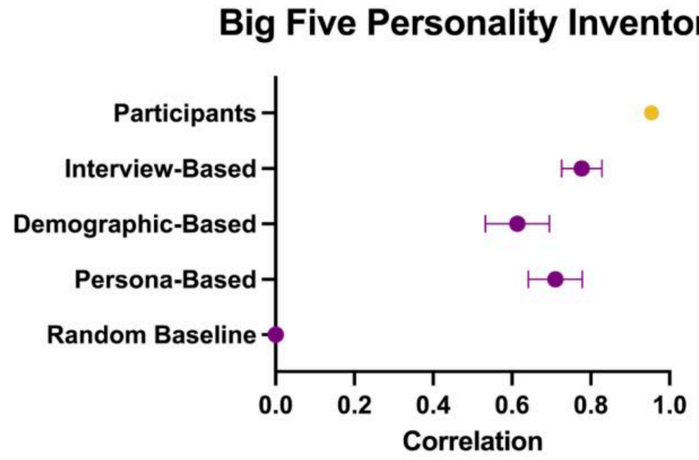

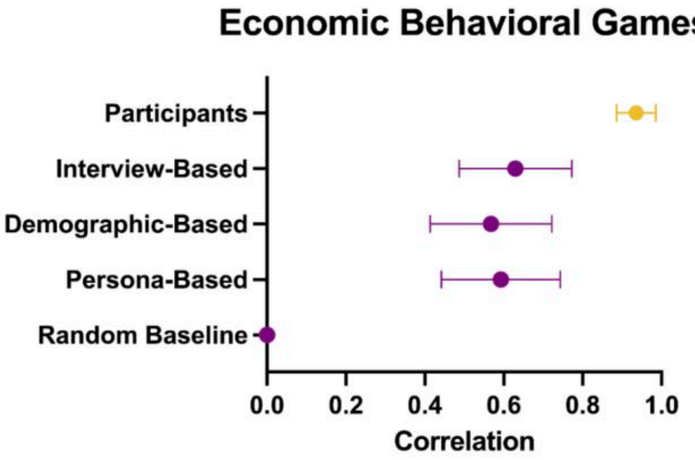

Summary

- When most people think about generative AI for visual computing today, they typically think about generating images, videos, 3D meshes, etc.

- But there's a growing body of work on generating plausible human behavior
  - Note this is different from generating "super-human behavior" (highly competent agents)

- Very interesting questions about "what is humanlike"? Are AI agents "aligned" with what humans would do? How can we use the results of agent similar to make reliable predictions about how a design action will affect or result in certain human behavior in the future?