

Re-classification of Information Seeking Tasks and Their Computational Solutions

Grace Hui Yang (Georgetown University)

@Task-focused IR Workshop, Microsoft Seattle
09/27/2023

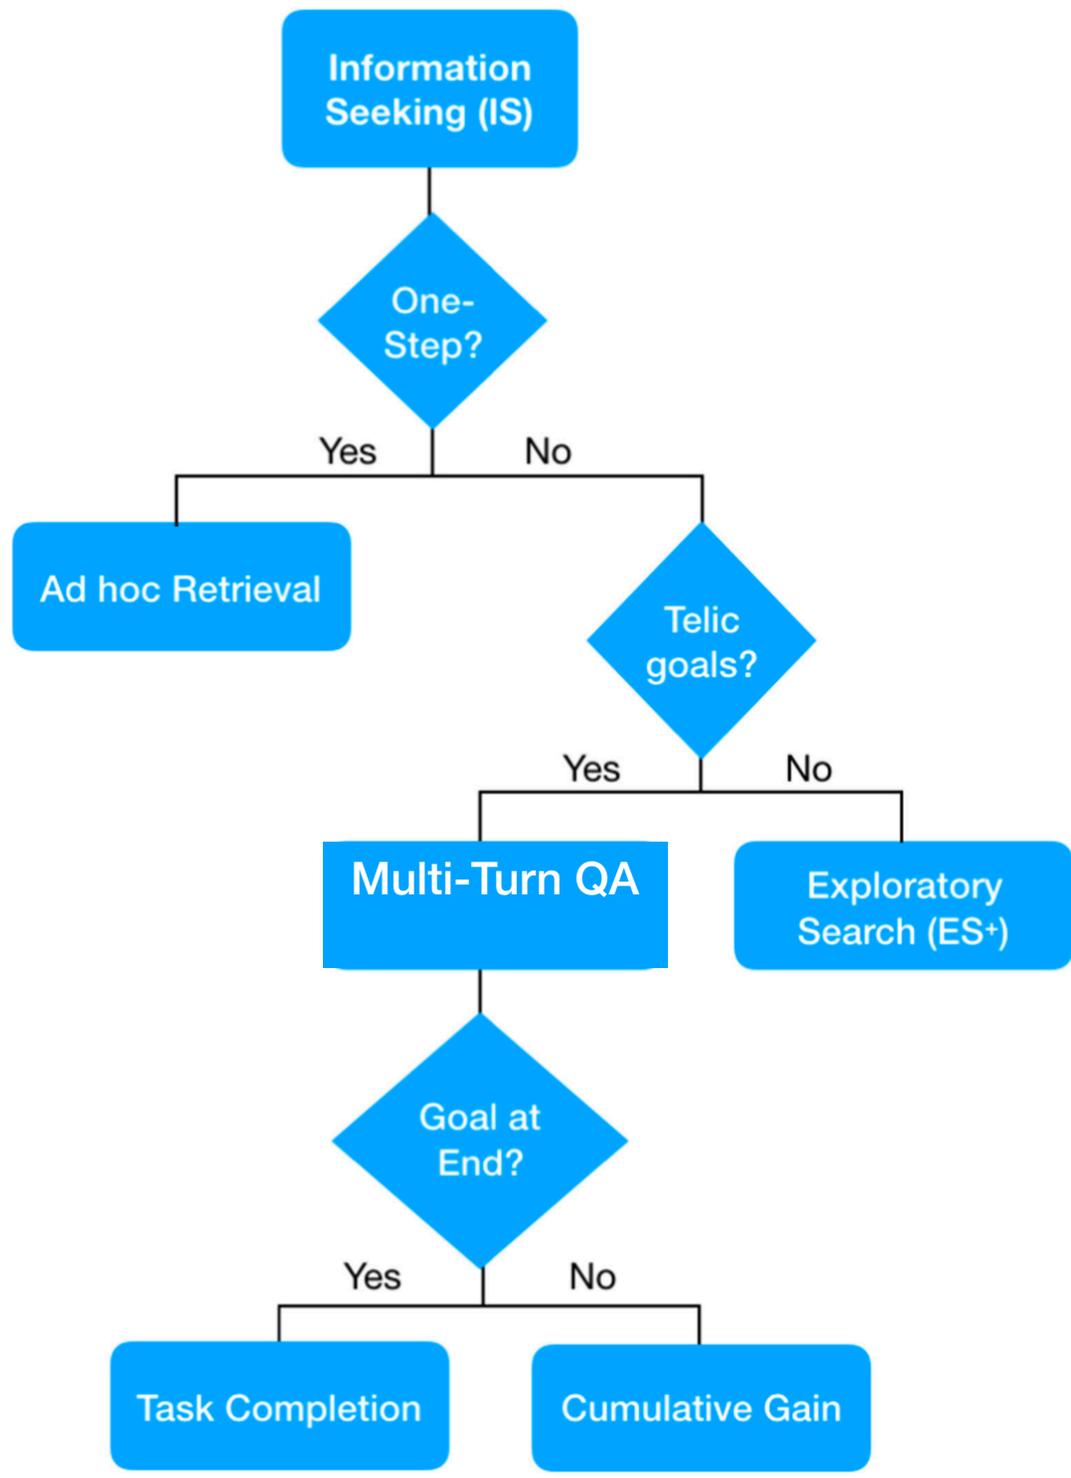
Existing Taxonomies of IS

- Gary Marchionini. 2006. Exploratory search: From finding to understanding. *Commun. ACM* 49, 4 (2006), 41–46.
 - Look-ups, Learning (e.g. literature review), Investigation (e.g. legal search)
- Ryen White and Resa Roth. *Exploratory Search: Beyond the Query-Response Paradigm*. Morgan & Claypool Publishers, 2009.
 - Exploratory search, sense-making, information forage, and berry-picking, information retrieval
- Athukorala, et al., 2016. Is exploratory search different? A comparison of information search behavior for exploratory and lookup tasks. *JAIST*. 67, 11 (2016), 2635–2651.
 - Look-ups vs. exploratory search

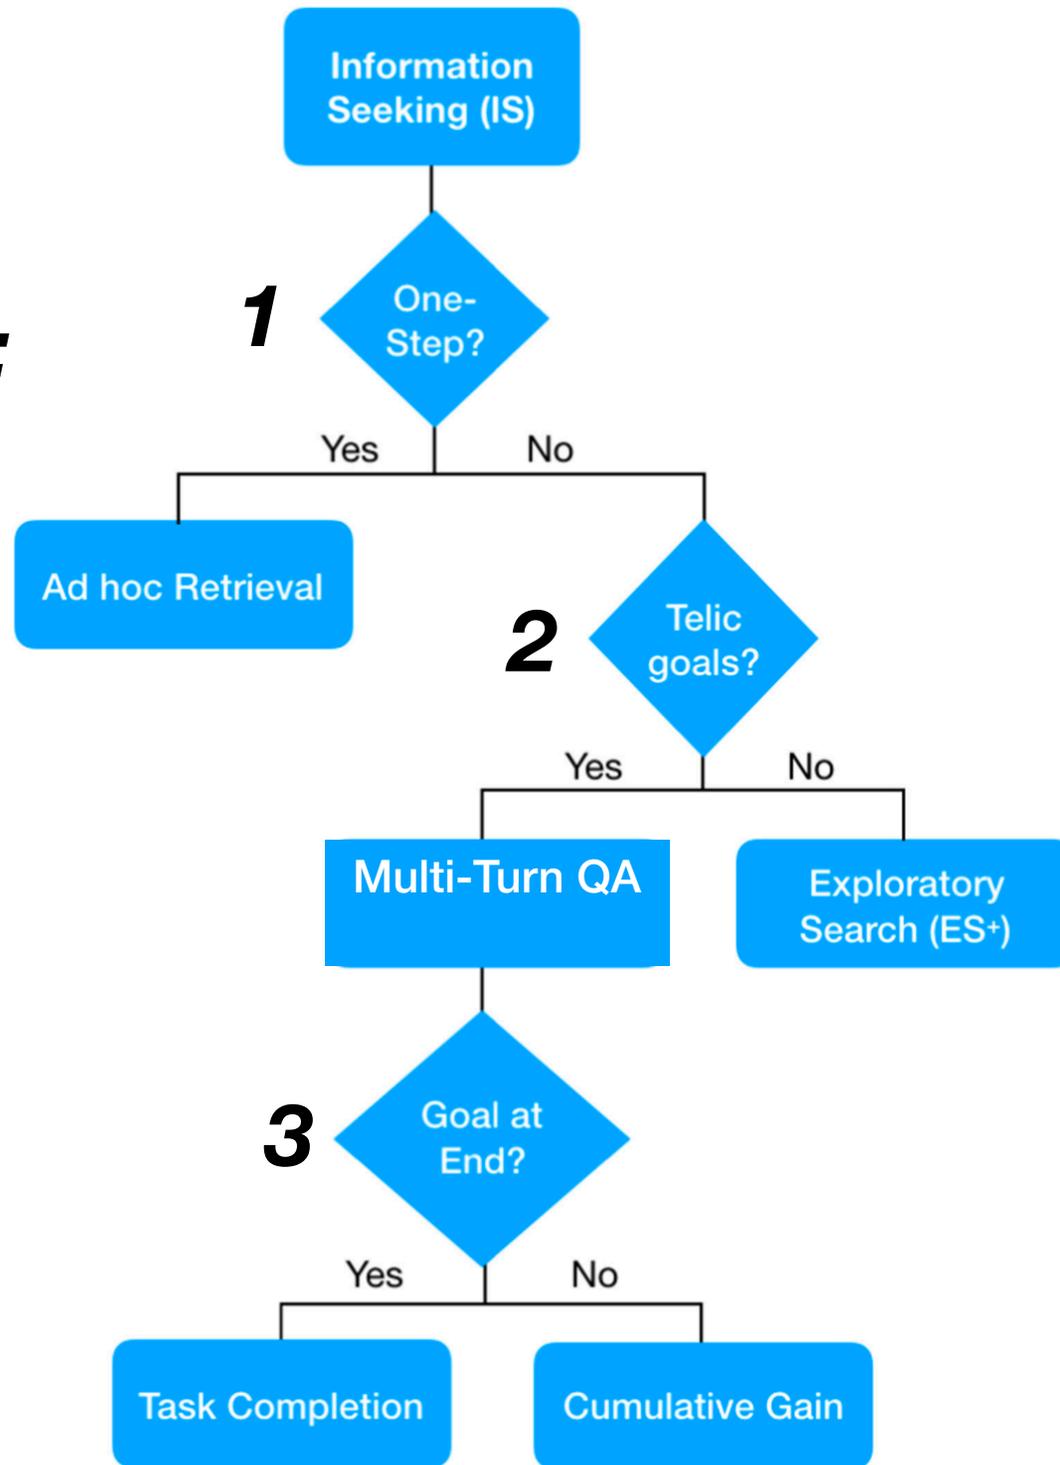
This Talk: A Re-classification of IS Tasks

- Take into account the tasks' innate *computational* nature
- “what it is” vs. “how it goes”
 - Previously, a “noun”-like definition
 - Now, a “verb”-like definition

Zhiwen Tang and Grace Hui Yang. A Re-classification of Information Seeking Tasks and Their Computational Solutions. *TOIS* 40(4): 80:1-80:32 (2022).

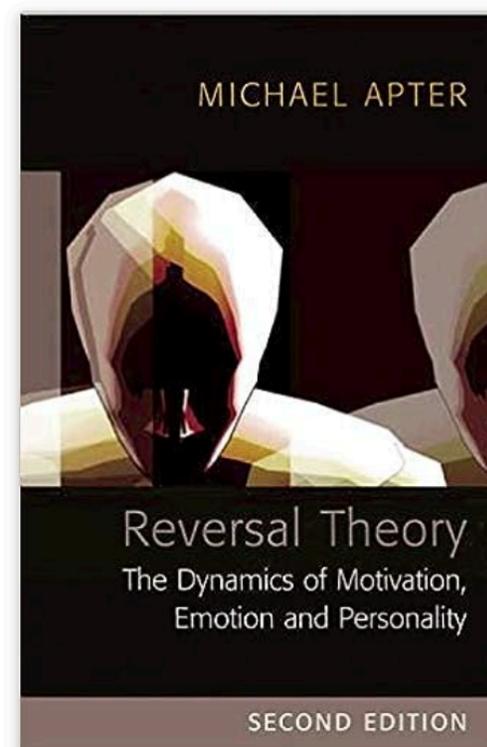


***New
Dimensions:***



2nd Dimension: Search Goal Types

- Based on Reversal Theory, a psychology study that defines four domains of human motivations
 - “means-ends,” “rules,” “transactions,” and “relationships”
- Mainly using the first domain (means-ends), search goals can be
 - *Telic*:
 - User is motivated by achievement, task completion, and fulfilling of goals
 - Serious. Focus on future goals and achievement. Tend to avoid arousal, risk & anxiety.
 - E.g., when people have telic goals, they run because they want to win a medal
 - *Paratelic*:
 - User is playful and seeks excitement and fun
 - Playful, passionate and fun. Focus on current moment. Seek excitement and entertainment.
 - E.g., when people have paratelic goals, they run because they enjoy running itself



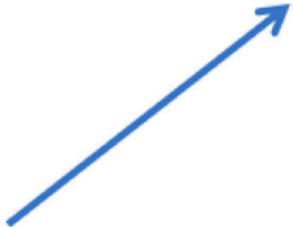
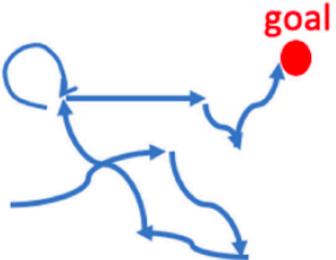
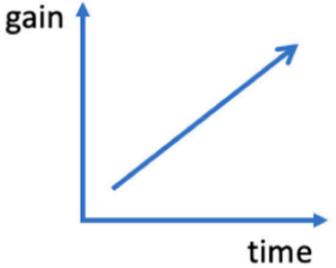
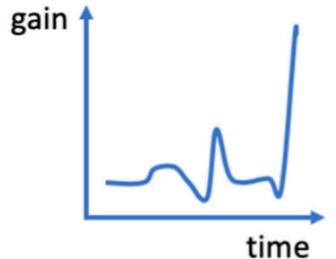
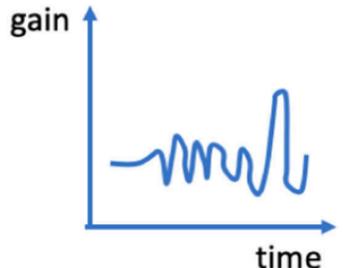
Motivational States



Michael J. Apter. 2005. Personality Dynamics: Key Concepts in Reversal Theory. Apter International.

3rd Dimension: Is Your Goal Converging?

- Converging goals:
 - Close a deal, make a purchase, follow a recommendation, book a ticket, win a game, etc ...
- Non-converging goals:
 - Be happy, get as much as rewards/relevant documents, watch more good movies, get this answer correctly answered and expect more to be answered correctly, etc ...

	IS w/ Telic Goals		IS w/ Paratelic Goals
	Being cumulative	Being successful	Being enjoyable
Example Research Problems	Dynamic Search Recommender Systems eDiscovery Online L2R	Dynamic Search Taskbots Task-based Dialogue Systems Conversational AI/Search Multi-turn QA	Exploratory Search Social bots
Task Illustration			
Gain Curves			
Suitable ML Methods	Bandits Algorithms Local Methods	Policy-Based RL Actor-Critic Model-Based RL	Value-Based RL Model-Based RL Open-Ended RL Local Methods
	Imitation Learning		

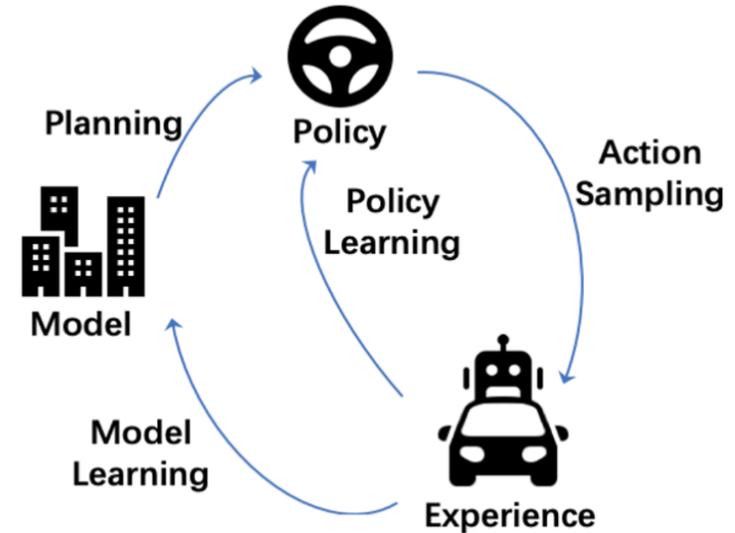
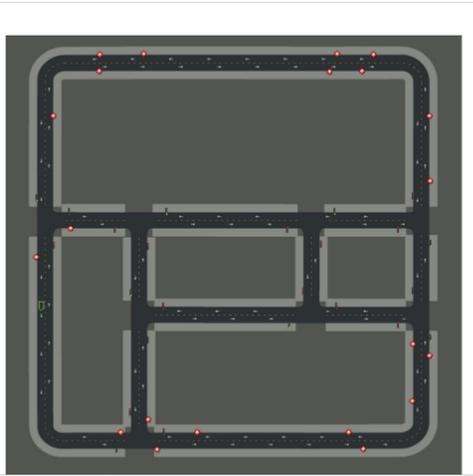
Local Methods

- They are stateless, commonly used IR methods
- Relevance feedback methods
 - E.g., Rocchio
- Query reformulation, query expansion, query-reweighing, query suggestion ...
- Active learning methods

Bandits Algorithms

- Stateless algorithms or with a limited local context
- Choose the best actions to max long-time return/rewards
- With a flavor of gambling
- Widely used in recommender systems

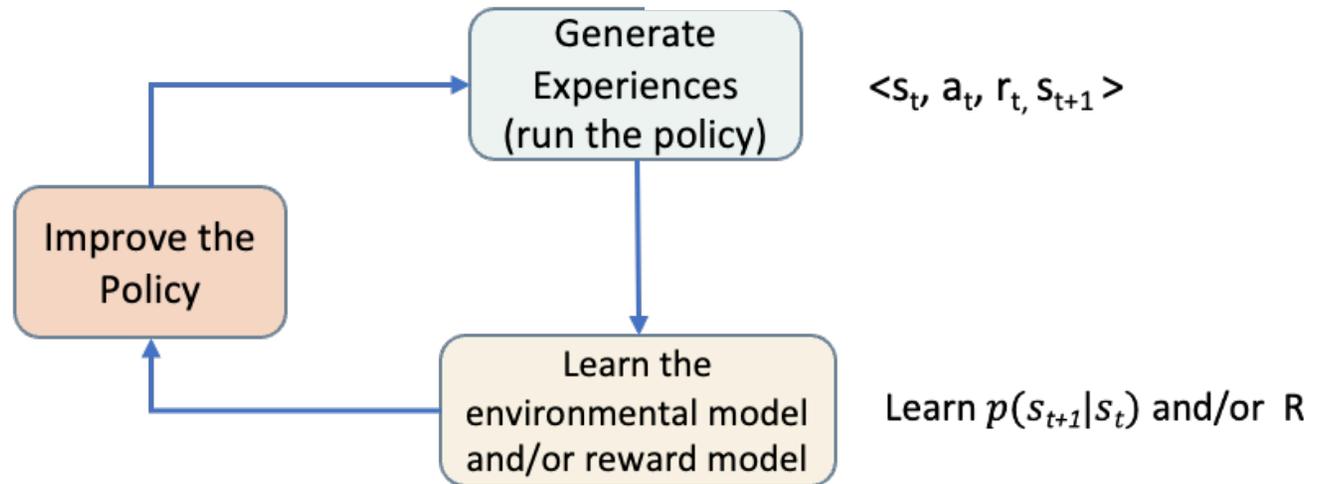
Model-Based RL



Planning or any model-free policy updates

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} E[R_{\pi_{\theta}}]$$

$$\pi(s) \leftarrow \underset{a}{\operatorname{argmax}} Q_{\mu}(s, a)$$



Open-Ended RL

- Task domain (environment) is unbounded
- Domain randomization & Augmentation methods
 - Limin Chen, Zhiwen Tang and Grace Hui Yang. Balancing Reinforcement Learning Training Experiences in Interactive Information Retrieval. SIGIR'20
- Evolutionary methods

A Hypercube-Based Indirect Encoding for Evolving Large-Scale Neural Networks

Accepted to appear in *Artificial Life journal* 15(2), Cambridge, MA: MIT Press, 2009

Kenneth O. Stanley (kstanley@cs.ucf.edu)

David D'Ambrosio (ddambro@cs.ucf.edu)

Jason Gauci (jgauci@cs.ucf.edu)

School of Electrical Engineering and Computer Science

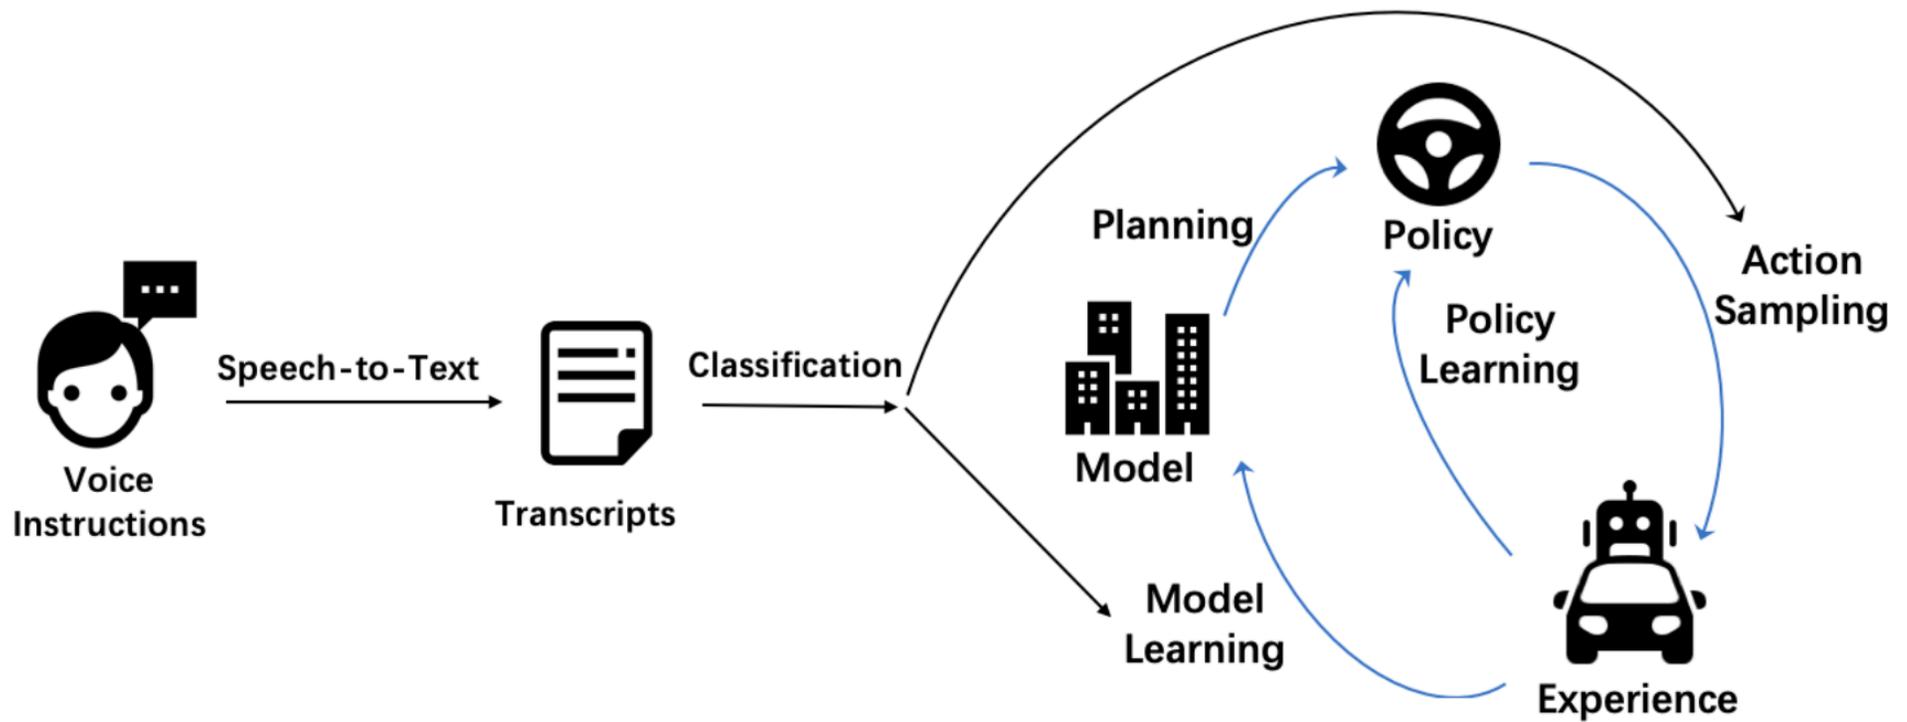
University of Central Florida

4000 Central Florida Blvd.

Orlando, FL 32816-2362 USA

Keywords: Compositional Pattern Producing Networks, CPPNs, HyperNEAT,
large-scale artificial neural networks, indirect encoding, generative and developmental systems

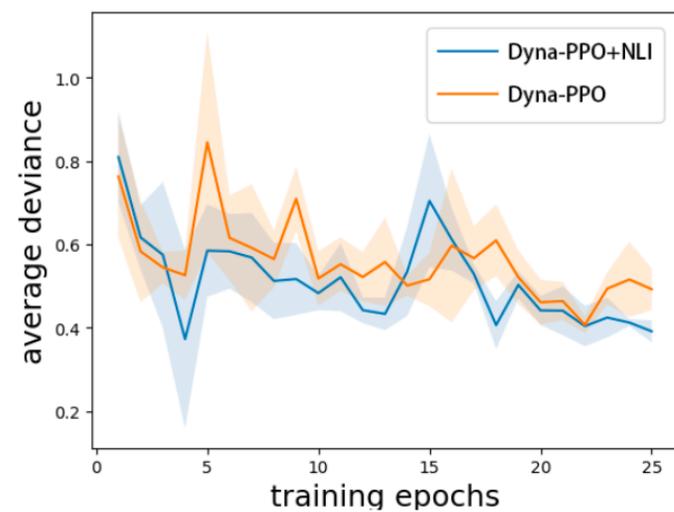
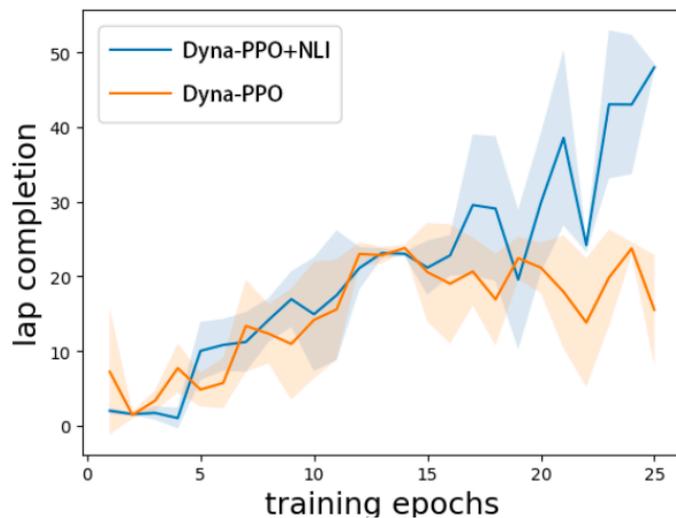
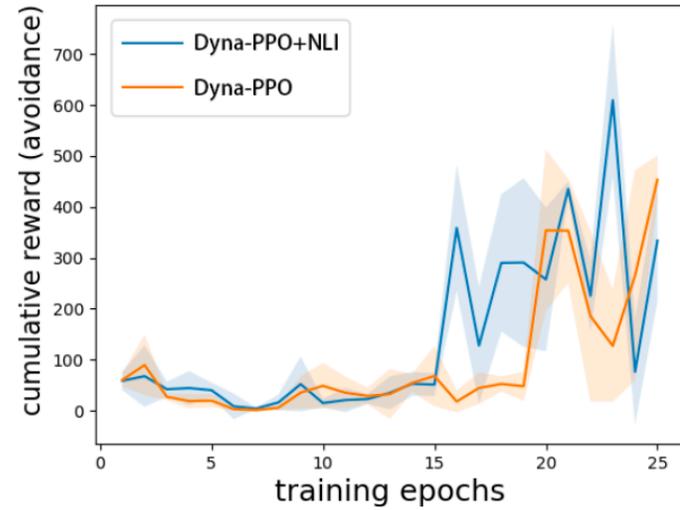
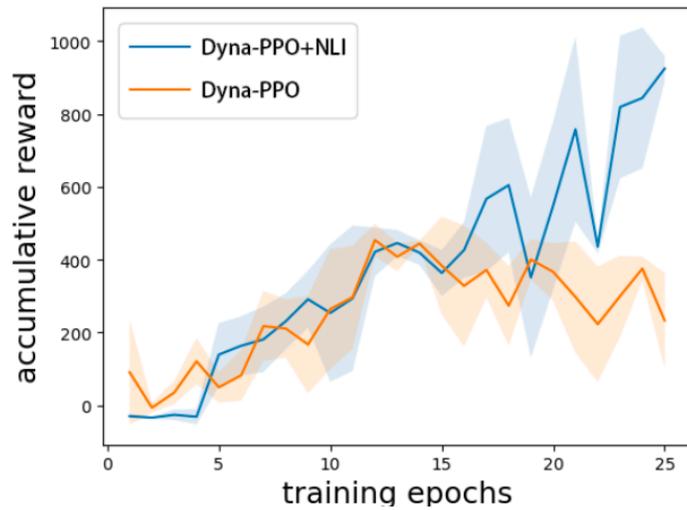
Human-in-the-loop Learning (HILL) to Coach RL Agents



Mingze Wang, Ziyang Zhang and Grace Hui Yang. Incorporating Voice Instructions in Model- Based Reinforcement Learning for Self-Driving Cars. *NeurIPS'21 Workshop on Machine Learning for Autonomous Driving*,

Incorporating Voice Instructions in Model-Based RL

Instructions in Model-Based RL



Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality

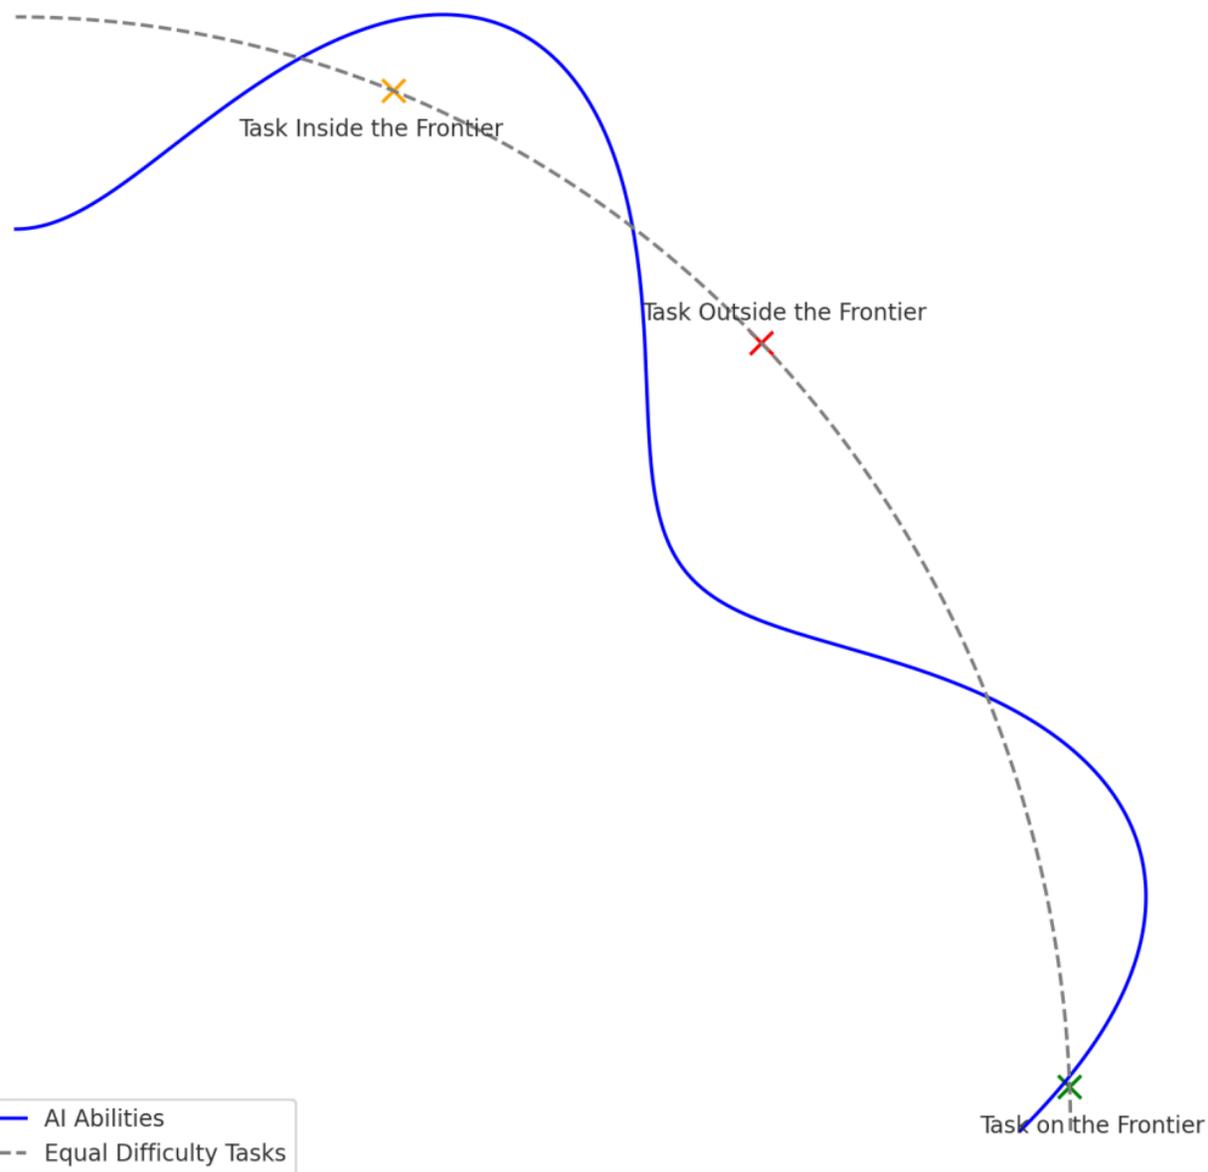
Fabrizio Dell'Acqua
Edward McFowland III
Ethan Mollick
Hila Lifshitz-Assaf
Katherine C. Kellogg

Saran Rajendran
Lisa Kraymer
François Cadelon
Karim R. Lakhani



Harvard
Business
School

Jagged Frontier of AI Capabilities



Notes: This figure displays the AI frontier as jagged. Tasks with the same perceived difficulty may be on one side or the other of the frontier. ChatGPT produced this image starting from the authors' prompts.

Conclusion

- A re-classification of IS tasks based on its computational (procedural) structure
 - New dimensions: # search iterations, search goal/ motivation types, procedures to reach these goals
- Suitable algorithmic solutions for each new class
- Ways to incorporate more human interferences into RL-Based Agents' learning process

Call for Sponsorship

SIGIR 2024@Washington, D.C.



- Fully in-person conference
 - In-person presentation + early registration/visa letter + proxy presenters + recording tech talks during conference + post-conference video releasing on Youtube or ACM DL
- Five Days of Events
 - Full/short papers, perspectives papers, resource & reproducibility papers, demos, doctoral consortium, tutorials, workshops
 - Award ceremony, industry track (SIRIP), government day, a special topic (LLMs) panel, and various DEI and social events