No AGI Without Inference-time Search

Simon Ouellette

<u>Main Claim:</u> human-like AI (AGI) is not possible with a learning-only algorithm.
An inference-time search component is necessary.

THEORY

OUTLINE OF ARGUMENT

- Open-world problem domains cannot be completely learned.
- Epistemic uncertainty: the learned model assigns the highest probability to an incorrect prediction.
- When a verifier is available (e.g. planning, abductive reasoning), an intelligent system uses search to resolve the uncertainty.
- ML systems have no intrinsic mechanism for trial-and-error (search)

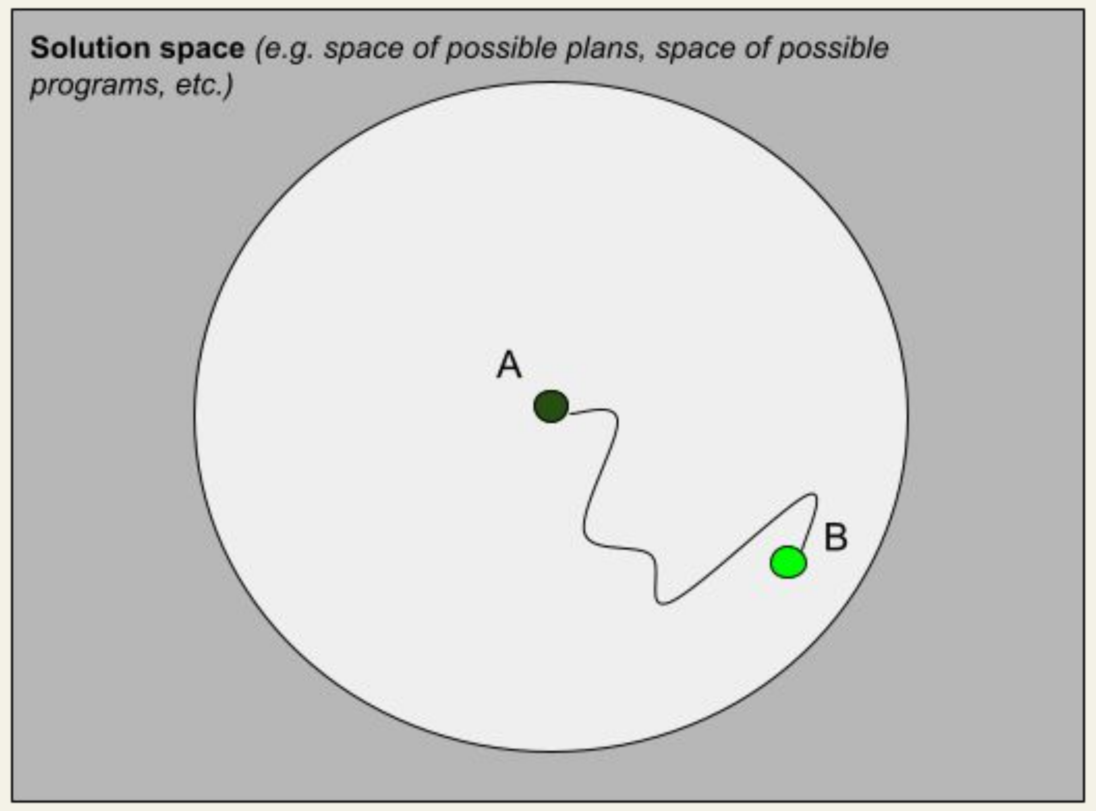

DUAL PROCESSING SYSTEM

- The deep learning module (System 1) emits an approximate (but incorrect) solution -- point A.
- Inference-time search (System 2) uses the verifier to search the space of nearby solution and land on exact solution -- point B.

EVIDENCE

CASE STUDY #1: Lake & Baroni (2023) [1]

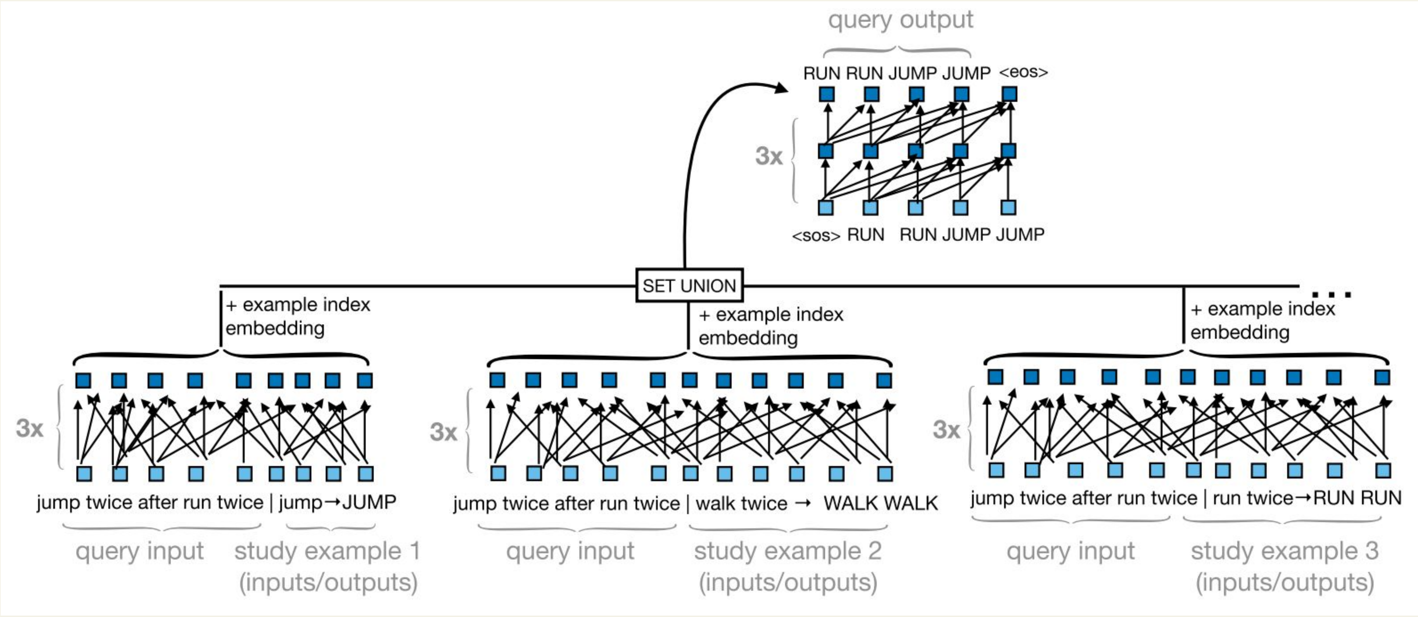

- Transformers do generalize well when the task distribution consists of surface-level re-mappings between the English words and the instruction atoms.
- **Transformers failed to generalize to structurally novel tasks**, such as longer output sequences on SCAN and more complex sentence structures in COGS, with error rates at 100%.
- **Reason: requires productivity** -- the ability to compositionally generate an unlimited set of possible solutions from axiomatic concepts and rules.
- But pure ML has no inherent mechanism for such trial-and-error (search).

CASE STUDY #2: Ouellette et al. (2024) [2]

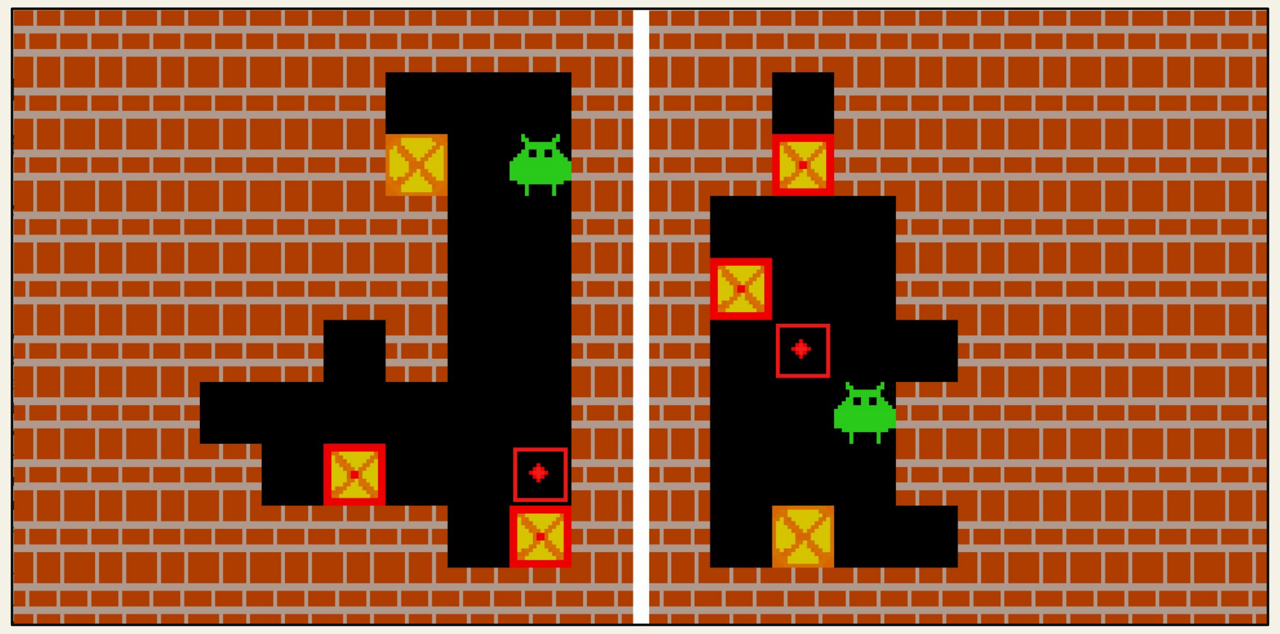

We found that a model-based reinforcement learning system equipped with **a neurally-guided search component can be 50 to 100 times more sample-efficient** on Sokoban than the state-of-the-art Transformer-based solution. Also outperformed all model-free reinforcement learning techniques in sample efficiency.

Relation to current argument: studying performance in a small data regime is analogical to studying performance in open-world problems. Only a relatively small percentage of the total possible space of states is ever seen during training.

CASE STUDY #3: ARC-AGI (Chollet, 2019) [3]

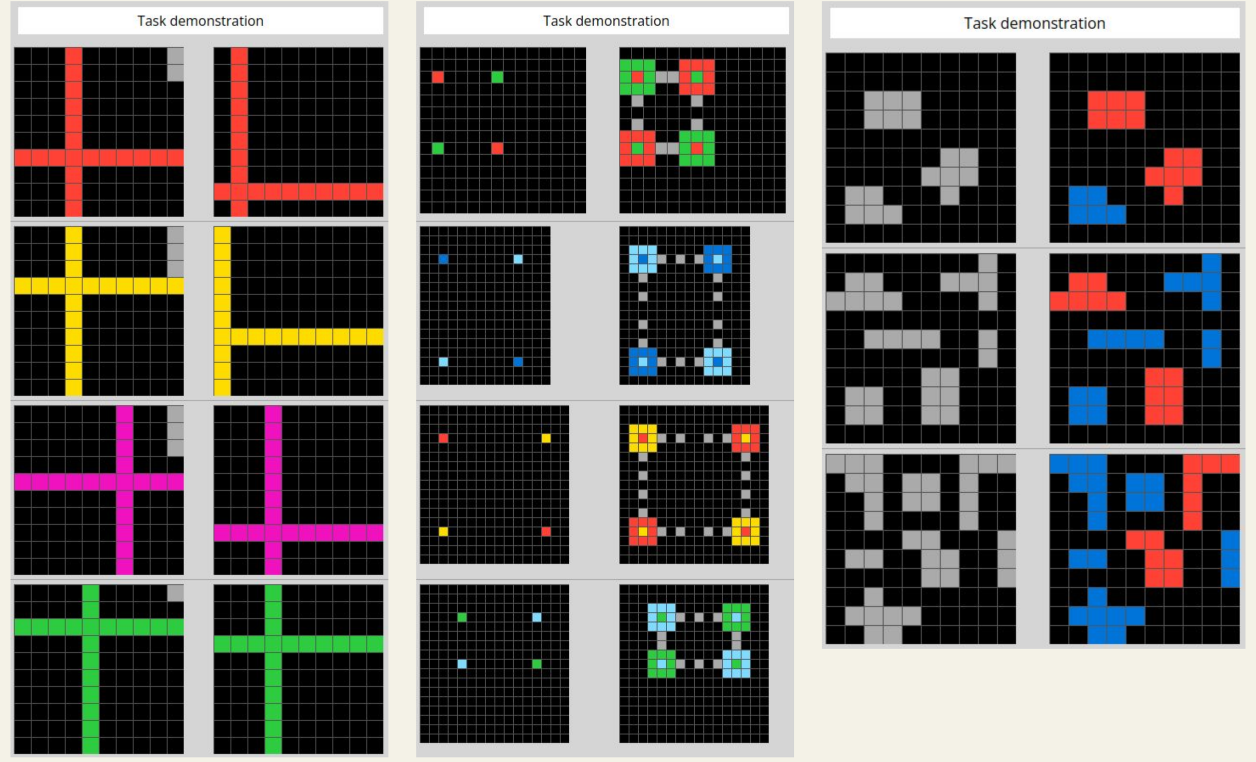

ARC-AGI: open-world visual reasoning dataset. Hidden test set contains qualitatively, structurally distinct tasks from any publicly available examples. Transformers perform poorly on these: they don't generalize well to the hidden test set.

Humans	MindsAI	Greenblatt	o1-preview	Claude 3.5	o1-mini
84%	47%	43%	18%	14%	9.5%

- MindsAI: **Gradient-descent search over weights at inference time.**
- Greenblatt: uses an LLM to generate and progressively refine thousands of Python programs to solve a given task = **Test-time discrete search over programs (program synthesis)**

References:
[1] Lake, B. M., & Baroni, M. (2023). Human-like systematic generalization through a meta-learning neural network. Nature, 623(7985), 115-121.
[2] Ouellette, S., Beaudry, E. et Bouguessa, M. (2024). Conviction-based planning for sparse reward reinforcement learning problems. ICAPS Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning.
[3] URL: arcprize.org