# Decomposed Inductive Procedure Learning:
## Learning Academic Tasks with Human-Like Data Efficiency

**Daniel Weitekamp[1], Christopher J. MacLellan[1], Erik Harpstead[2], Kenneth Koedinger[2]**
[1]Georgia Institute of Technology, [2]Carnegie Mellon University
weitekamp@gatech.edu

**Decomposing learning** into multiple distinct mechanisms **significantly improves data efficiency**, bringing it in line with **human learning**

## Decomposing Learning Mechanisms

- Ablation from 1-mechanism (**reinforcement learning**) to 3-mechanism (**Decomposed Inductive Procedure Learning**)

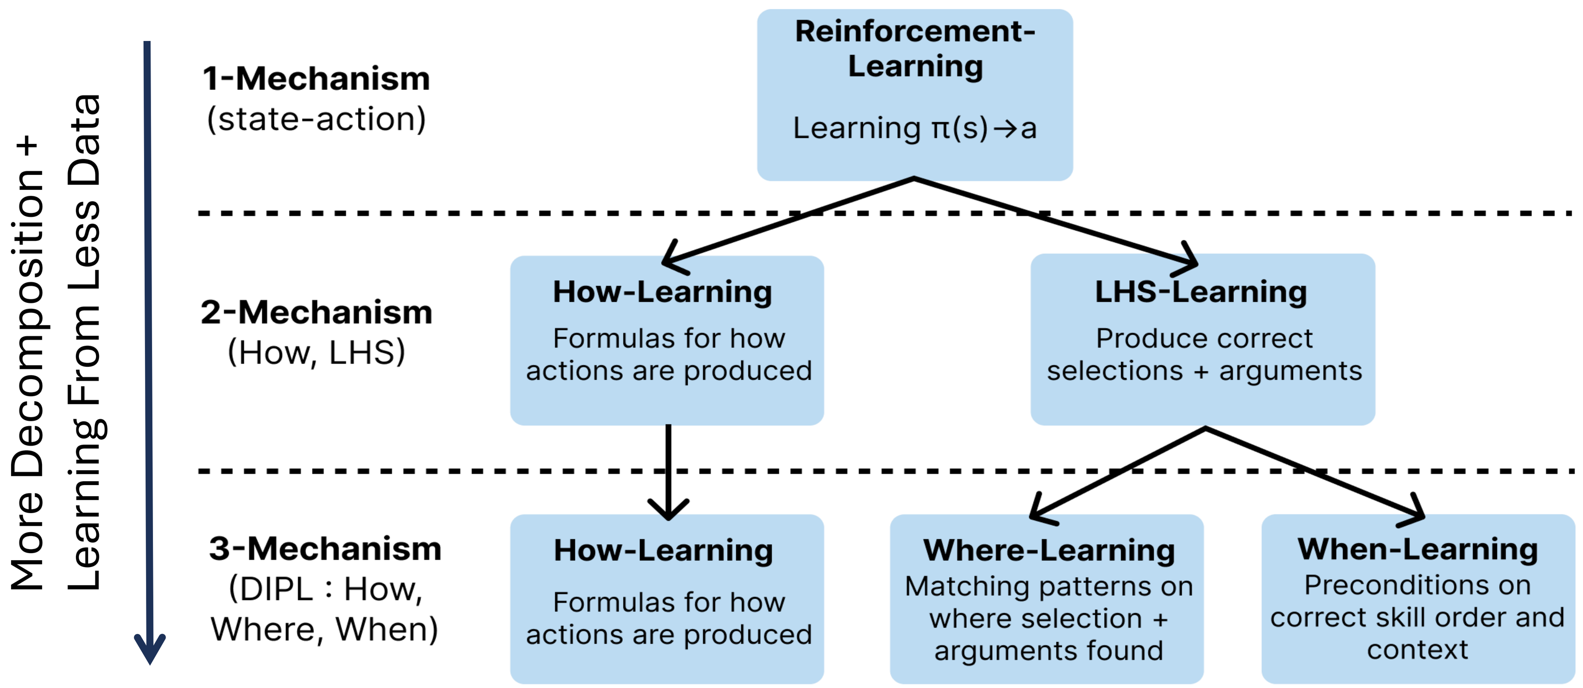


- RL: `π(s) → a`
- DIPL: `When(s,Where(s)) → How(Where(s)) = a`
  - Multiple symbolic pieces (How, Where, When)

## Motivation

- **Human Learning** is orders of magnitude **faster than** data-driven machine learning (ML) like **reinforcement learning**, which relies on **gradient-descent**
- Humans rely on specialization, **distinct cognitive mechanisms** working together to enable **rapid learning**.
- **Simulated learner systems:** Sierra[1], SimStudent[2], the Apprentice Learner (AL) architecture[3] and AI2T[4] match human learning rates in tutoring systems, and from human instruction.
- Are these systems **faster learners** because of their **symbolic learning mechanisms**, or because **learning mechanisms specialization improves learning efficiency**?

Paper Link Here!

## Two Tutoring System Tasks

Agents solve step-by-step with immediate step-correctness feedback

### Multi-Column Addition / Fraction Arithmetic

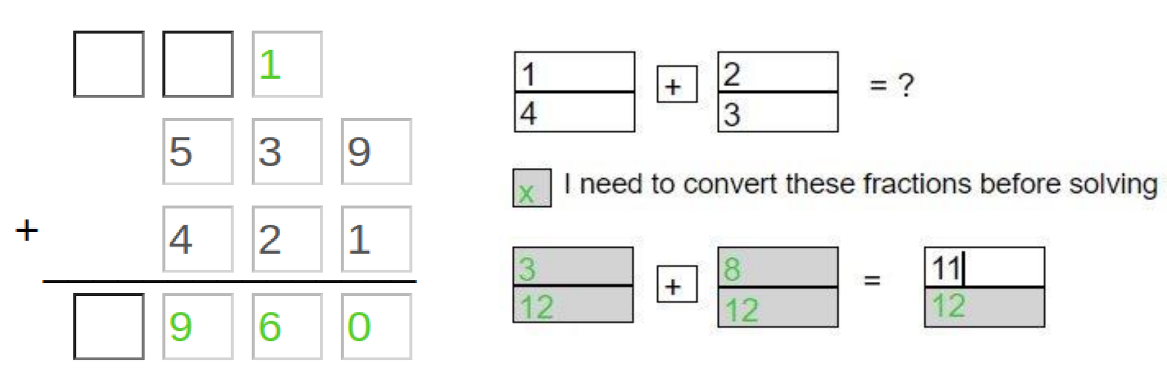


Long addition over 3-by-3 digit numbers

Adding fractions with different and same denominators and multiplying fractions

## Methods

**1-Mechanism:**
- Reinforcement Learning: PPO and DQN
- Decision Tree

**2-Mechanism (How+LHS):**
- Rules learned as compositions of functions to execute actions + decision tree to gate application

**3-Mechanism** (DIPL)
- With and without *relative featurization:* features are re-expressed relative to variables in rules.

**Training:** All models are given immediate feedback as they work step-by-step, and +Demos models are given a demo worked example after each incorrect action.

**RL Agents:** Given an enumerated action and state space. 2,702 unique actions in fractions, 71 in multi-column addition

## Results

- Each level of decomposition improves data-efficiency. Thus, **learning decomposition** is more **essential to achieving human-like learning** than symbolic learning mechanisms.
- **DIPL can learn with ~500x less data than the best RL models.**
- **DIPL shows similar learning rates as seen in human data.**
- Relative featurization improves performance.

| | | Fractions | MC Addition |
|---|---|---|---|
| 1-mech | PPO | Not Converge | 30,642 |
| | DQN+Demos | 11,315 | 9,496 |
| | DT+Demos | 1,944 | 7,816 |
| 2-mech | How+LHS | 17 | 270 |
| 3-mech | DIPL (no rel. feat.) | 33 | 38 |
| | DIPL | 20 | 19 |
| | Human Data | ~9-14 | N/A |

Table 1: Number of problems before $< 10\%$ average error.

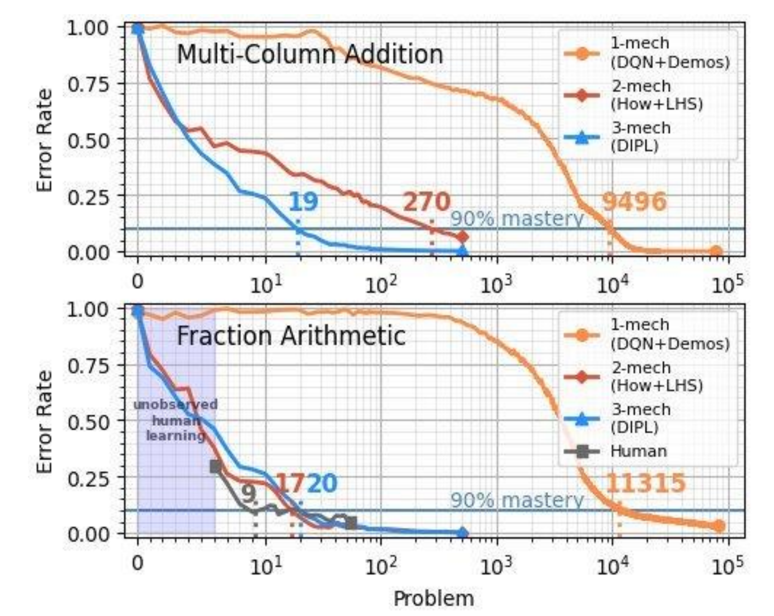


## Future Work

- Yet more decomposition could produce yet better data-efficiency
- For instance, inducing Hierarchical Task Networks (HTNs) could simplify the role of when-learning.

## How-Learning

Determines *how* skills apply actions using an abductive process, that searches for compositions of prior knowledge functions that self-explain a worked example.

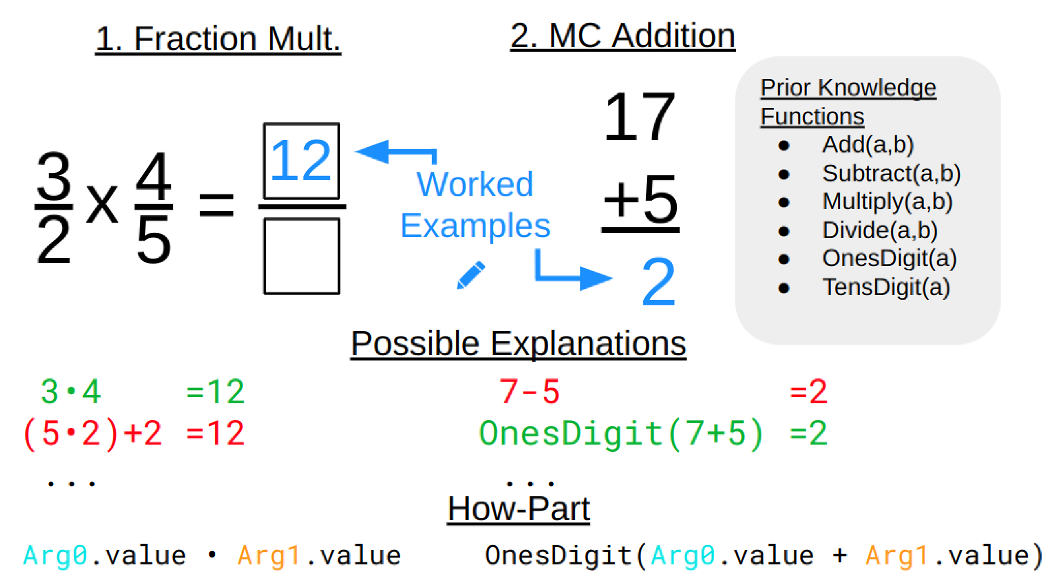


## Where-Learning

Discovers matching patterns to determine *where* skills can be applied. Where-learning builds spatial generalizations that identify where it may be possible to apply a skill.

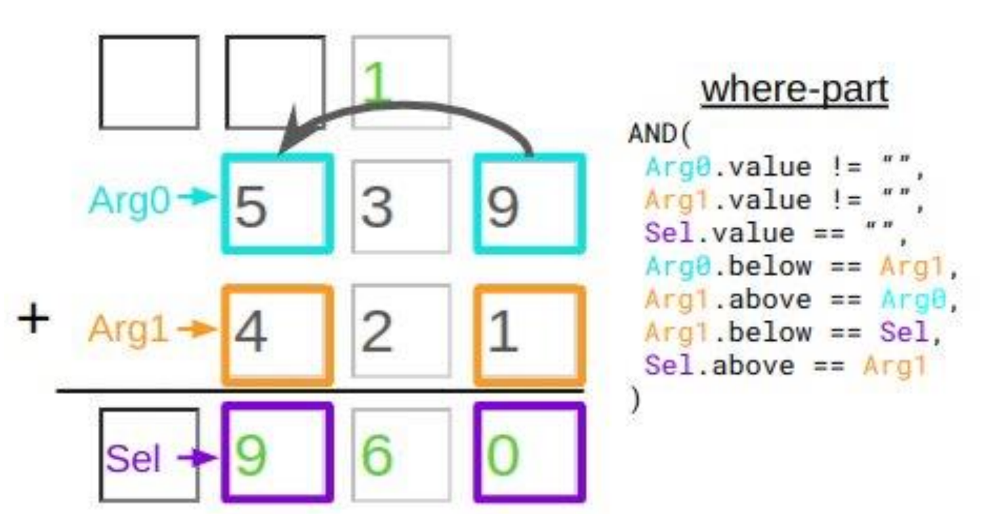


## When-Learning

Learns preconditions for skills to determine *when* they can be applied. Learns inductively from positive and negative examples

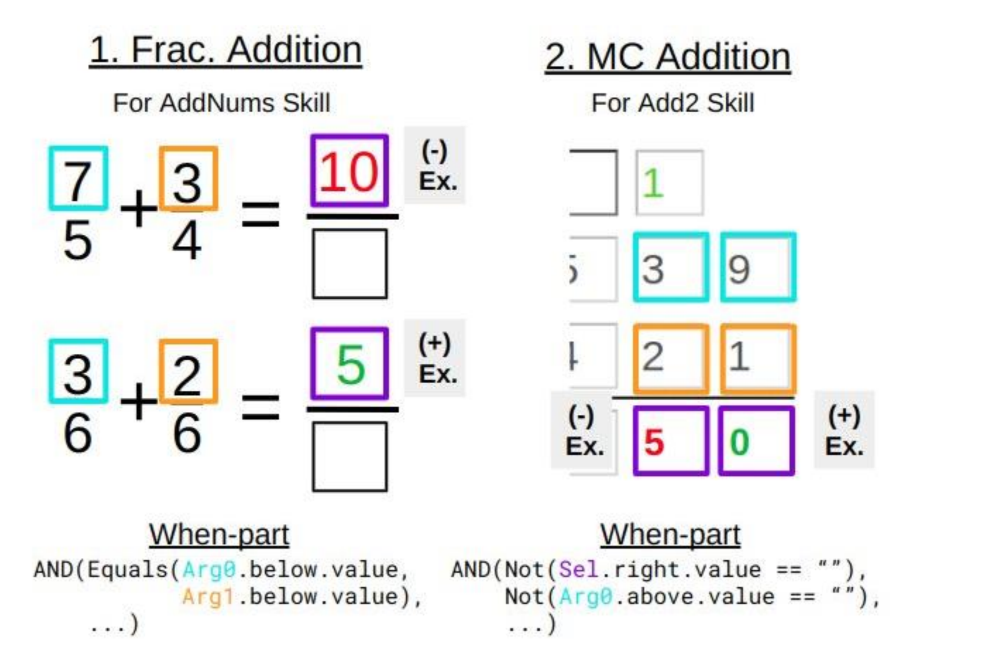


Georgia Institute of Technology

TAIL Teachable AI Lab

## References

1. VanLehn, K. (1990). Mind bugs: The origins of procedural misconceptions. MIT press.
2. Matsuda, N., Cohen, W. W., & Koedinger, K. R. (2015). Teaching the teacher: Tutoring simstudent leads to more effective cognitive tutor authoring. International Journal of Artificial Intelligence in Education, 25(1), 1–34
3. MacLellan, C. J., & Koedinger, K. R. (2020). Domain-general tutor authoring with apprentice learner models. International Journal of Artificial Intelligence in Education, 1–42
4. Weitekamp, D., Harpstead, E., & Koedinger, K. (2024). AI2T: Building trustable ai tutors by interactively teaching a self-aware learning agent. arXiv preprint arXiv:2411.17924.

SCAN FOR PAPER