Learning Fast and Slow: A Redux of Levels of Learning in General Autonomous Intelligent Agents

Shiwali Mohan[1] **and John E. Laird**[2]

[1] SRI International, Palo Alto, California, USA
[2] Center for Integrated Cognition, Ann Arbor, Michigan, USA
shiwali.mohan@sri.com, laird@cic.iqmri.org

The world that general autonomous intelligent agents (e.g., humans) operate in is complex, dynamic, partially observable, and unknown. We must continually react to their environment, focusing our computational resources on making the best decision for the current situation using all our available knowledge. We need to learn everything we can from our experience, building up our knowledge so that we are prepared to make the best decisions in the future. To maximize our available knowledge, we have evolved to be effective and efficient learners.

We revisit the thesis we proposed originally in Laird and Mohan (2018). Building upon the dual process theory (Kahneman 2011), our thesis posits that in human-like agents, learning can be split into two levels and is inspired by research across neuroscience, psychology, and cognitive architecture (Laird, Lebiere, and Rosenbloom 2017). Level 1, or L1 learning, encompasses *architectural learning mechanisms* that are innate, automatic, online, and effortless, such as the temporal difference update mechanism in reinforcement learning. Level 2, or L2, encompasses *deliberate learning strategies* that are realized through knowledge and controlled by an agent. These strategies are essentially tasks that the agent adopts, and they compete with other tasks for mental (and physical) resources, and they depend on L1 mechanisms to do the actual "learning." In addition, they themselves can be learned. A simple example in humans is deciding to explicitly rehearse a phone number to memorize it. Deliberately repeating the number several times aloud (or to one's self) creates experiences that are consolidated by automatic L1 memory mechanisms, making the number available for later recall. L2 strategies may include additional learning mechanisms and take advantage of the underlying L1 mechanisms to extract regularities and record knowledge structures from the generated experiences.

Our thesis is based on our experiences in building artificial agents that have complex learning behaviors. Interactive task learning (Mohan and Laird 2014; Laird et al. 2017; Kirk and Laird 2016; Mininger and Laird 2018) enables an agent to learn from natural human teaching. Open-world learning (Mohan et al. 2023; Piotrowski et al. 2023) enables an agent to learn in a continually evolving, unknown world. In these agent architectures, we studied and developed L2 learning strategies that generate experiences for L1 learning mechanisms that are continually and incrementally learning. Our talk will discuss the design decisions in developing L2 mechanisms, their reliance on L1 mechanisms, and their impact on the agent's learning behaviors.

A discussion on levels of learning is pertinent in the current age of intelligent systems built upon deep learning and generative processes. While these systems have shown powerful, flexible behavior and in-context learning, they lack the human capability of *online* learning and knowledge acquisition, both volitional automatic. This discussion will point the way for robust, human-like adaptive intelligent systems.

References

Kahneman, D. 2011. *Thinking, Fast and Slow*. Macmillan.

Kirk, J. R.; and Laird, J. E. 2016. Learning General and Efficient Representations of Novel Games Through Interactive Instruction. *Advances in Cognitive Systems*, 4.

Laird, J.; Lebiere, C.; and Rosenbloom, P. 2017. A Standard Model of the Mind: Toward a Common Computational Framework Across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine*, 38(4).

Laird, J.; and Mohan, S. 2018. Learning Fast and Slow: Levels of Learning in General Autonomous Intelligent Agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Laird, J. E.; Gluck, K.; Anderson, J.; Forbus, K. D.; Jenkins, O. C.; Lebiere, C.; Salvucci, D.; Scheutz, M.; Thomaz, A.; Trafton, G.; et al. 2017. Interactive Task Learning. *IEEE Intelligent Systems*, 32(4): 6–21.

Mininger, A.; and Laird, J. E. 2018. Interactively Learning a Blend of Goal-Based and Procedural Tasks. In *AAAI*.

Mohan, S.; and Laird, J. E. 2014. Learning Goal-Oriented Hierarchical Tasks from Situated Interactive Instruction. In *AAAI*, 387–394.

Mohan, S.; Piotrowski, W.; Stern, R.; Grover, S.; Kim, S.; Le, J.; and De Kleer, J. 2023. A Domain-Independent Agent Architecture for Adaptive Operation in Evolving Open Worlds. *arXiv preprint arXiv:2306.06272*.

Piotrowski, W.; Stern, R.; Sher, Y.; Le, J.; Klenk, M.; deKleer, J.; and Mohan, S. 2023. Learning to Operate in Open Worlds by Adapting Planning Models. *International Conference on Autonomous Agents and Multiagent Systems*.