Class Discussion

Gene Lewis, SAIRG

"Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which **simulates the child's**?

If this were then subjected to an appropriate course of education one would **obtain the adult brain.**"

- Alan Turing

Challenges in Intelligence

a woman riding a horse on a dirt road

an airplane is parked on the tarmac at an airport

a group of people standing on top of a beach

Karpathy, Fei-Fei, “Deep Visual Semantic Alignments for Generating Image Descriptions”

Lake et al, “Building Machines that Learn and Think Like People”

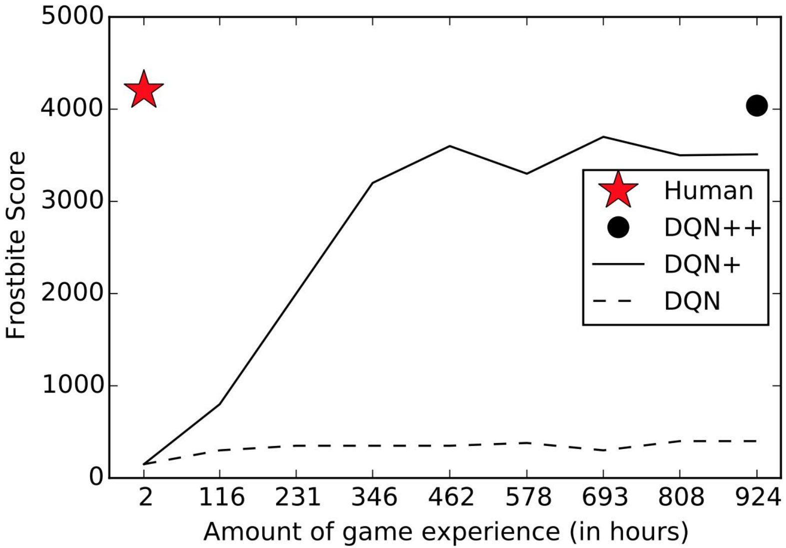

Mnih et al, “Playing Atari with Deep Reinforcement Learning”

Lake et al, “Building Machines that Learn and Think Like People”

- Get the lowest possible score.
- Get closest to 100, or 300, or 1000, or 3000, or any level, without going over.
- Beat your friend, who's playing next to you, but just barely, not by too much, so as not to embarrass them.
- Go as long as you can without dying.
- Die as quickly as you can.
- Pass each level at the last possible minute, right before the temperature timer hits zero and you die (i.e., come as close as you can to dying from frostbite without actually dying).
- Get to the furthest unexplored level without regard for your score.
- See if you can discover secret Easter eggs.
- Get as many fish as you can.
- Touch all the individual ice floes on screen once and only once.
- Teach your friend how to play as efficiently as possible.

Lake et al, "Building Machines that Learn and Think Like People"

Grounding as Extracted Structure

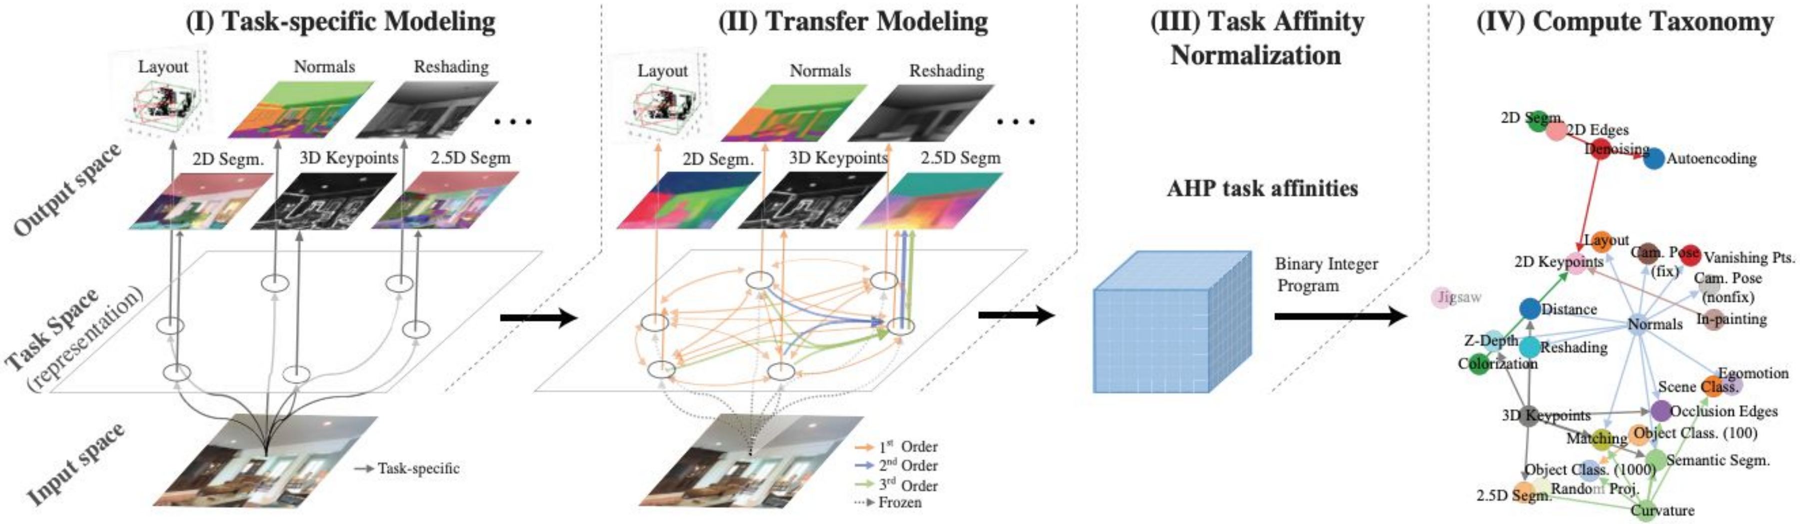

Zamir et al, “Taskonomy: Disentangling Task Transfer Learning”

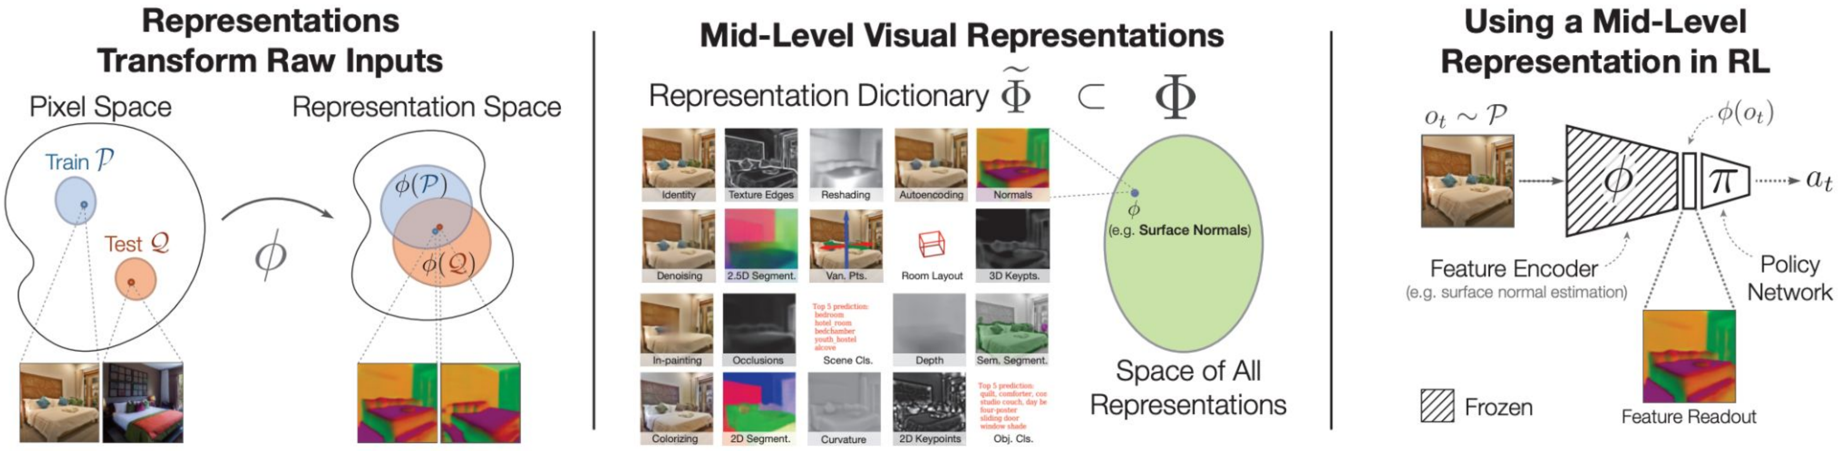

Sax et al, “Learning to Navigate Using Mid-Level Visual Priors”

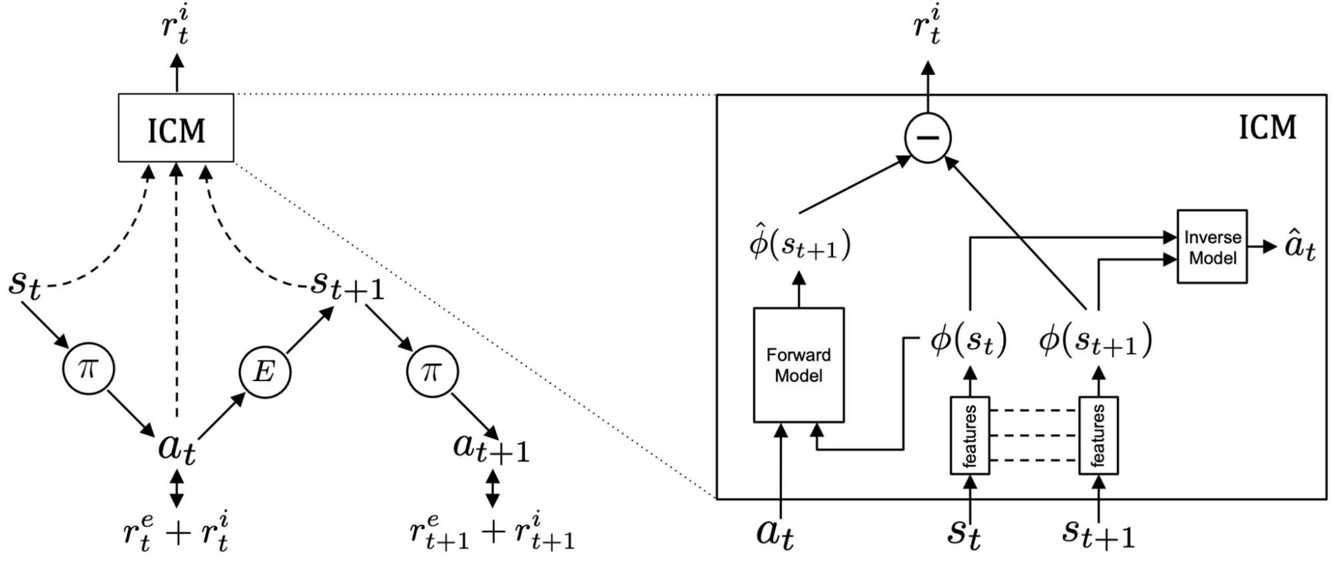

Pathak et al, "Curiosity-driven Exploration by Self-supervised Prediction"

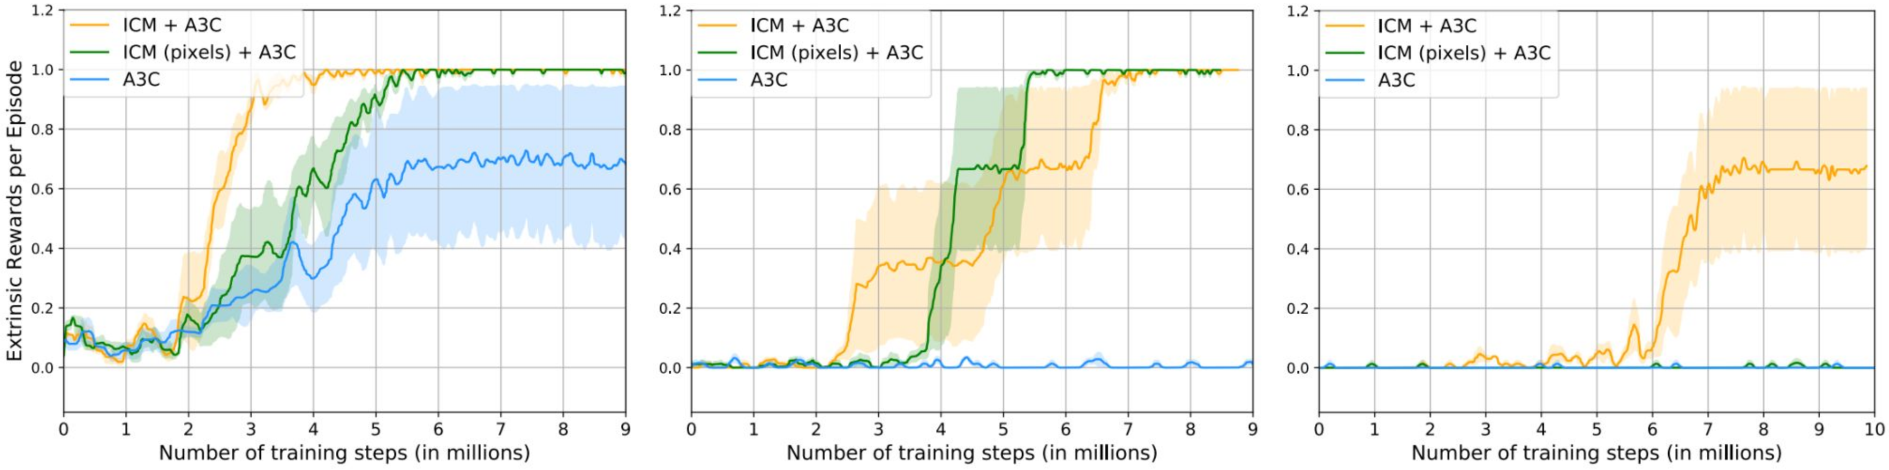

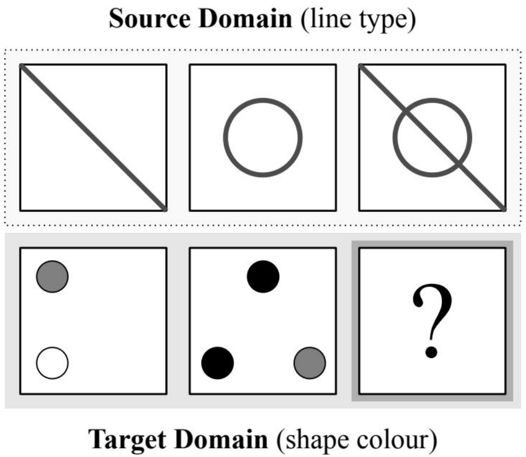

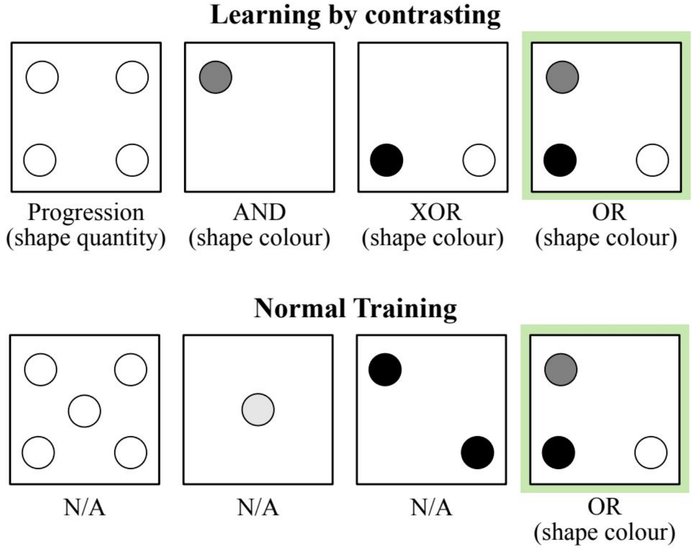

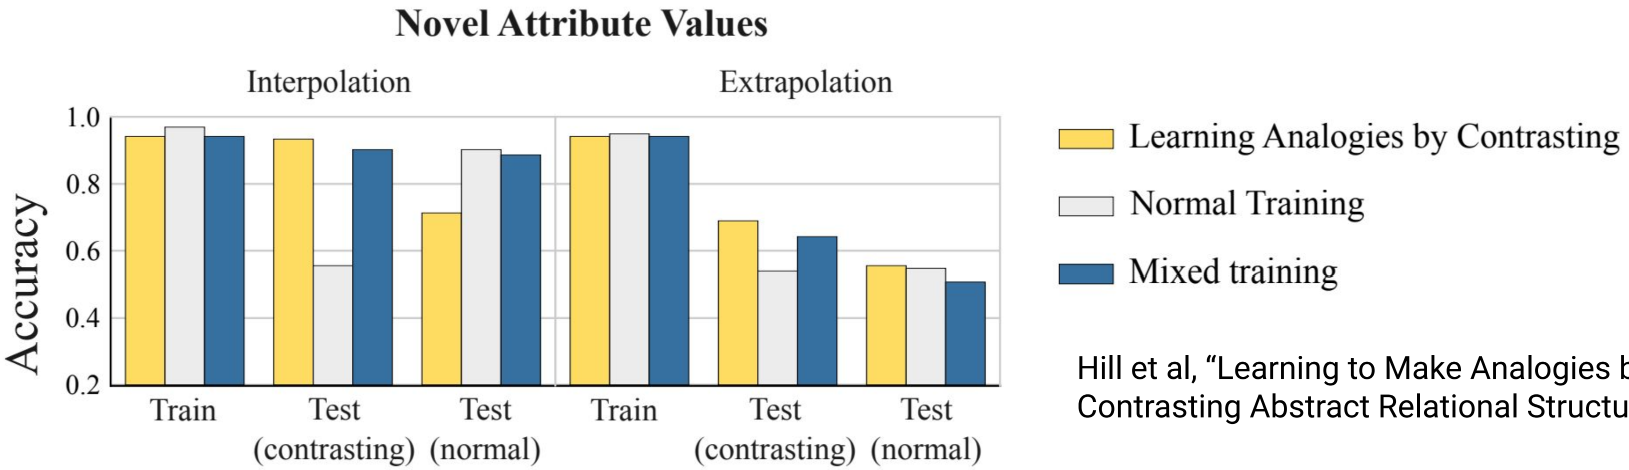

Hill et al, “Learning to Make Analogies by Contrasting Abstract Relational Structure”

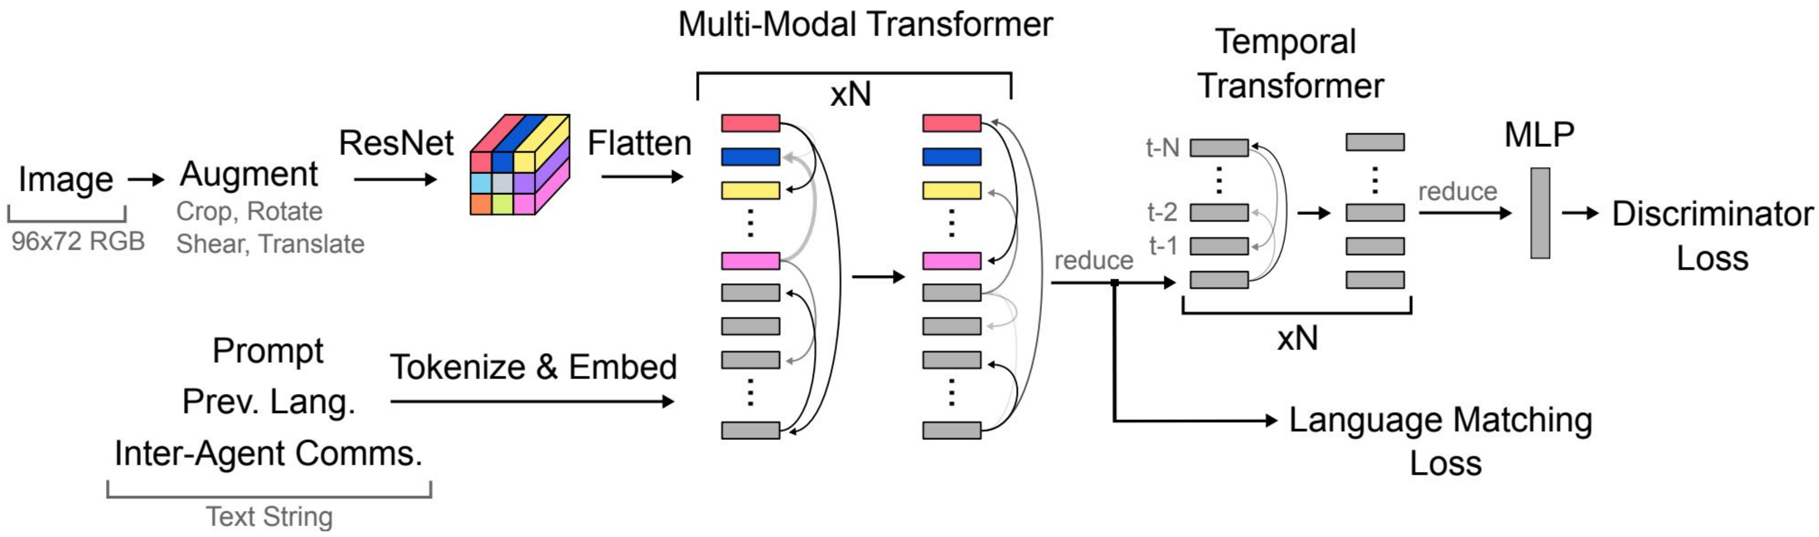

Abramson et al, “Imitating Interactive Intelligence”

Structured Problem Solving

partially observed Markov decision process $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{E}, r\}$

$\mathcal{S}$ – state space states $s \in \mathcal{S}$ (discrete or continuous)

$\mathcal{A}$ – action space actions $a \in \mathcal{A}$ (discrete or continuous)

$\mathcal{O}$ – observation space observations $o \in \mathcal{O}$ (discrete or continuous)

$\mathcal{T}$ – transition operator (like before)

$\mathcal{E}$ – emission probability $p(o_t|s_t)$

r – reward function $r : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$

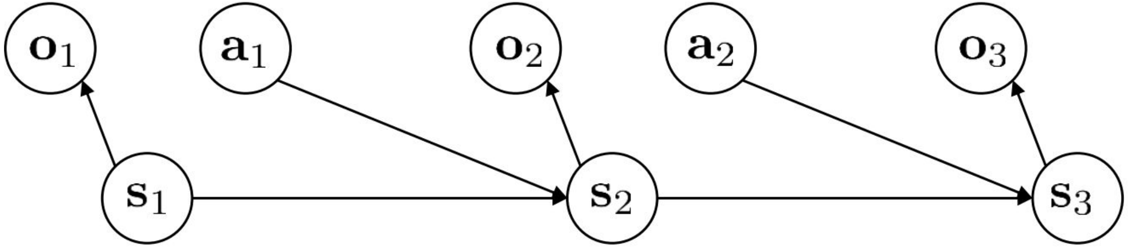

Levine et al, CS285 Lecture 4 Slides: “Introduction to Reinforcement Learning”

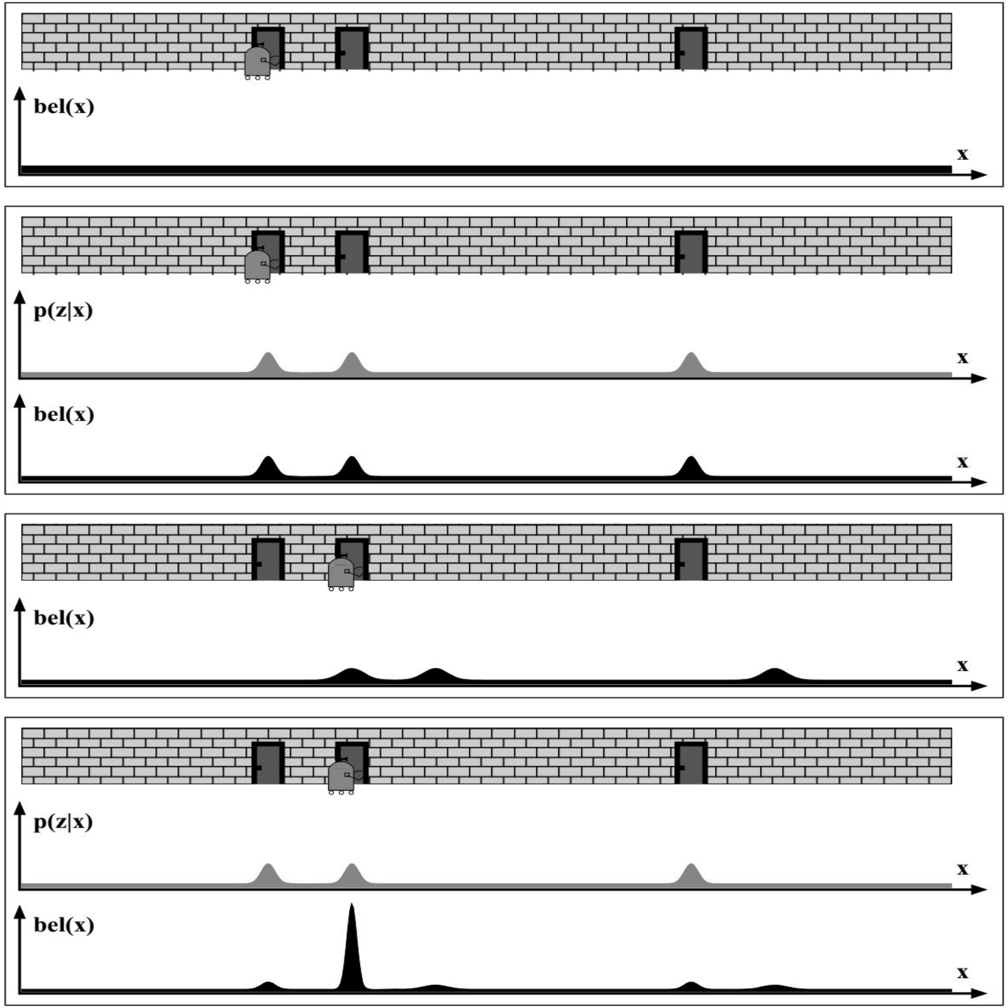

$$p(x \mid y) = \frac{p(y \mid x)\, p(x)}{p(y)}$$

$$= \frac{p(y \mid x)\, p(x)}{\int p(y \mid x')\, p(x')\, dx'}$$

Thrun et al, *Probabilistic Robotics*

Complementary Learning Systems

Goals:	Remember Specifics	Extract Generalities
Example:	Where is car parked?	Best parking strategy?
Need to:	Avoid interference	Accumulate experience
	Solution:	
1.	Separate reps (keep days separate) 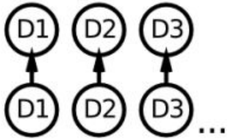	Overlapping reps (integrate over days) 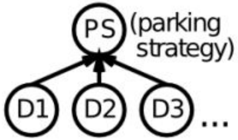
2.	Fast learning (encode immediately)	Slow learning (integrate over days)
3.	Learn automatically (encode everything)	Task-driven learning (extract relevant stuff)
	These are incompatible, need two different systems:	
System:	Hippocampus	Neocortex

O'Reilly et al, *Computational Cognitive Neuroscience 4th ed.*

Key Figure

Complementary Learning Systems (CLS) and their Interactions.

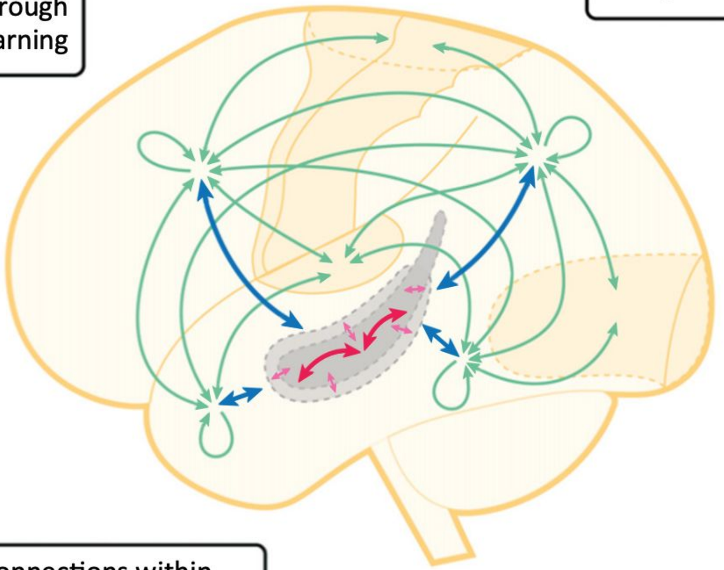

Kumaran et al, “What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated”

c. MERLIN

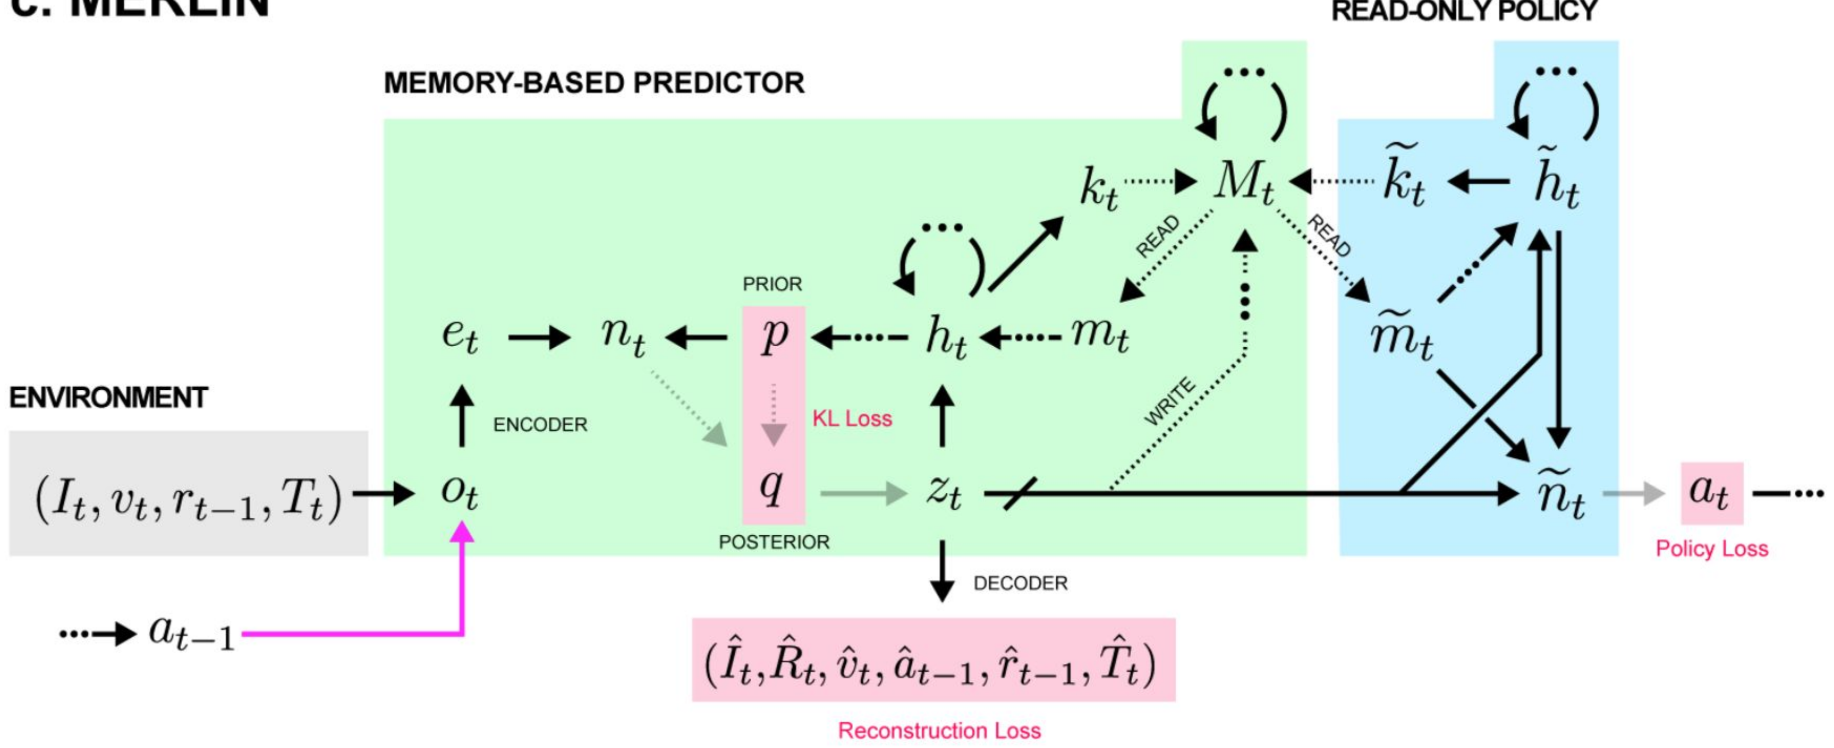

Wayne et al, “Unsupervised Predictive Memory in a Goal-Directed Agent”

“A **human being should be able to** change a diaper, plan an invasion, butcher a hog, conn a ship, design a building, write a sonnet, balance accounts, build a wall, set a bone, comfort the dying, take orders, give orders, cooperate, act alone, solve equations, analyse a new problem, program a computer, cook a tasty meal, fight efficiently, die gallantly. **Specialization is for insects.**”

- Robert Heinlein

“A **human being should be able to** change a diaper, plan an invasion, butcher a hog, conn a ship, design a building, write a sonnet, balance accounts, build a wall, set a bone, comfort the dying, take orders, give orders, cooperate, act alone, solve equations, analyse a new problem, **program a computer**, cook a tasty meal, fight efficiently, die gallantly. **Specialization is for insects.**”

- Robert Heinlein

Thank You