Tutorial at IJCAI, August 19th /20th 2021

Neural Machine Reasoning

Lecture 2: Dual system of reasoning

Truyen Tran, Vuong Le, Hung Le and Thao Le
{truyen.tran,vuong.le,thai.le,thao.le}@deakin.edu.au

https://neuralreasoning.github.io

Lecture 2: Sub-topics

- **Problem with current DL**
- Dual system theories
- Existing ideas for System 2
- Theory of mind

DL has been fantastic, but …

- It is great at interpolating
  - → data hungry to cover all variations and smooth local manifolds
  - → fail to handle change of distributions
  - → little systematic generalization (novel combinations)
- Lack of human-perceived reasoning capability
  - Lack natural mechanism to incorporate prior knowledge, e.g., common sense
  - No built-in causal mechanisms
  - → Have trust issues!
- To be fair, may of these problems are common in statistical learning!

Problem: Systematic generalization

- Refers to the ability to work robustly with new combinations with zero probability in training data.
  - E.g., if we understand ‘John loves Mary’, then we can also understand ‘Mary loves John’, but machine may fail due to zero probability of the latter if not done properly.
- Current DL has a major problem with it.
  - This is not new: Has been argued for 30+ years!
- Much research is needed on multiple fronts (e.g., syntax, indirection, datasets, measuring)

Bahdanau, Dzmitry, et al. "Systematic generalization: what is required and can it be learned?." *arXiv preprint arXiv:1811.12889* (2018).

Fodor, Jerry A., and Zenon W. Pylyshyn. "Connectionism and cognitive architecture: A critical analysis." *Cognition* 28.1-2 (1988): 3-71.

Problem: Out-of-distribution

- Data, context change, both life-long and life-wide, sometimes rapidly (e.g., context switch), sometimes slowly (e.g., aging)
- Other agents in the play → non-stationaries
- Continual learning is needed → need to handle catastrophic forgetting.

Lecture 2: Sub-topics

- Problem with current DL
- **Dual system theories**
- Existing ideas for System 2
- Theory of mind

Hypothesis: We need System 2

- Decoupling from perception/representation (which deep learning does well)
- Holds hypothetical thought
  - Enabling mental travels & imagination.
- Slow. Deliberative. Conscious.
- Needs working memory. But the size is not essential. Its attentional control is.

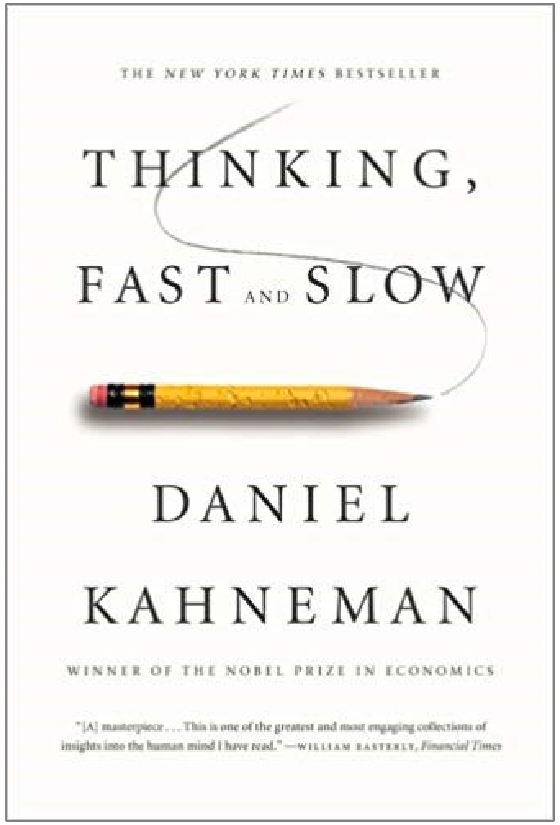

References	System 1	System 2
Fodor (1983, 2001)	Input modules	Higher cognition
Schneider & Schiffrin (1977)	Automatic	Controlled
Epstein (1994), Epstein & Pacini (1999)	Experiential	Rational
Chaiken (1980), Chen & Chaiken (1999)	Heuristic	Systematic
Reber (1993), Evans & Over (1996)	Implicit/tacit	Explicit
Evans (1989, 2006)	Heuristic	Analytic
Sloman (1996), Smith & DeCoster (2000)	Associative	Rule based
Hammond (1996)	Intuitive	Analytic
Stanovich (1999, 2004)	System 1 (TASS)	System 2 (Analytic)
Nisbett et al. (2001)	Holistic	Analytic
Wilson (2002)	Adaptive unconscious	Conscious
Lieberman (2003)	Reflexive	Reflective
Toates (2006)	Stimulus bound	Higher order
Strack & Deustch (2004)	Impulsive	Reflective

Evans, Jonathan St BT. "Dual-processing accounts of reasoning, judgment, and social cognition." *Annu. Rev. Psychol.* 59 (2008): 255-278.

System 2 may have two layers: Reflective and Algorithmic

Photo credit: mumsgrapevine

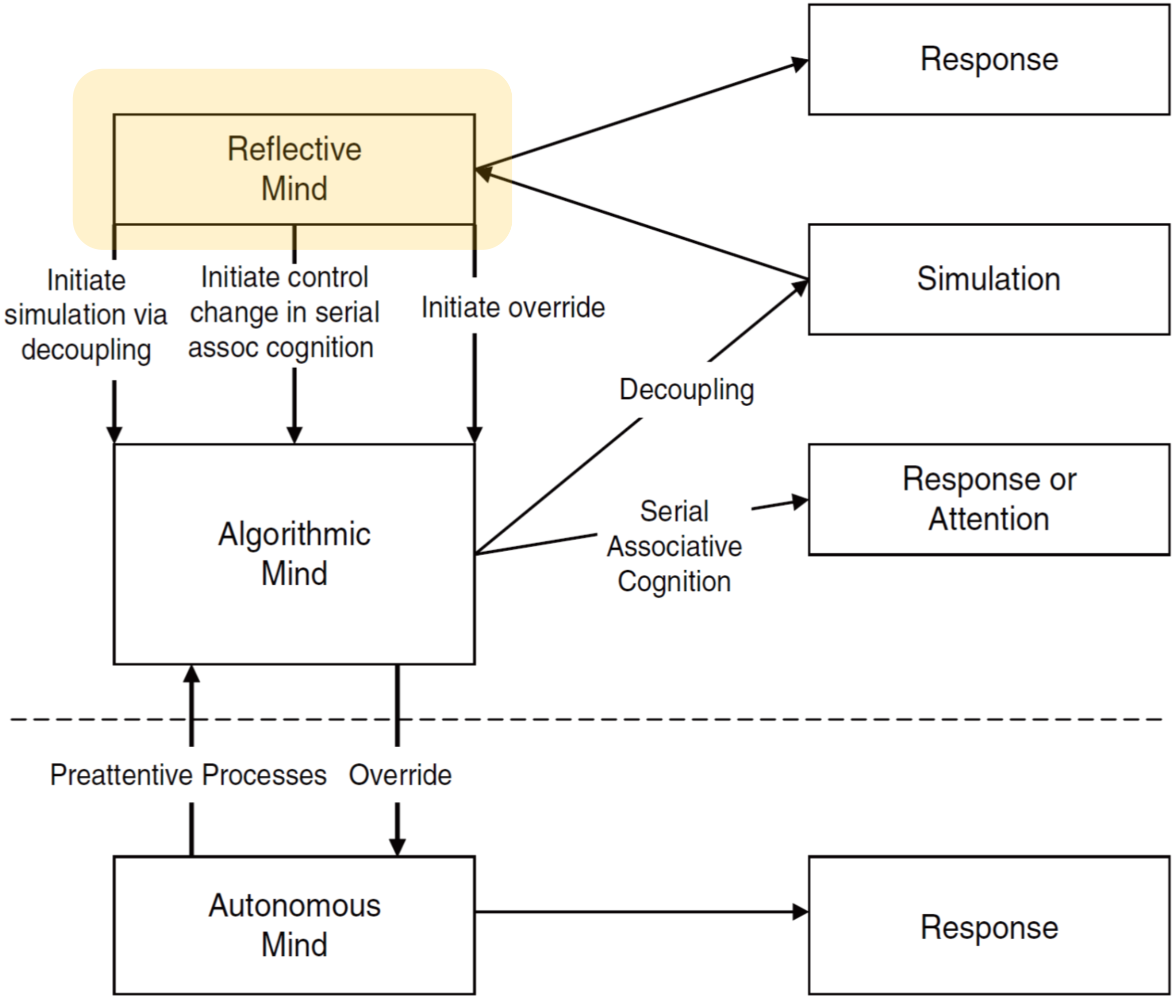

Stanovich, K. E. (2009). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory. *In two minds: Dual processes and beyond*, 55-88.

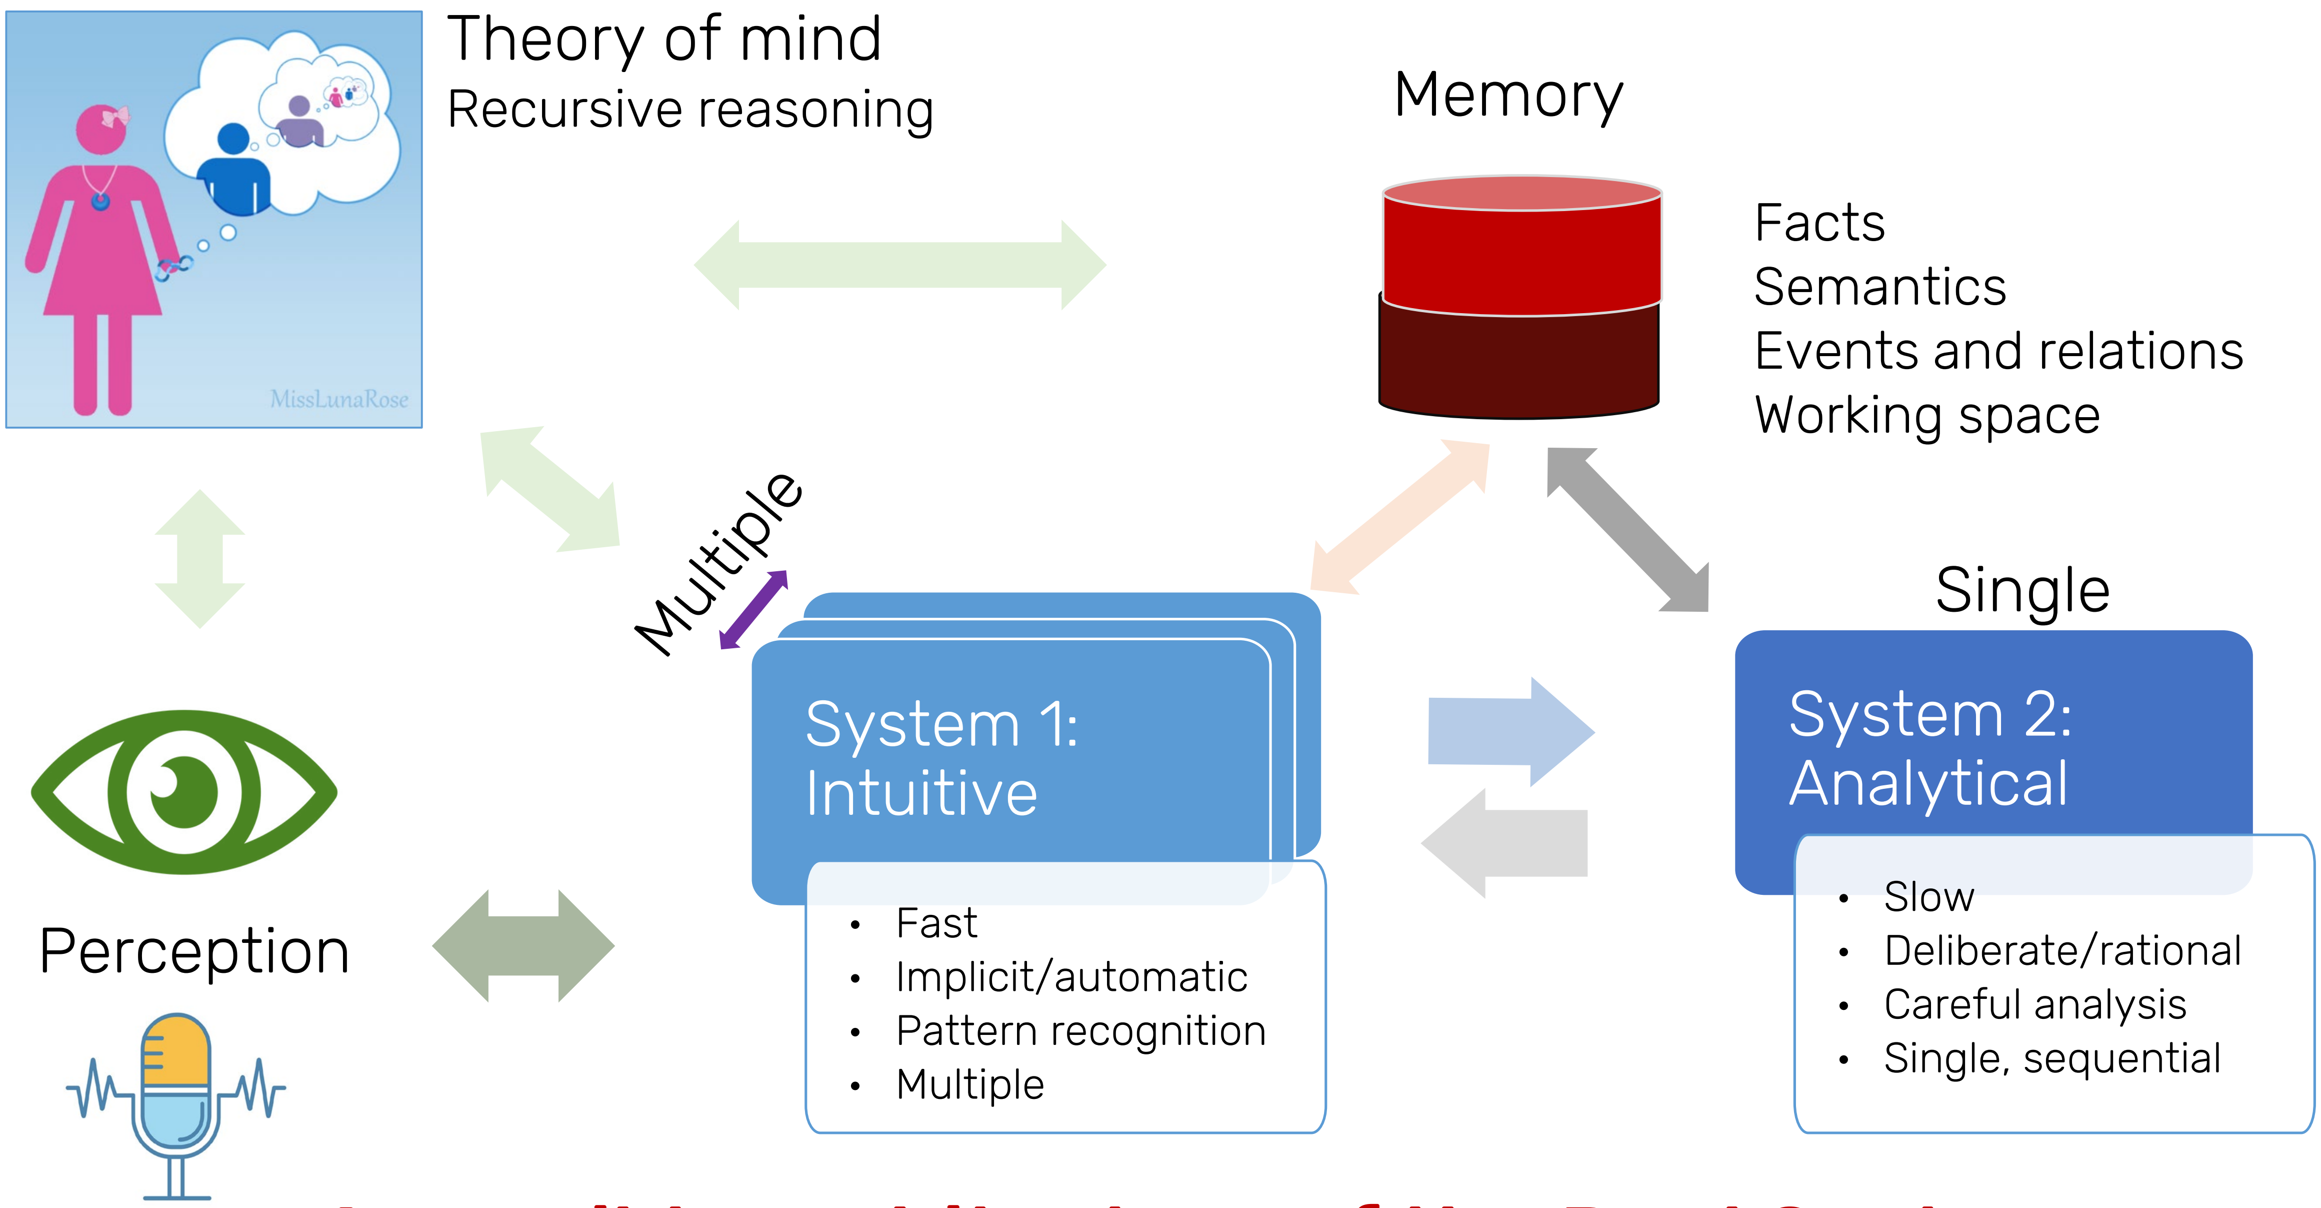

A possible architecture of the Dual System

Image credit: VectorStock | Wikimedia

Lecture 2: Sub-topics

- Problem with current DL
- Dual system theories
- **Existing ideas for System 2**
- Theory of mind

More reading: Greff, Klaus, Sjoerd van Steenkiste, and Jürgen Schmidhuber. "On the binding problem in artificial neural networks." *arXiv preprint arXiv:2012.05208* (2020).

Object-concept binding

- Perceived data (e.g., visual objects) may not share the same semantic space with high-level concepts.
- Binding between concept-object enables reasoning at the concept level

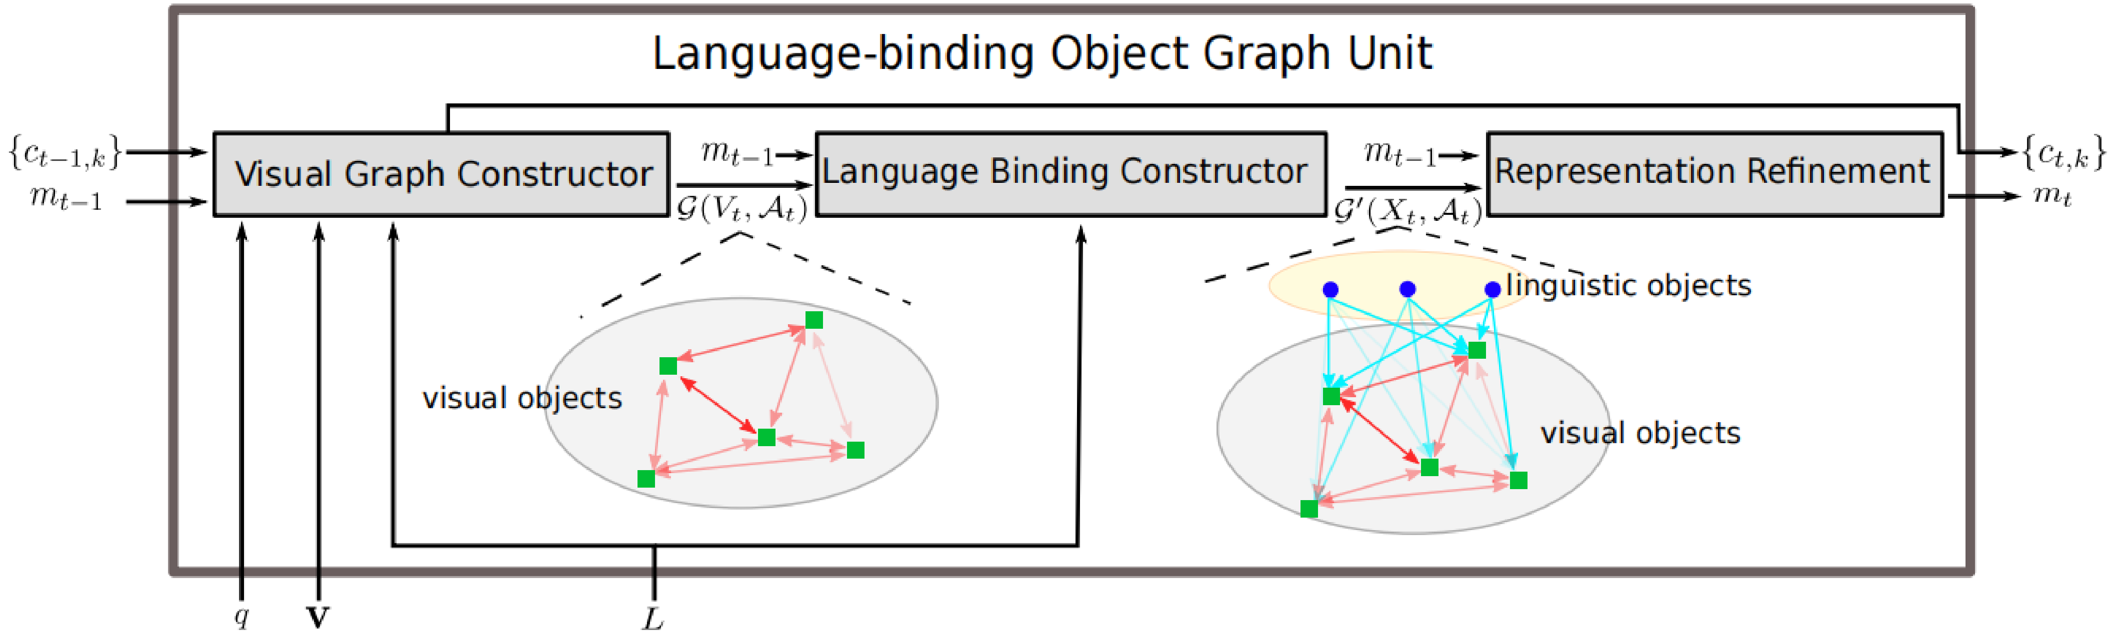

Example of concept-object binding in LOGNet (Le et al, IJCAI'2020)

Object-concept binding (cont.)

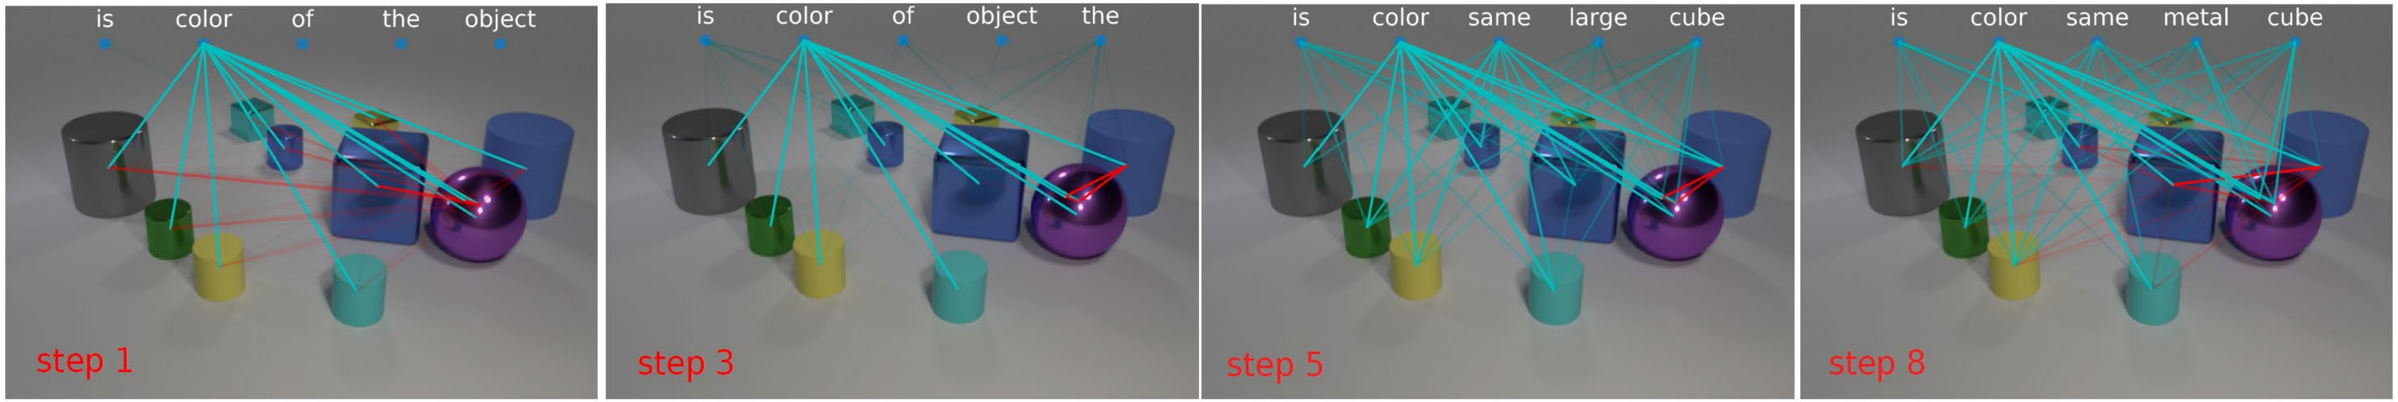

Question: Is the color of the big matte object the same as the large metal cube?
Prediction: yes **Answer**: yes

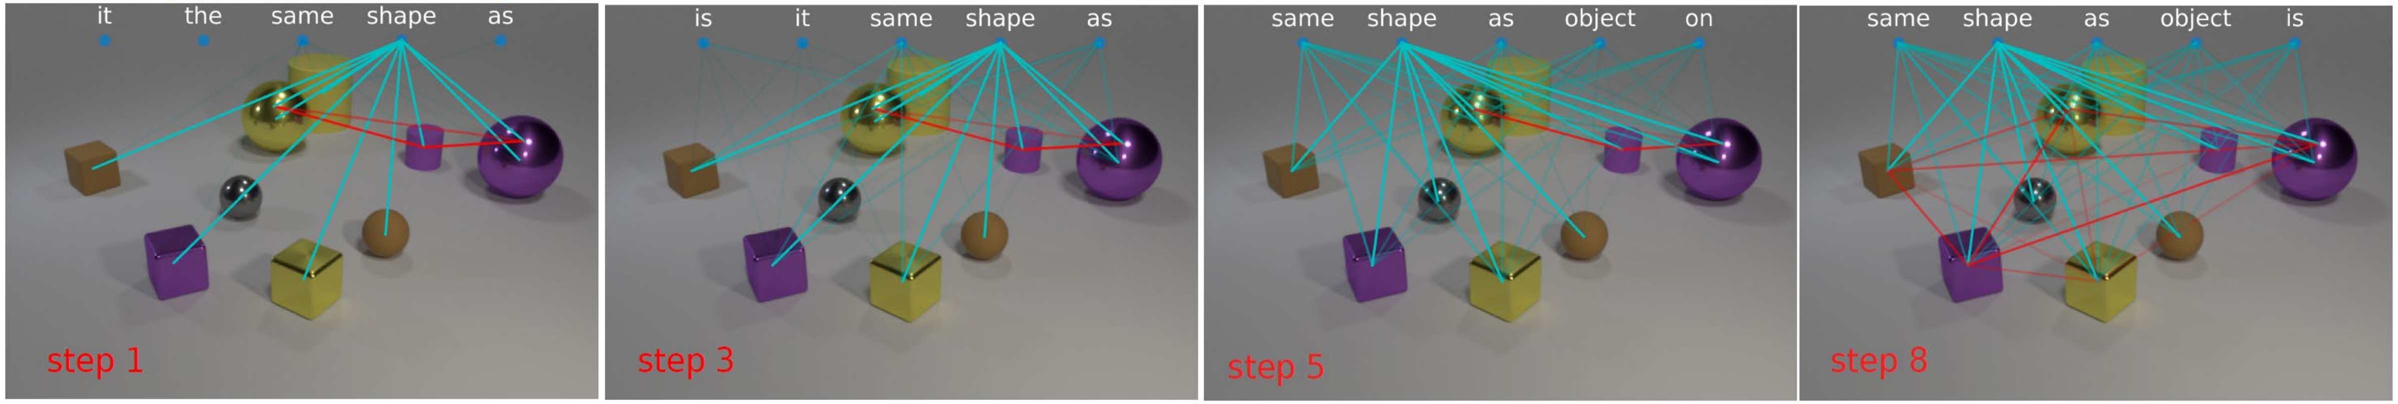

Question: There is a tiny purple rubber thing; does it have the same shape as the brown object that is on the left side of the rubber sphere?
Prediction: no **Answer**: no

Attention & Indirection

- Focus on the most relevant pieces for each reasoning step.
  - Piece = item, relation & sub-program/module.
- When piece is pointer to others, we have indirection, a powerful way to generalize to different representations if the “names” of items & relations remain.
- May need ability to “zoom in” – coarse to fine attention.
  - E.g., face detection → eye detection → eye corners

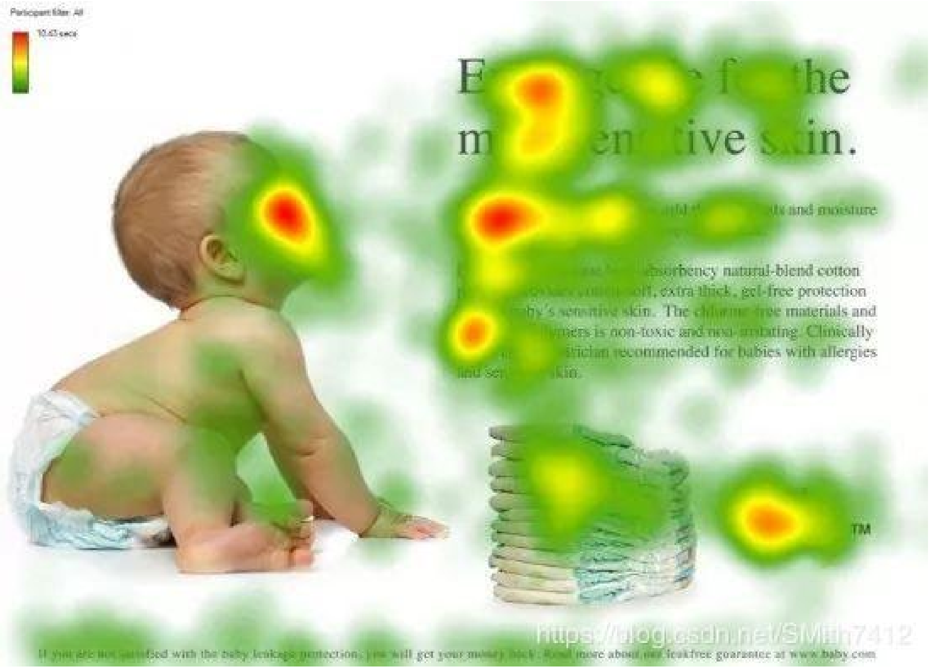

Photo: programmersought

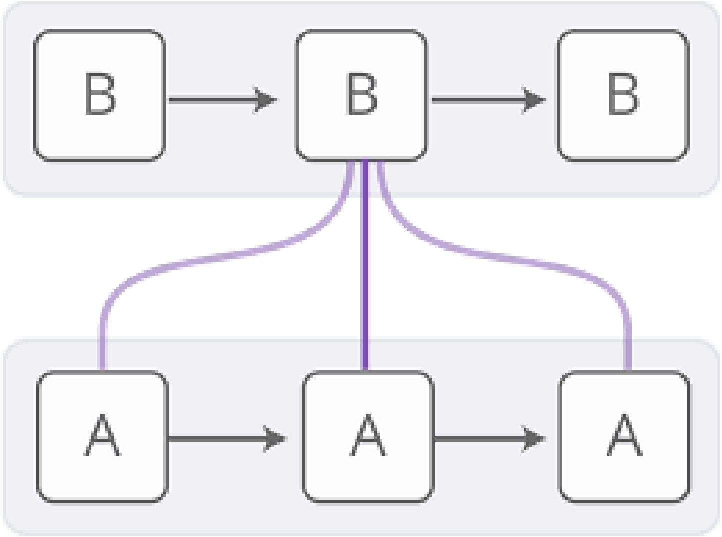

http://distill.pub/2016/augmented-rnns/

Iterative message passing in BP

- It iteratively computes "beliefs" of unobserved variables based on evidences from observed variables.
- Known result in 2001-2003: BP minimises Bethe free-energy minimization.

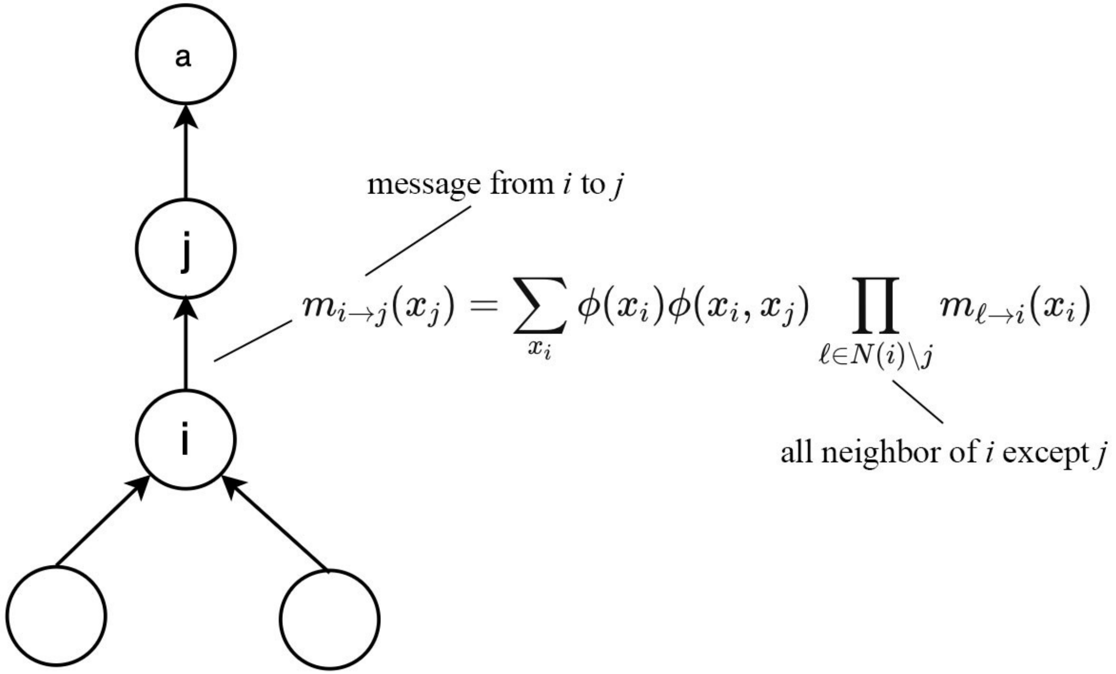

- Does BP qualify as a deliberative mechanism for System 2?

Figure credit: Jonathan Hui

Heskes, Tom. "Stable fixed points of loopy belief propagation are local minima of the bethe free energy." *Advances in neural information processing systems*. 2003.

Neural graph message passing

Relation graph

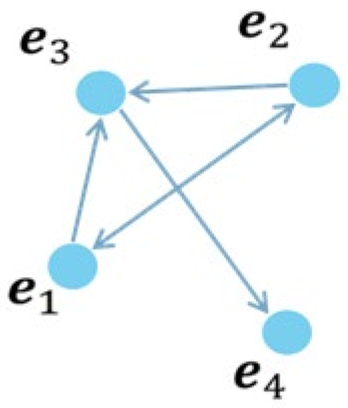

GCN update rule, vector form

$$h_{v_i}^{(l+1)} = \sigma\left(\sum_j \frac{1}{c_{ij}} h_{v_j}^{(l)} W^{(l)}\right)$$

GCN update rule, matrix form

$$f(H^{(l)}, A) = \sigma\left(\hat{D}^{-\frac{1}{2}}\hat{A}\hat{D}^{-\frac{1}{2}} H^{(l)} W^{(l)}\right)$$

Generalized message passing

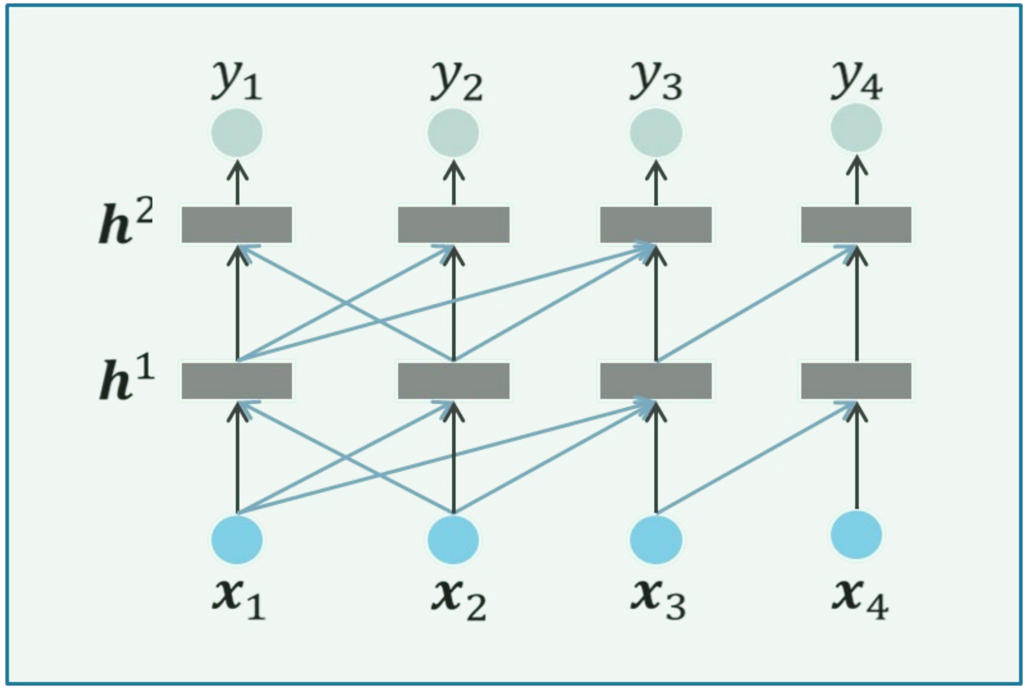

#REF: Pham, Trang, et al. "Column Networks for Collective Classification." *AAAI*. 2017.

What we have in store: Modular recurrences

- RIM: Recurrent Independent Mechanisms

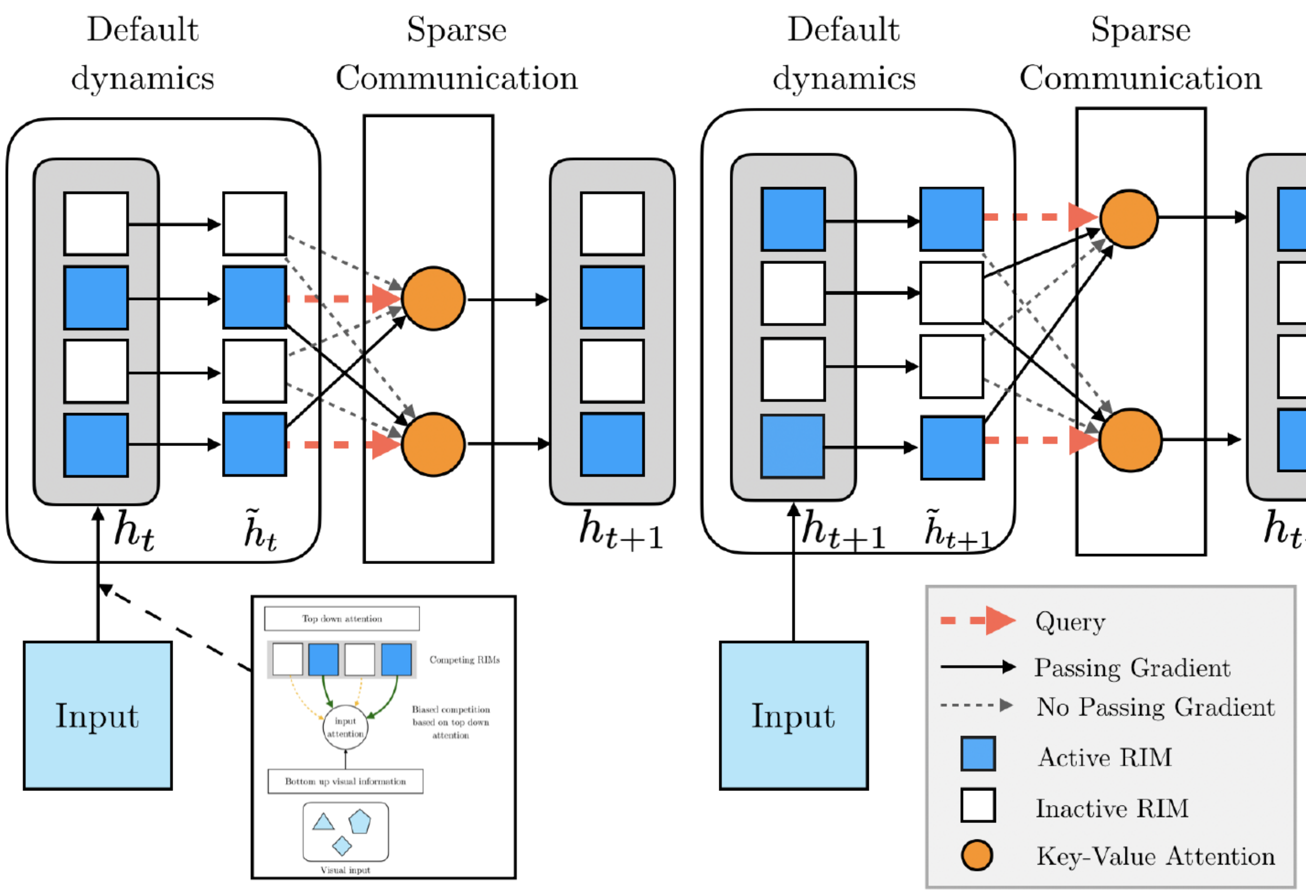

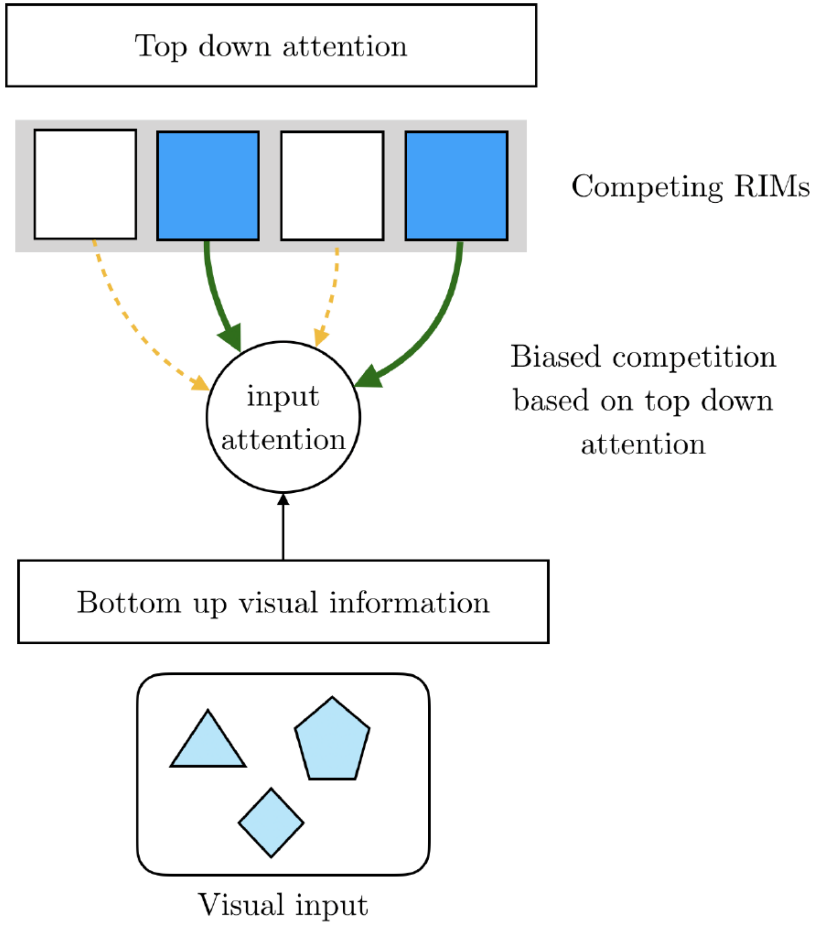

Goyal, Anirudh, et al. "Recurrent independent mechanisms." *arXiv preprint arXiv:1909.10893* (2019).

Self-attentive associative memories (SAM)

Learning relations automatically over time

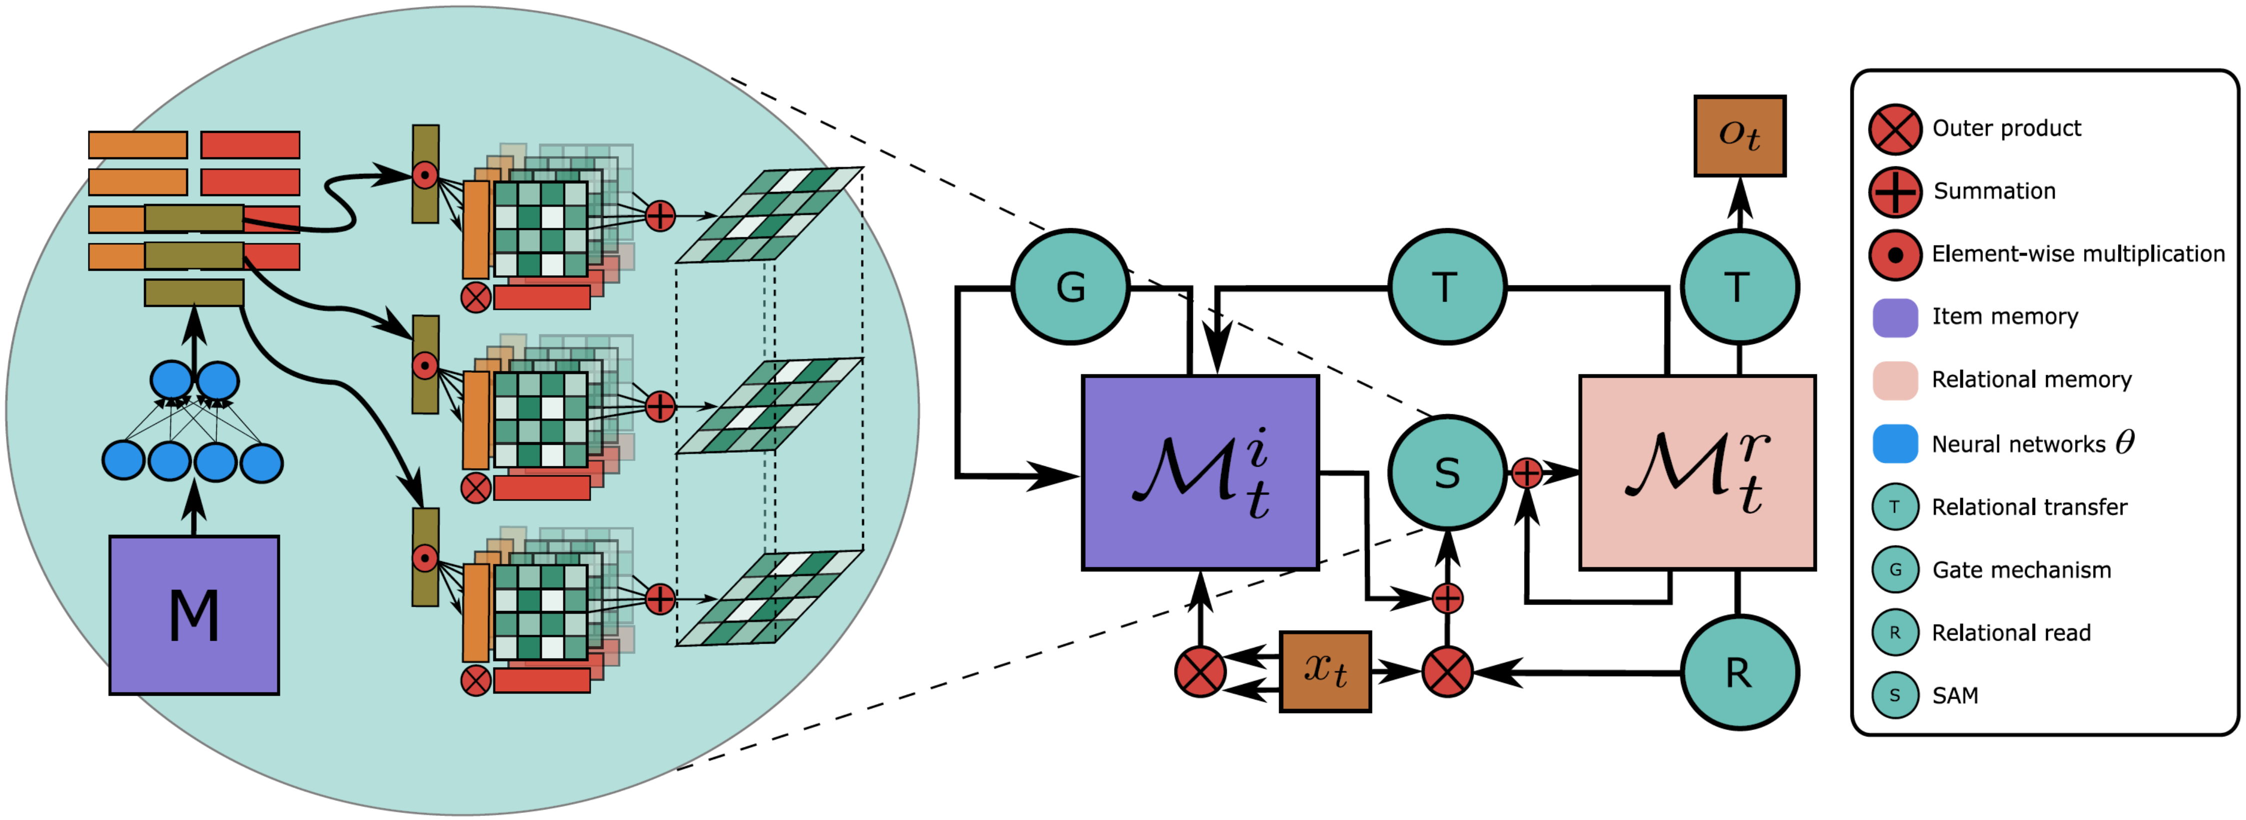

Hung Le, Truyen Tran, Svetha Venkatesh, "Self-attentive associative memory", *ICML'20*.

Memory of Programs in Neural Universal Turing Machine

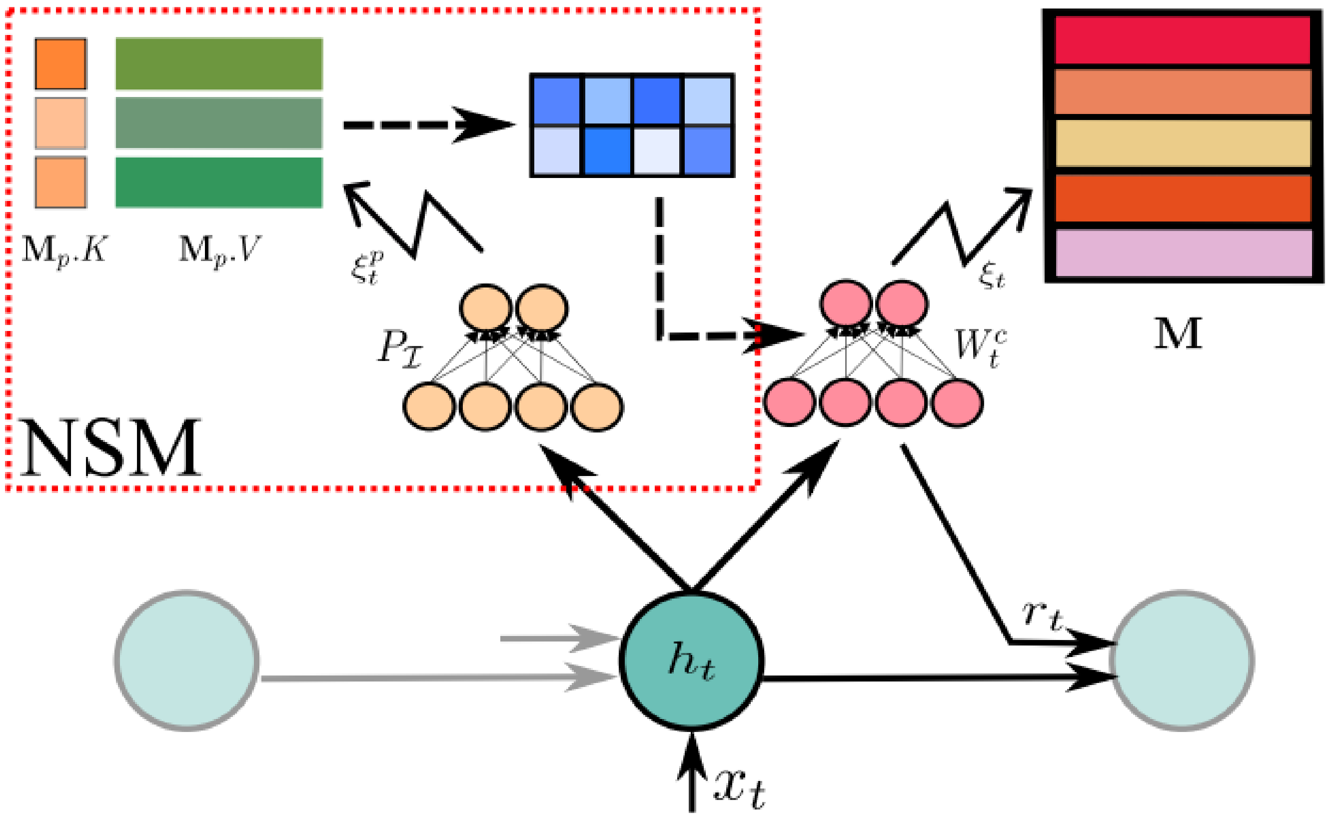

Figure 1: Introducing NSM into MANN. At each timestep, the program interface ($P_{\mathcal{I}}$) receives input from the state network and queries the program memory $\mathbf{M}_p$, acquiring the working weight for the interface network (W_t^c). The interface network then operates on the data memory $\mathbf{M}$ as normal.

Attention priors with syntax

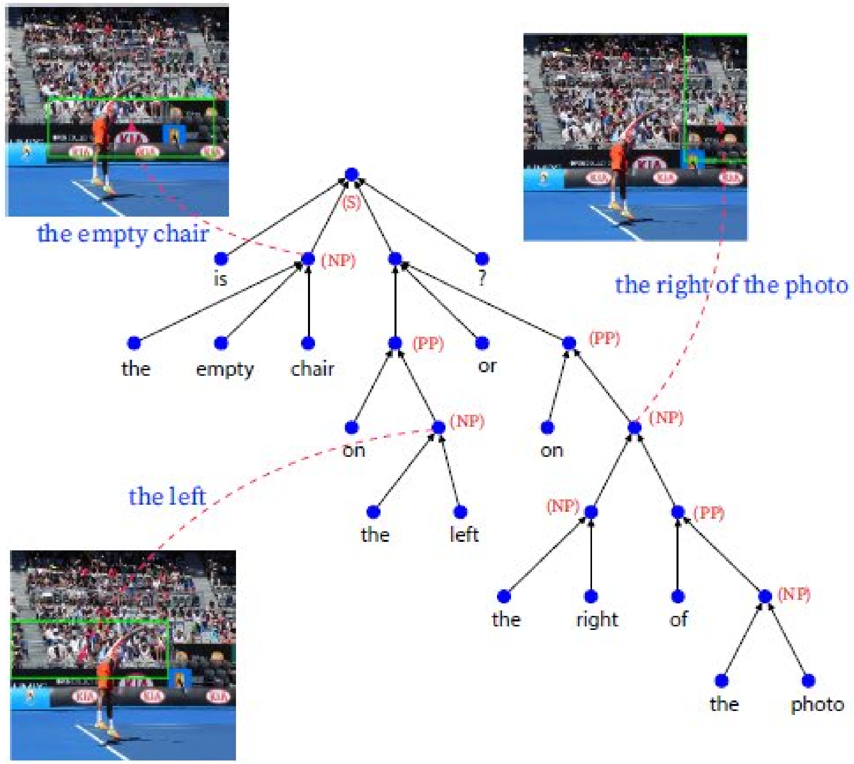

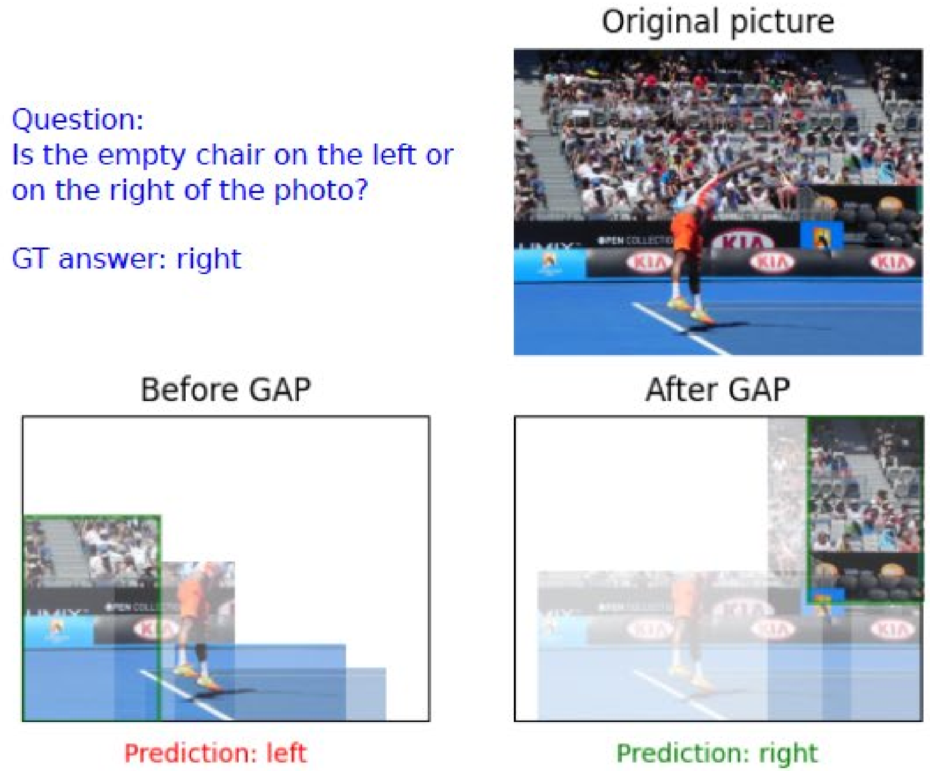

(Thao Le et al, *on going*)

A simple test: Separate reasoning process from perception

- Video QA: **inherent dynamic nature** of visual content over time.
- Recent success in visual reasoning with **multi-step inference** and handling of **compositionality**.

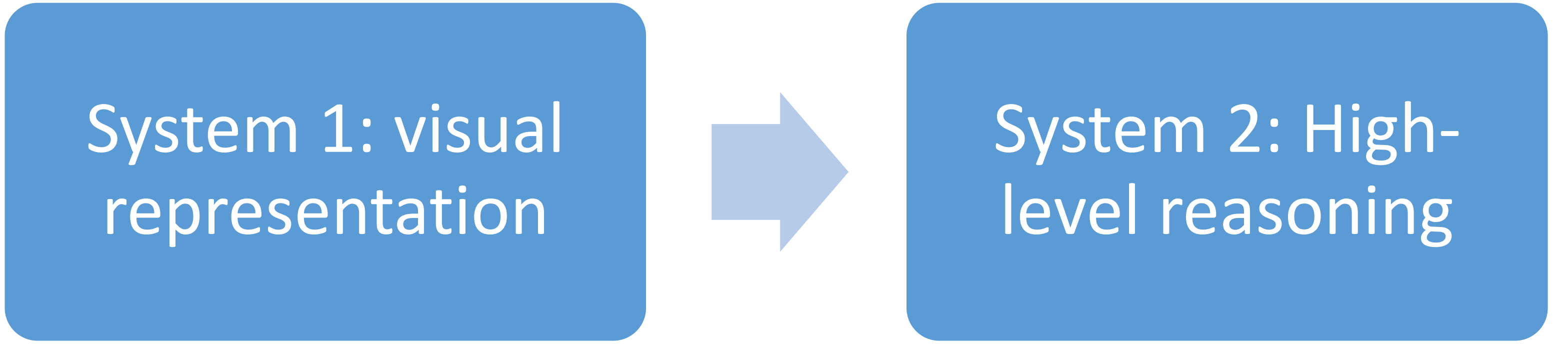

Le, Thao Minh, Vuong Le, Svetha Venkatesh, and Truyen Tran. "Neural Reasoning, Fast and Slow, for Video Question Answering." *IJCNN* (2020).

Separate reasoning process from perception (2)

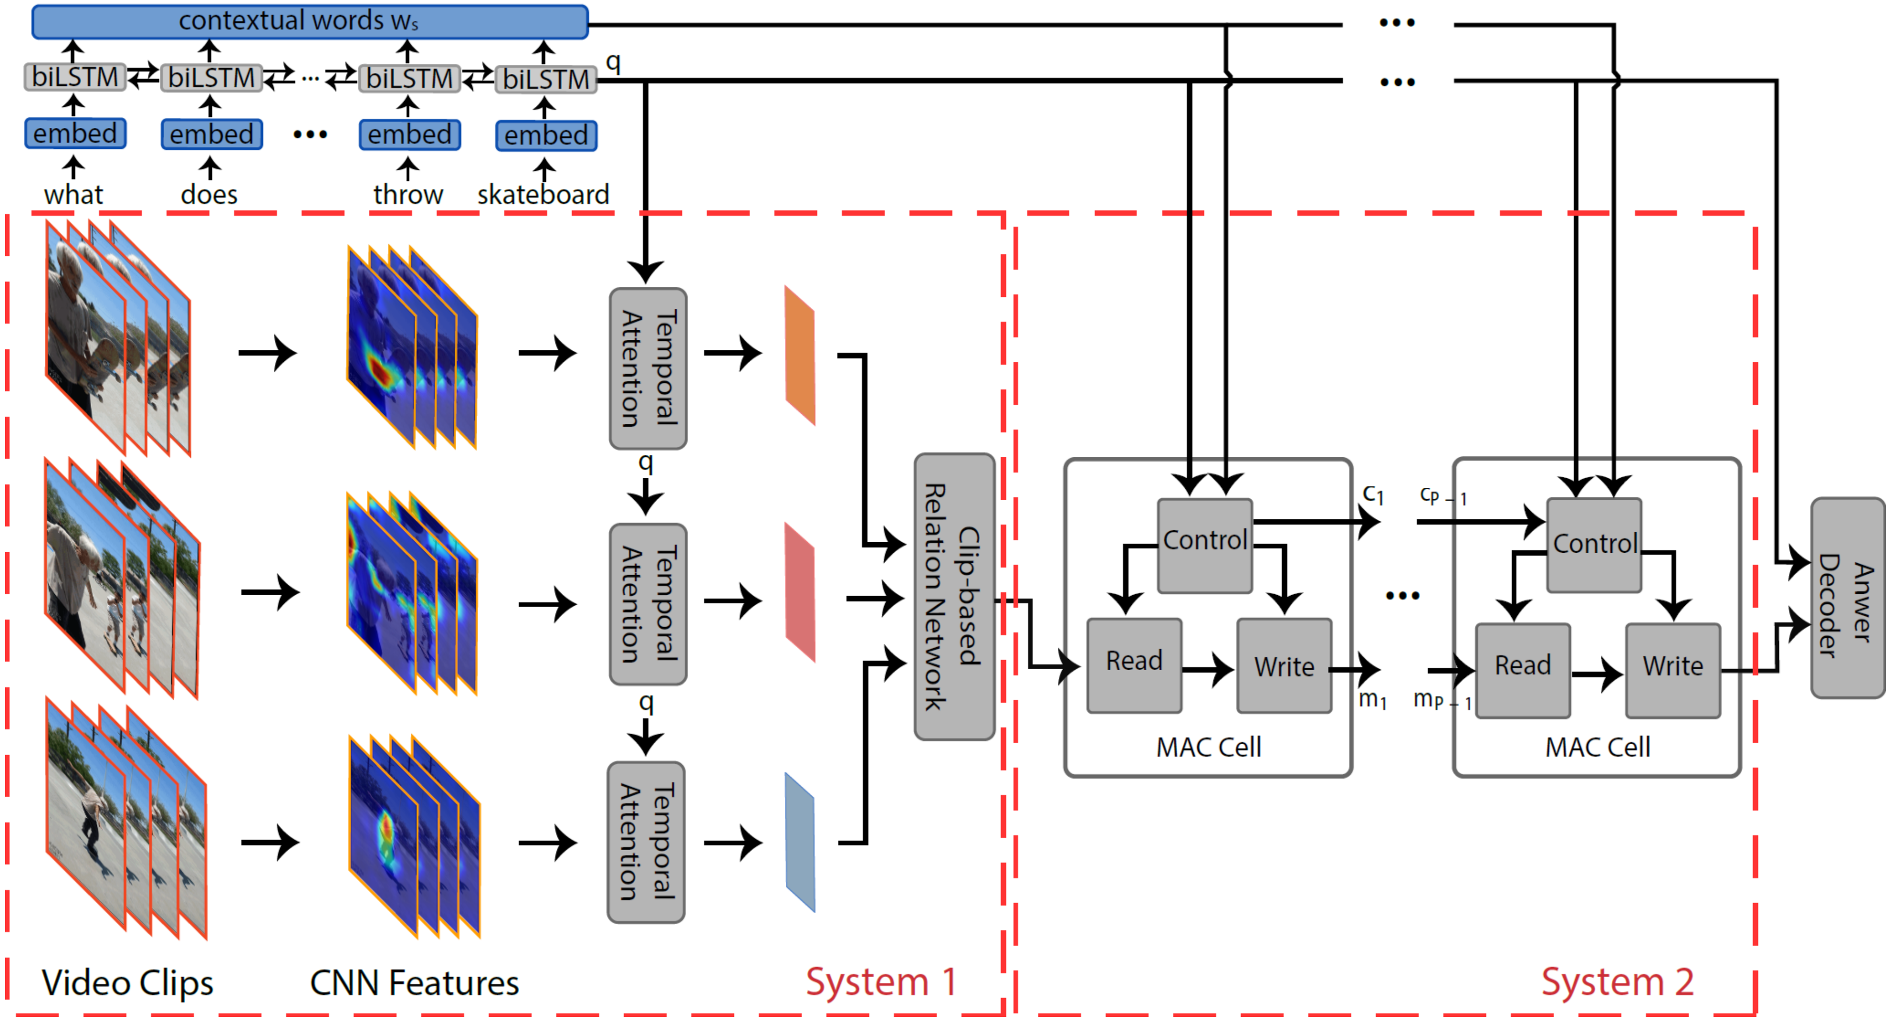

System 1: Clip-based Relation Network

Why temporal relations?

- Situate an event/action in relation to events/actions in the past and formulate hypotheses on future events.
- Long-range sequential modeling.

$$R^{(k)}(C) = h_{\Phi}\left(\sum_{i_1<i_2 \ldots <i_k} g_{\theta}\left(\bar{C}_{i_1}, \bar{C}_{i_2}, \ldots, \bar{C}_{i_k}\right)\right)$$

For $k = 2,3,\ldots,K$ where h_ϕ and g_θ are linear transformations with parameters ϕ and θ, respectively, for feature fusion.

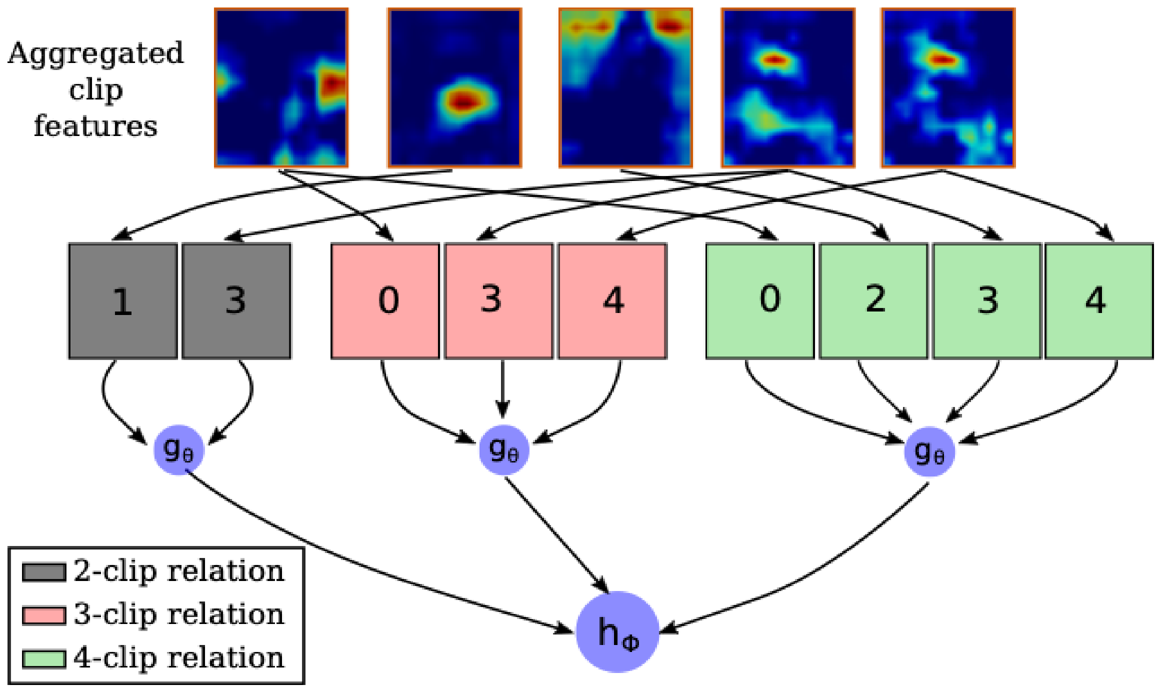

System 2 Candidate: MAC Net

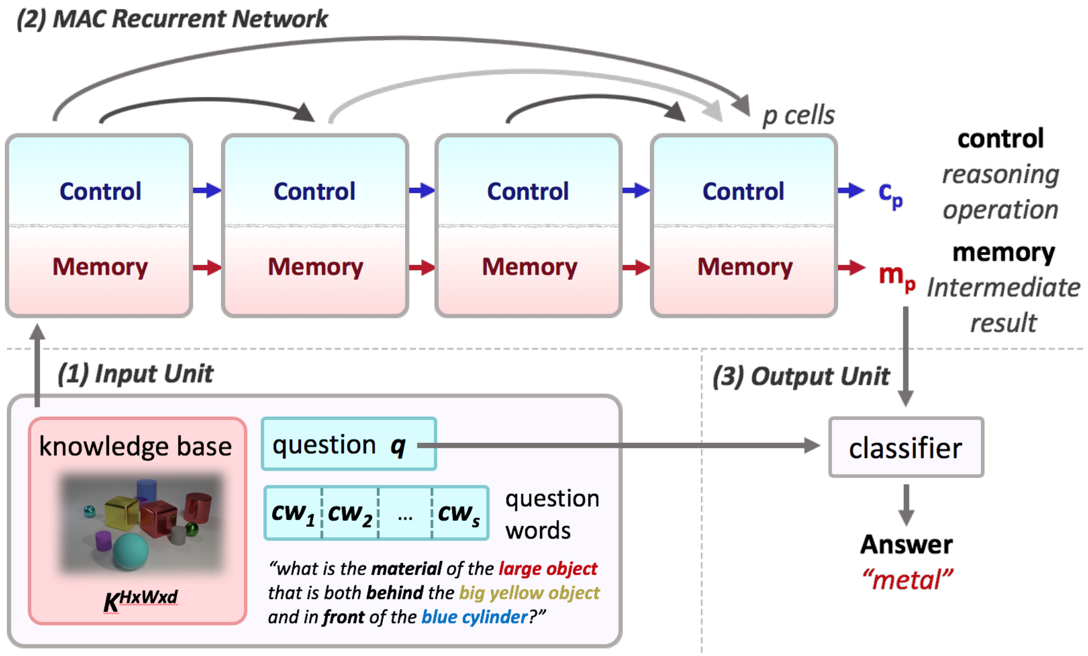

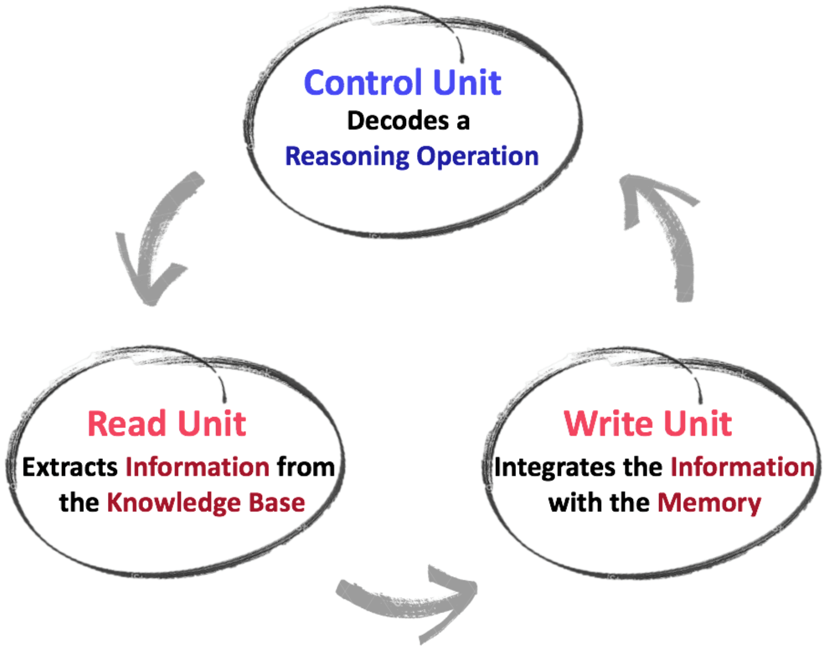

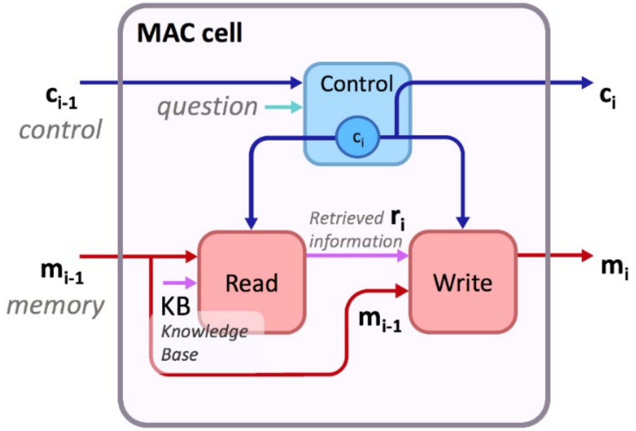

Hudson, Drew A., and Christopher D. Manning. "Compositional attention networks for machine reasoning." *ICLR* 2018.

Lecture 2: Sub-topics

- Problem with current DL
- Dual system theories
- Existing ideas for System 2
- **Theory of mind**

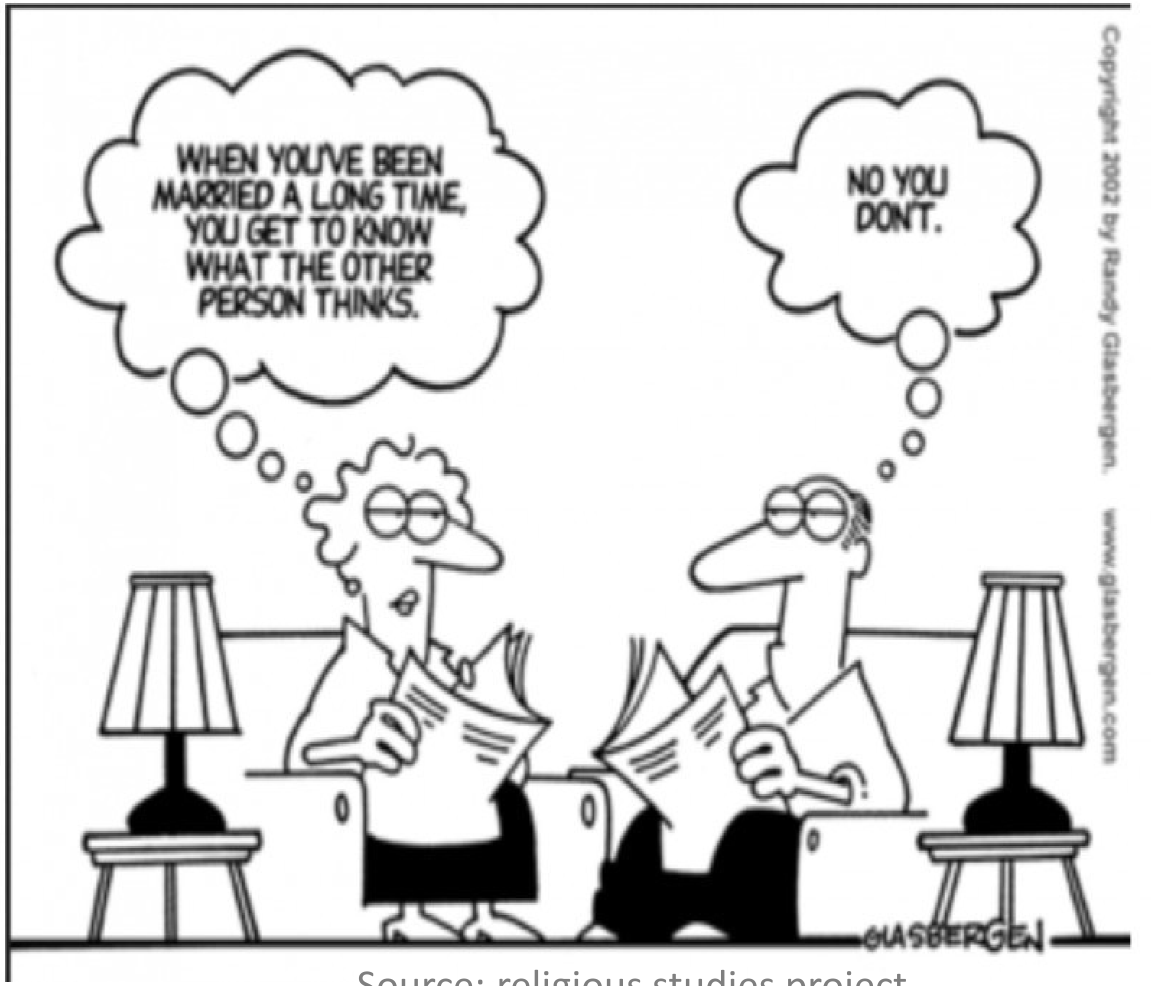

Source: religious studies project

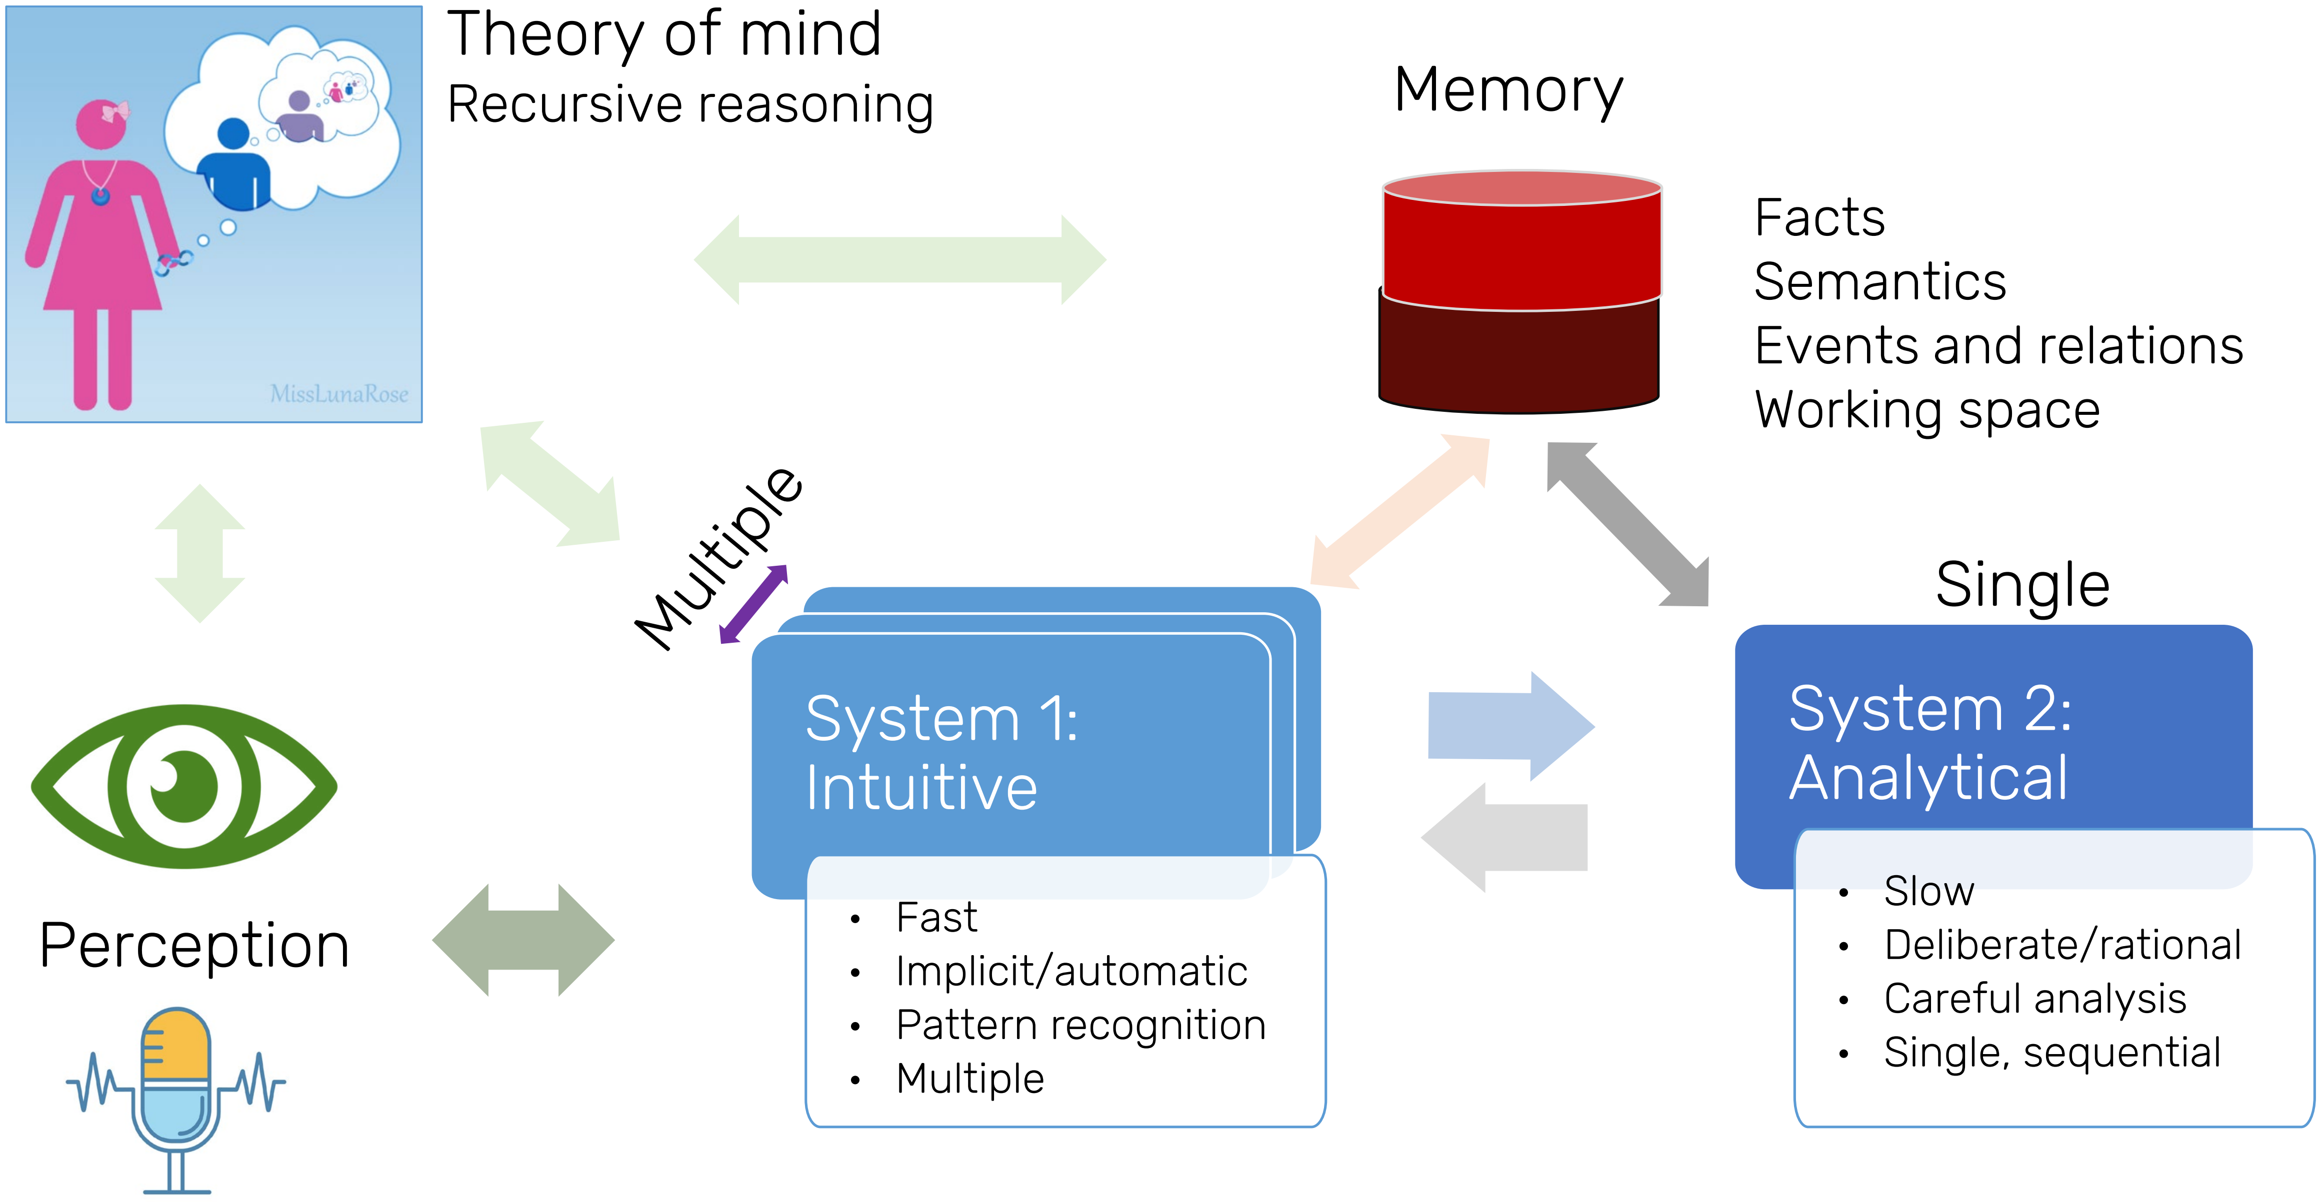

Where would ToM fit in?

Image credit: VectorStock | Wikimedia

Contextualized recursive reasoning

- Thus far, QA tasks are straightforward and objective:
  - Questioner: I will ask about what I don’t know.
  - Answerer: I will answer what I know.
- Real life can be tricky, more subjective:
  - Questioner: I will ask only questions I think they can answer.
  - Answerer 1: This is what I think they want from an answer.
  - Answerer 2: I will answer only what I think they think I can.

→ **We need Theory of Mind to function socially.**

Social dilemma: Stag Hunt games

- **Difficult decision**: individual outcomes (selfish) or group outcomes (cooperative).
  - Together hunt Stag (both are cooperative): Both have more meat.
  - Solely hunt Hare (both are selfish): Both have less meat.
  - One hunts Stag (cooperative), other hunts Hare (selfish): Only one hunts hare has meat.
- **Human evidence**: Self-interested but considerate of others (cultures vary).
- **Idea**: Belief-based guilt-aversion
  - One experiences loss if it lets other down.
  - Necessitates Theory of Mind: reasoning about other's mind.

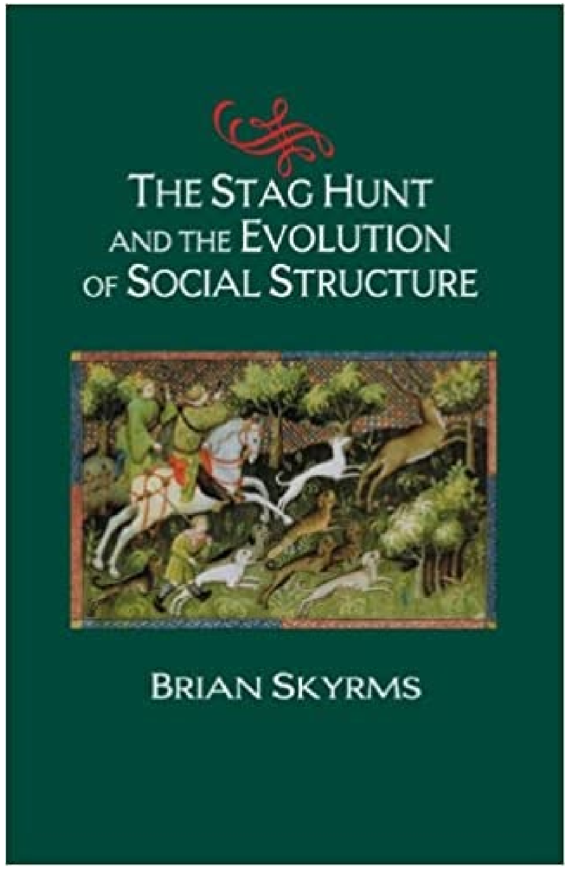

Theory of Mind Agent with Guilt Aversion (ToMAGA)

Update Theory of Mind

- Predict whether other's behaviour are cooperative or uncooperative
- Updated the zero-order belief (what other will do)
- Update the first-order belief (what other think about me)

Guilt Aversion

- Compute *the expected material reward* of other based on Theory of Mind
- Compute *the psychological rewards*, i.e. "feeling guilty"
- Reward shaping: subtract the expected loss of the other.

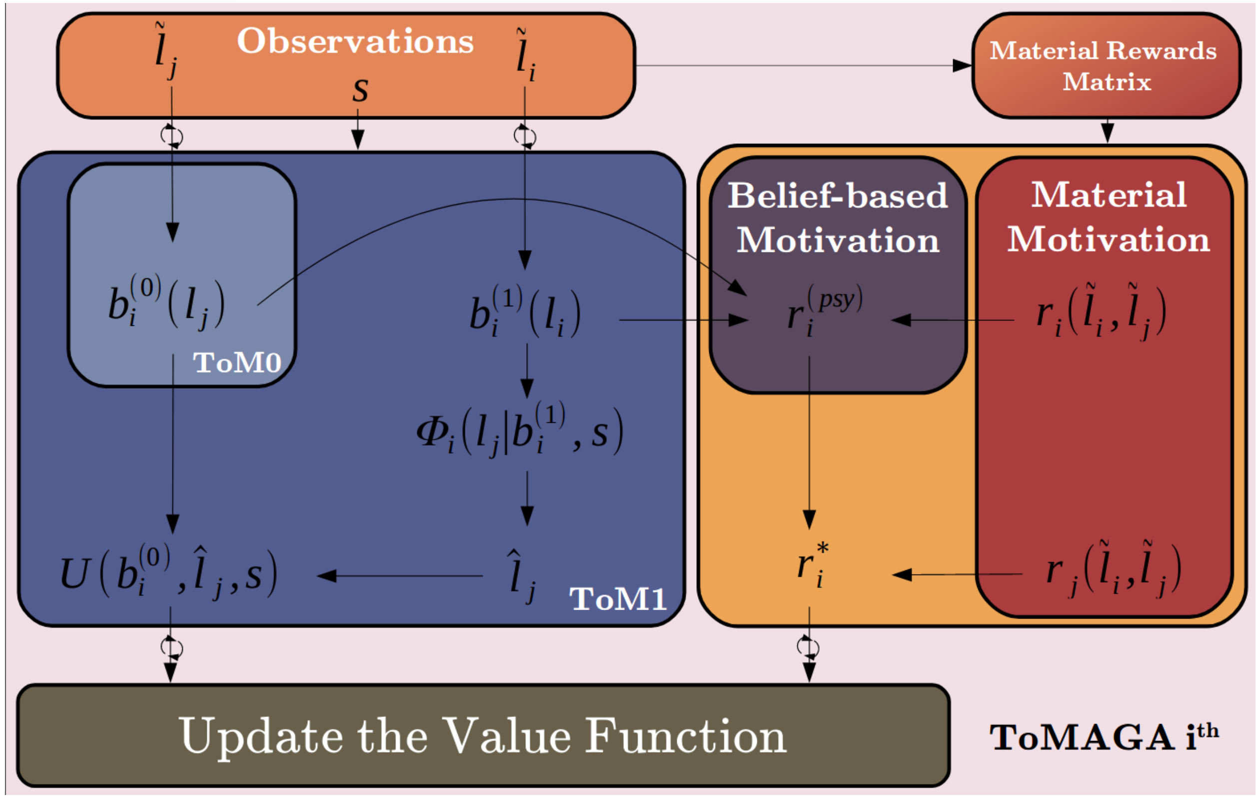

Nguyen, Dung, et al. "Theory of Mind with Guilt Aversion Facilitates Cooperative Reinforcement Learning." *Asian Conference on Machine Learning*. PMLR, 2020.

[Slide credit: Dung Nguyen]

Rabinowitz, Neil, et al. "Machine theory of mind." *International conference on machine learning*. PMLR, 2018.

Machine ToM Architecture (inside the Observer)

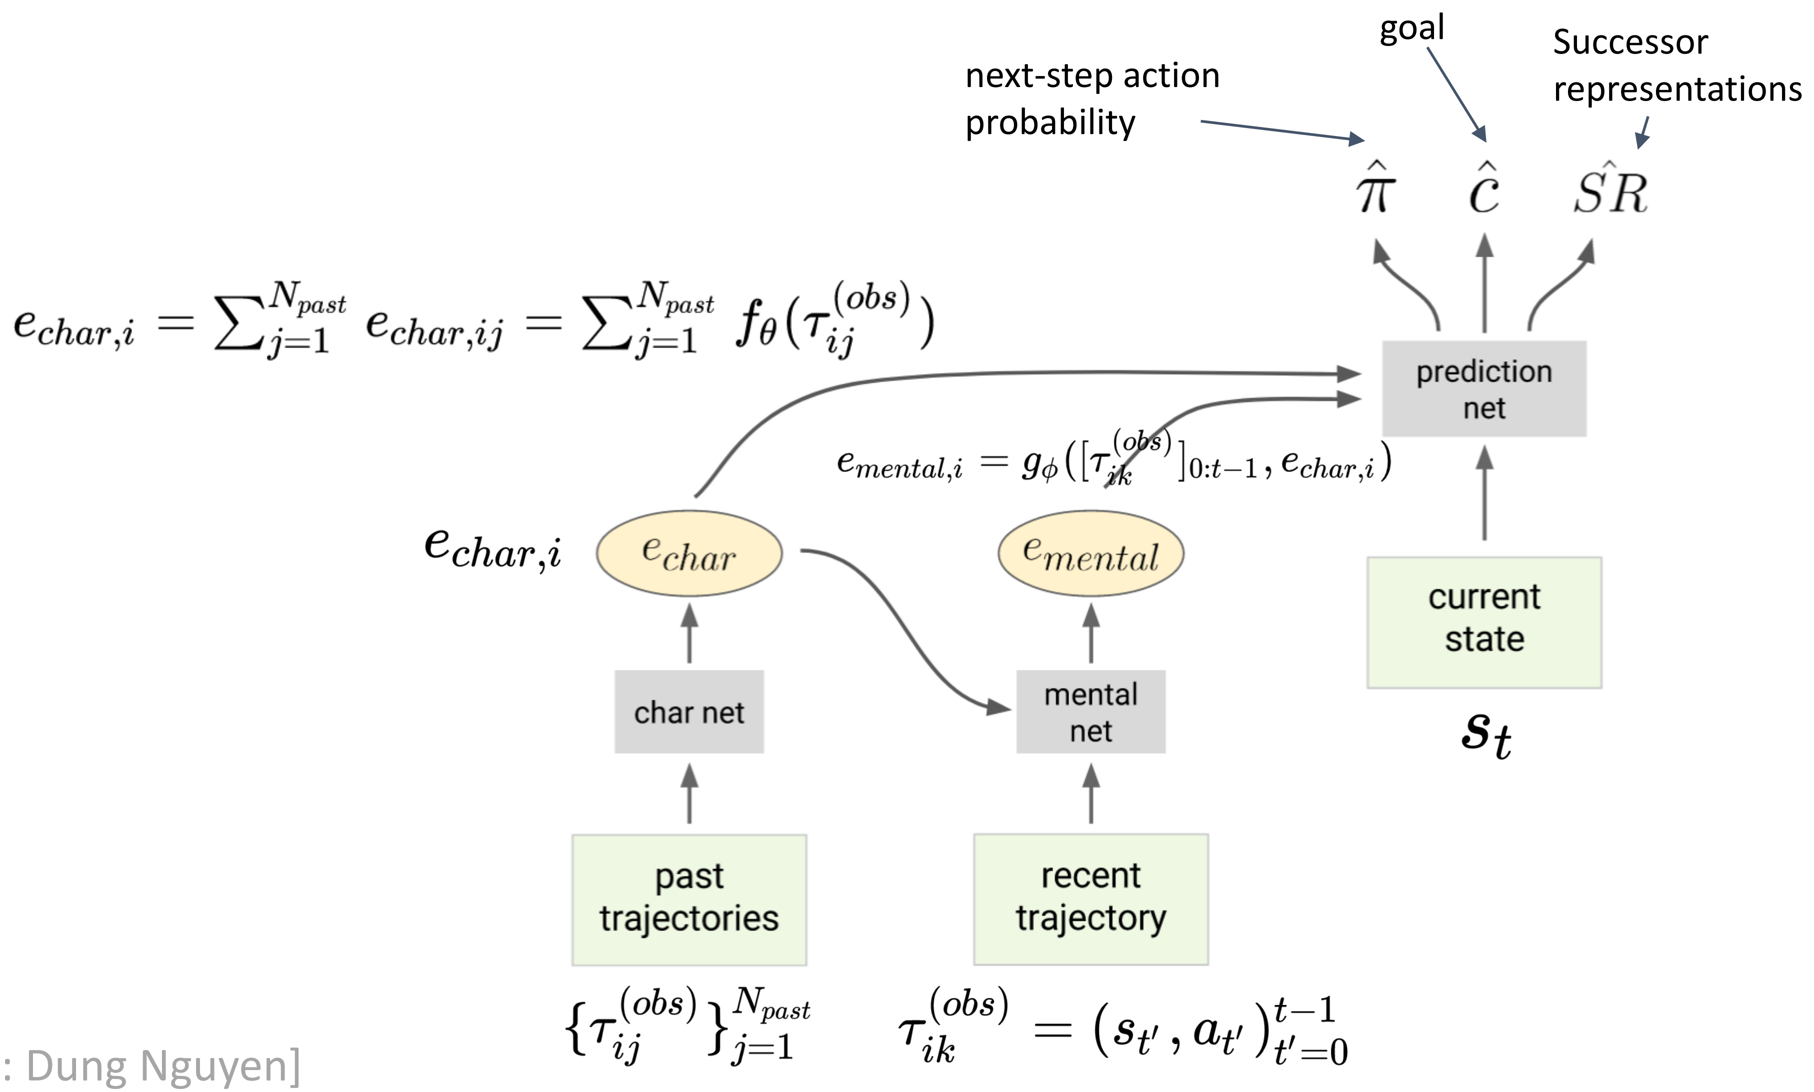

[Slide credit: Dung Nguyen]

A ToM architecture

- Observer maintains memory of previous episodes of the agent.
- It theorizes the “traits” of the agent.
  - Implemented as Hyper Networks.
- Given the current episode, the observer tries to infer goal, intention, action, etc of the agent.
  - Implemented as memory retrieval through attention mechanisms.

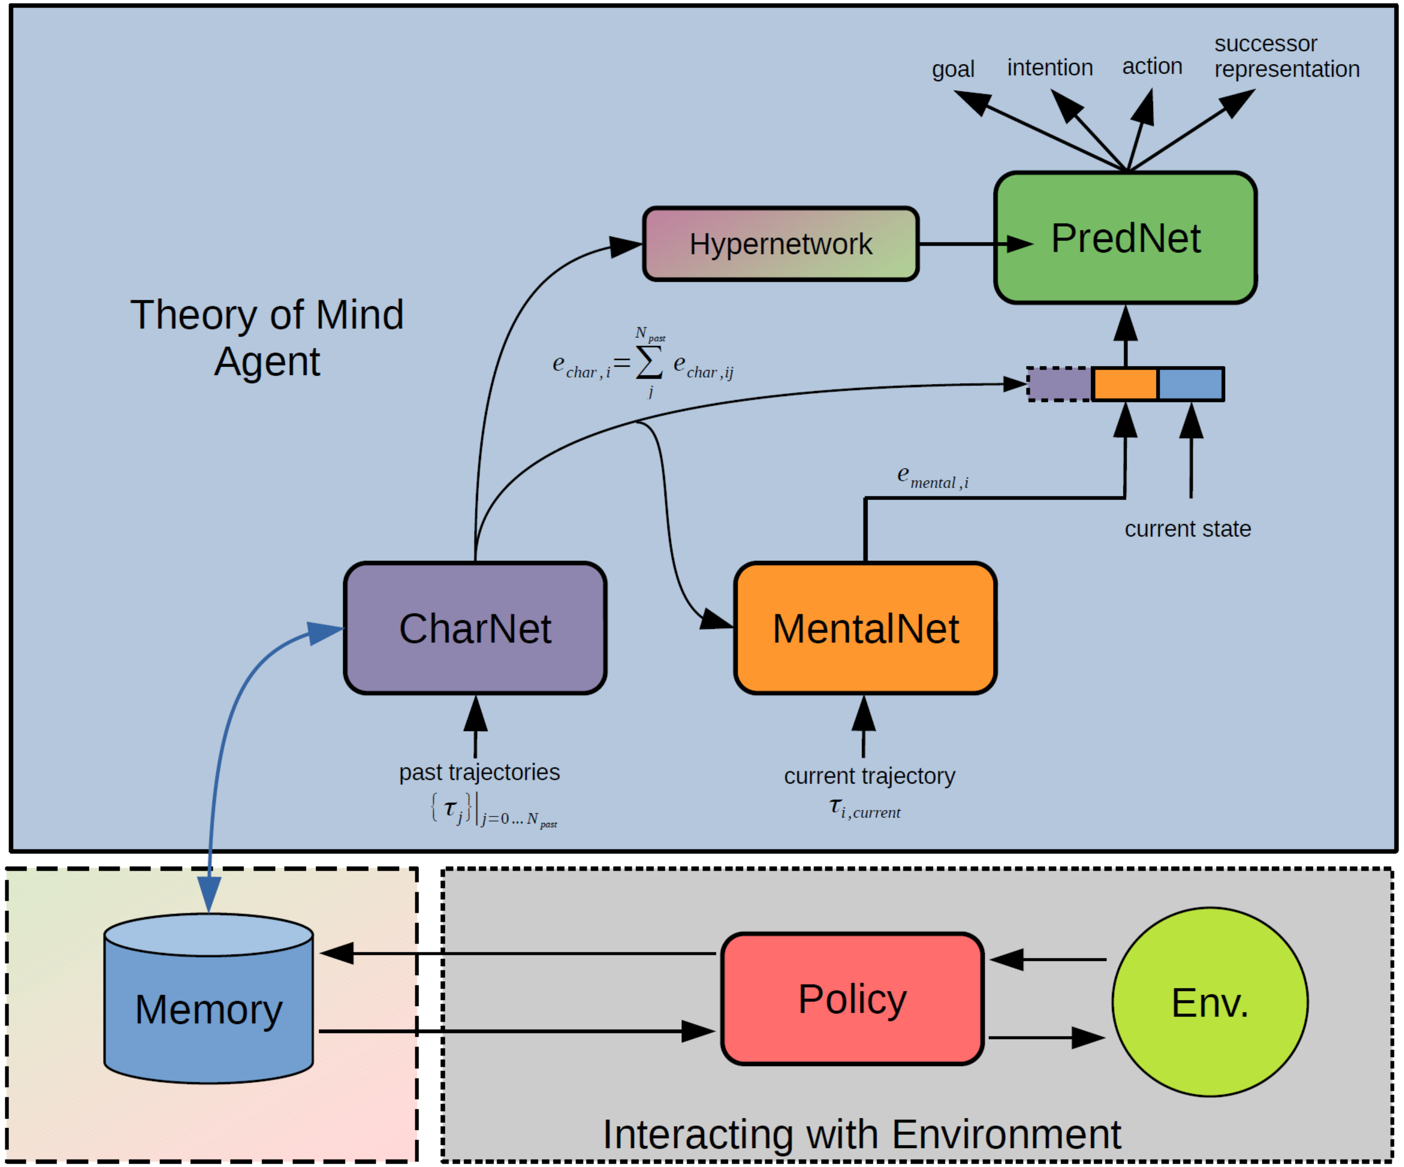

End of Lecture 2

https://neuralreasoning.github.io