OONP Appendix

A Reader

We show in Figure 1 a particular instance of Reader, including the components and the information flow among them. This particular Reader has three symbolic processors (namely, Symbolic Matching, Symbolic Reasoner, Symbolic Analyzer) and a Neural Net Controller (with Policy-net as the subcomponent). All the components in Reader are coupled through intensive exchange of information as shown in Figure 1. Below is a snapshot of the information processing at time t in Reader

- **STEP-1:** let the processor Symbolic Analyzer to check the Action History ($\mathsf{M}^t_{\text{act}}$) to construct some symbolic features for the trajectory of actions;
- **STEP-2:** access Matrix Memory ($\mathsf{M}^t_{\text{mat}}$) to get an vectorial representation for time t, denoted as $\mathbf{s}_t$;
- **STEP-3:** access Inline Memory ($\mathsf{M}^t_{\text{inl}}$) to get the symbolic representation $\mathbf{x}_t^{(s)}$ (through location-based addressing) and distributed representation $\mathbf{x}_t^{(d)}$ (through location-based addressing and/or content-based addressing);
- **STEP-4:** feed $\mathbf{x}_t^{(d)}$ and the embedding of $\mathbf{x}_t^{(s)}$ to Neural Net Controller to fuse with $\mathbf{s}_t$;
- **STEP-5:** get the candidate objects (some may have been eliminated by $\mathbf{x}_t^{(s)}$) and let them meet $\mathbf{x}_t^{(d)}$ through the processor Symbolic Matching for the matching of them on symbolic aspect;
- **STEP-6:** get the candidate objects (some may have been eliminated by $\mathbf{x}_t^{(s)}$) and let them meet the result of STEP-4 in Neural Net Controller ;
- **STEP-7:** Policy-net combines the result of STEP-6 and STEP-5, to issue actions;
- **STEP-8:** update $\mathsf{M}^t_{\text{obj}}$, $\mathsf{M}^t_{\text{mat}}$ and $\mathsf{M}^t_{\text{inl}}$ with actions on both symbolic and distributed representations;
- **STEP-9:** put $\mathsf{M}^t_{\text{obj}}$ through the processor Symbolic Reasoner for some high-level reasoning and logic consistency.

Note that we consider only single action for simplicity, while in practice it is common to have multiple actions at one time step, which requires a slightly more complicated design of the policy as well as the processing pipeline.

B Experiments: Logical consistency

Suppose at time t, the ontology in $\mathsf{M}^t_{\text{obj}}$ contains the following three facts (among others)

- fact-1: `John` (a PERSON-object) is in `kichten` (a LOCATION-object);
- fact-2: `John` carries `apple` (an ITEM-object);
- fact-3: `John` drops `apple`;

where fact-3 is just established by Policy-net at t. Symbolic Reasoner will add a new $\texttt{is-located-at}_B$ link between `apple` and `kitchen` based on domain logic*.

*The logic says, an item is not "in" a location if it is held by a person.

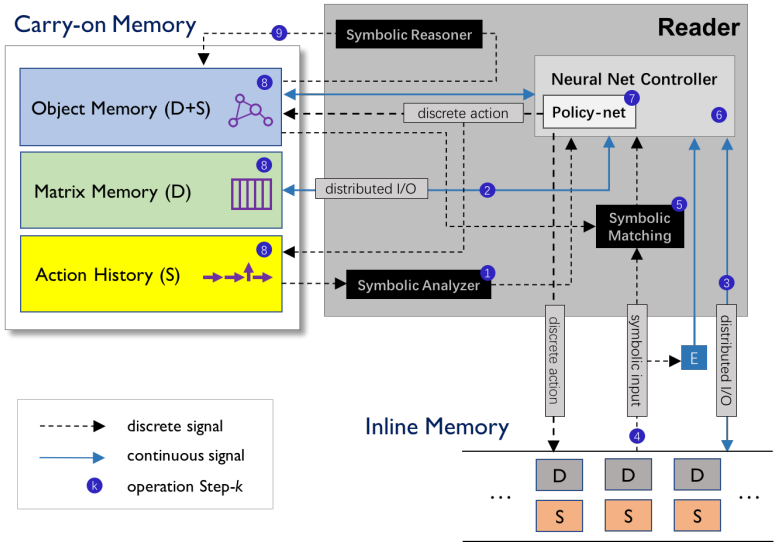

Figure 1: A particular implementation of Reader in a closer look, which reveals some details about the entanglement of neural and symbolic components. Dashed lines stand for continuous signal and the solid lines for discrete signal