Machine Reading as Model Construction

Peter Clark
Allen Institute for AI (AI2)
Seattle, WA
peterc@allenai.org

1 WHAT IS MACHINE READING?

With the advent of large datasets of paragraphs + questions, e.g., SQuAD [4], TriviaQA [3], there has been renewed interest in general-purpose "reading comprehension" (RC) systems, capable of answering questions against those paragraphs, e.g., [5, 6]. These systems have become remarkably effective at factoid QA. However, they require extensive training data, and can still struggle with queries requiring complex inference [1]. The extent to which these systems have truely read and understood the paragraph remains unclear [2].

At the other end of the spectrum, AI has also developed sophisticated formalisms for modeling the world, e.g., situation calculus, event calculus, qualitative modeling. These frameworks allow systems to represent facts which are known, and infer facts which are unknown. Models built with these frameworks constitute an understanding of the world, in that sense that they are predictive: If the model's computational clockwork moves in a way similar to the world, then the model can predict how the world will behave, constituting a degree of understanding of the world. In this context, machine reading can be viewed as the task of constructing such models from text, given a particular modeling framework in which to express those models.

While it is possible that a neural system might eventually be able to infer a predictive, neural model of the world solely from large numbers of examples, we do not believe this is likely in the near future. Rather, we see the way forward as combining the pattern-learning techniques of neural systems with the modeling capabilities of structured representations. AI modeling frameworks provide a set of primitives for constructing predictive models, and neural systems can help construct models within those frameworks that best fit data. The grand challenge for machine reading, going forward, is combining these two technologies together to do this.

2 MACHINE READING ABOUT PROCESSES

At AI2 we have been pursuing a specific genre of machine reading along these lines, namely reading paragraphs describing processes (e.g., photosynthesis). Our goal is not to simply answer lookup questions, but also answer questions that go beyond the text, in particular about the states that exist during a process. Such questions are challenging because those world states are often implicit, making questions hard to answer from surface cues alone.

For example, consider the following paragraph about photosynthesis:

> *Chloroplasts in the leaf of the plant trap light from the sun. The roots absorb water and minerals from the soil. This combination of water and minerals flows from the stem into the leaf. Carbon dioxide enters the leaf. Light, water and minerals, and the carbon dioxide all combine into a mixture. This mixture forms sugar (glucose) which is what the plant eats.*

While reading comprehension (RC) systems can reliably answer lookup questions such as:

(1) What do the roots absorb? (A:water, minerals)

they struggle when answers are not explicit, e.g.,

(2) Where is sugar produced? (A:in the leaf)

For example, the RC system BiDAF [5] answers "glucose" to this second question. This question requires knowledge and inference: If carbon dioxide *enters* the leaf (stated), then it will be *at* the leaf (unstated), and as it is then used to produce sugar, the sugar production will be at the leaf too. This is the kind of inference that our system, ProComp ("process comprehension"), is able to model, using a structured representation of events and states.

Our approach is illustrated in Figure 1, and we briefly summarize it here. First, ProComp extracts a Process Graph from the paragraph, representing the event sequence in the process. It then performs a STRIPS-like simulation of the process, using a set of precondition/effect rules about events, mined from VerbNet. Finally, a small set of answer procedures operate over that simulation, allowing several classes of questions about change to be answered (e.g., "Where is X at step Y?", "What entities change size during the process?"). Although our initial work has used largely traditional techniques, it is still able to outperform RC systems on questions about change, and thus illustrates the importance of modeling in machine reading.

3 INTEGRATING NEURAL METHODS

Our initial system uses three basic operations:

- (Event extraction) Given a sentence describing an event, identify the event and the participants within it.
- (State prediction) Given a sentence describing an event, and an entity mentioned in the sentence, predict the state of the entity before/after the event (where the state of the entity is a set of properties associated with it, selected from a predefined set).
- (State inference) Given a partial description of the entities and their states during the process (i.e., a partially filled Participant Grid), fill in the remaining states.

To date, we have collected a large number of hand-annotated examples of these predictions to evaluate our system ProComp. However, clearly this data can also be used for learning, to train a system to make these inferences. Note that this does not obviate the need for ontology design - the appropriate dimensions of modeling still need to be selected. However, it does offer an example-based means

K-CAP2017 Workshops and Tutorials Proceedings, 2017

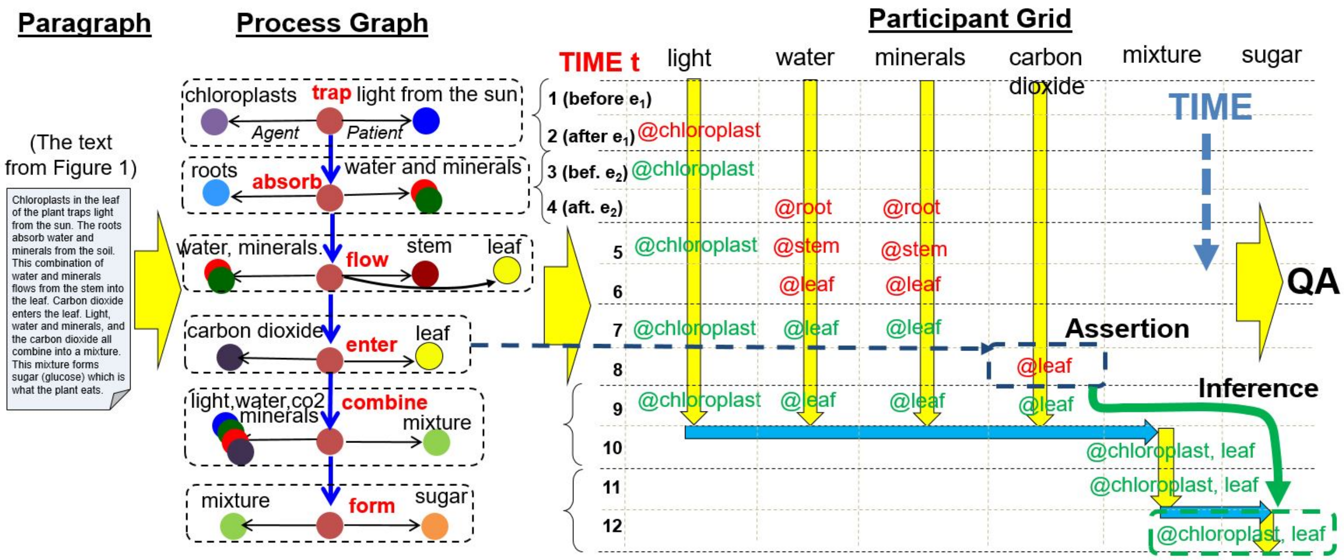

Figure 1: An illustration of machine reading as model construction and inference: Here the constructed model is a process graph (sequence of events), and inference is state-space simulation, presented graphically as a "Participant Grid". Each row in the Grid is a state (time vertically downwards), each column is a process participant, and each cell shows facts true of a participant in a state. For brevity, @ denotes is-at(), yellow lines denotes exists(), red denotes a direct consequence of an event, green an inferred consequence. For example, at line 8 in the Grid (labelled "Assertion"), the "CO2 enters leaf" step asserts that CO2 is therefore @leaf after the event. By inference, the sugar must therefore be produced at the leaf too (green box), a fact not explicitly stated in the text.

for connecting that ontology and the reasoning to data. This is an exciting direction we are pursuing.

4 SUMMARY

Unlike much recent work, we view machine reading as the task of constructing a model from text using a particular modeling framework. The framework provides the building blocks for modeling a certain class of phenomena, and the task of reading is to construct a model within that framework. We have illustrated this for reading text about processes, using a state-based modeling framework.

There is a symbiotic relationship between text and modeling frameworks:

- Text suggests which modeling framework is appropriate (e.g., the text appears to be describing a process, so use a framework suitable for processes)
- The modeling framework provides expectations about what to look for in the text (e.g., given it's a process, expect to see events and their participants)

This approach does not remove the need for learning, rather it provides a scaffolding within which learning can take place, and a mechanism for then supporting inference and prediction - activities that truly demonstrate that the machine has understood what it has read.

REFERENCES

[1] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*. 1693–1701.

[2] Robin Jia and Percy Liang. 2017. Adversarial Examples for Evaluating Reading Comprehension Systems. In *Proc. EMNLP'17*.

[3] Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *Proc. ACL'17*.

[4] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proc. EMNLP'16*.

[5] Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional Attention Flow for Machine Comprehension. In *Proc. ICLR'17*.

[6] Junbei Zhang, Xiaodan Zhu, Qian Chen, Lirong Dai, and Hui Jiang. 2017. Exploring Question Understanding and Adaptation in Neural-Network-Based Question Answering. *arXiv preprint arXiv:1703.04617* (2017).