

Question Answering in a Graphical Cognitive Architecture

Sanjay Raveendran, Paul S. Rosenbloom and Anton Leuski

Department of Computer Science/Institute for Creative Technologies, University of Southern California
12015 Waterfront Dr., Playa Vista, CA 90094, Ph: (310) 448-5341, web: <http://www-scf.usc.edu/~sraveend/>
sraveend@usc.edu, rosenbloom@usc.edu, leuski@ict.usc.edu

Cognitive Architectures

A cognitive architecture embodies a hypothesis concerning the fixed structures underlying intelligent behavior, whether in natural or artificial systems. One of their core applications is building lifelike characters, called virtual humans, capable of understanding and generating speech and natural language, planning, reasoning, inference and learning. Some of these tasks are primarily symbolic (e.g., planning, reasoning), others are currently best approached statistically (e.g., learning and natural language), and others are primarily signal oriented (e.g., speech). This is reflected in the fact that the virtual humans built by institutes like the University of Southern California Institute for Creative Technologies (ICT) use two distinct types of cognitive architectures – symbolic and statistical – depending on the nature of the character to be created, with separate signal processing modules also included for more peripheral capabilities, like speech (Kenny et al. 2007).

NPCEditor is one such statistical architecture that is used for virtual humans capable of understanding questions within a specific domain and responding with an appropriate answer (Leuski and Traum 2010a). Using NPCEditor, a character designer can set up one or more characters, along with their responses on a particular topic and sample questions linked to those responses. When a user's question is received by the system, it uses a probabilistic text classifier to rank the responses in order of appropriateness and returns the best match. It is successfully being used in a variety of applications, such as virtual museum guides, virtual patients for training clinicians, and cultural trainers for soldiers in the U.S. Army (Leuski and Traum 2010b).

While the system is capable of some reasoning by building semantic frames based on its model of the user and

itself (Leuski and Traum 2008), it is not built for general-purpose cognition. Conversely, general purpose symbolic architectures such as Soar (Laird 2008) cannot, as yet, effectively perform tasks such as statistical classification. However, by leveraging the ability of graphical models (Koller and Friedman 2009) to produce and integrate state-of-the-art capabilities for symbol, probability and signal processing, it is possible to develop new forms of hybrid (discrete and continuous) mixed (relational and probabilistic) architectures that show the potential for bringing all of these threads together (Rosenbloom, 2011a).

We are exploring the possibility of cleanly and efficiently embedding the text classification functionality of NPCEditor into a unified graphical architecture. This has the potential to extend the functionality of conversational virtual humans while retaining the strengths of robust statistical processing and simple creation. This could also pave the way for a blending of ICT's two distinct classes of architectures for virtual humans.

From Semantic Memory to Text Classification

The graphical cognitive architecture under development is eventually intended to yield a hybrid mixed variant of Soar. In its present state, it combines a broad memory capability (Rosenbloom 2010) with the ability to make decisions (Rosenbloom 2011b). The memory capability combines a hybrid working memory with a hybrid mixed long-term memory capable of supporting, via a single general factor graph and the summary product algorithm (Kschischang, Frey and Loeliger 2001), the kinds of procedural rule-based memories and declarative semantic and episodic memories found in state-of-the-art architectures such as Soar and ACT-R (Anderson et al. 2004). At the core of the reimplemention of the NPCEditor is the semantic memory, which when cued by a subset of object features predicts the object's categorization and any missing feature values via a naive Bayes model that combines prior proba-

bilities on the categories with conditional probabilities of feature values given the categorization.

The text classifier in NPCEditor views question answering as an information retrieval task. The system models each text utterance (a user’s question or a system’s response) as a probability distribution over individual words – a language model. It estimates a language model of a system’s answer given a user’s question

$$P(a|Q) = \frac{\sum_s \pi_{A_s}(a) \prod_{i=1}^m \pi_{Q_s}(q_i)}{\sum_s \prod_{i=1}^m \pi_{Q_s}(q_i)}, \quad (1)$$

compares it to language models of individual responses in the system’s database using Kullback-Liebler divergence

$$D(P(Q)||P(A)) = \sum_a P(a|Q) \log \frac{P(a|Q)}{\pi_A(a)}, \quad (2)$$

and returns the best match (Leuski and Traum, 2010b). In these equations, Q is the user’s question, A is the system’s response, a is an answer word, m is the number of words in the question, s iterates over all question-response pairs $\{Q_s, A_s\}$ in the system’s database, and π_x is a language model of utterance x . The latter is normally computed using maximum-likelihood estimation over words.

Ignoring the normalization constant, we can rewrite equation (1) as the following multivariate function:

$$P(a, Q) = \sum_s \varphi(a, s) \prod_i \phi(s, q_i) \quad (3)$$

We can view this as the posterior probability for a category (in this case the answer word) given the prior (φ) and the conditional probabilities (ϕ), summed out for s . Since the values of ϕ and φ will not change for a given question and answer pair (they are dependent only on the words in the questions and answers), we can calculate them in advance and encode them as conditional probabilities in factor functions. This directly maps to how the architecture handles semantic memory.

Our current efforts are focused on computing the Kullback-Liebler divergence of each response to the question – equation (2) – for which there is no direct analogue in the architecture. We are exploring a range of approaches that exploit pre-calculation, approximation, extensions to factor graphs and alternative comparisons. If this issue can be resolved, then the closest matching response to the user’s question can be determined in a manner similar to how decisions are implemented (Rosenbloom 2010).

Conclusion and Future

Results to date have shown how the semantic memory within our graphical cognitive architecture can subsume key portions of a state-of-the-art text classification engine

with minimal changes, and its potential to support the rest of it. ICT uses both statistical systems like NPCEditor and more general-purpose symbolic architectures such as Soar in building its virtual humans. Being able to achieve the functionality of both within a single system holds promise for a unified “broad spectrum” architecture that can address both these individual needs plus handle various blends of them. This could in turn greatly increase the capabilities of coming generations of virtual humans.

Acknowledgements

This effort has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred. We would like to thank David Traum for helpful comments on this work.

References

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. 2004. An integrated theory of the mind. *Psychological Review* 111, (4): 1036-1060.
- Kenny, P., Hartholt, A., Gratch, J., Swartout, W., Traum, D., Marsella, S., & Piepol, D. 2007. Building Interactive Virtual Humans for Training Environments. In *Interservice/Industry Training, Simulation and Education Conference (IITSEC)* (7105): 1-16.
- Koller, D. and Friedman, N. 2009. *Probabilistic Graphical Models: Principles and Techniques (Adaptive Computation and Machine Learning)*. The MIT Press.
- Kschischang, F. R., Frey, B. J., and Loeliger, H. 2001. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory* 47: 498-519.
- Laird, J. E. 2008. Extending the Soar cognitive architecture. In *Artificial General Intelligence 2008: Proceedings of the First AGI Conference*. IOS Press.
- Leuski, A., and D. Traum. 2008. A statistical approach for text processing in virtual humans. In *Proceedings of the 26th Army Science Conference*.
- Leuski, A., and D. Traum. 2010a. NPCEditor: A tool for building question-answering characters. In *Proceedings of The 7th International Conference on Language Resources and Evaluation*.
- Leuski, A., and D. Traum. 2010b. Practical language processing for virtual humans. In *22nd Annual Conference on Innovative Applications of Artificial Intelligence*.
- Rosenbloom, P. S. 2010. Rethinking cognitive architecture via graphical models, In *Cognitive Systems Research*. In Press.
- Rosenbloom, P. S. 2011a. Combining procedural and declarative knowledge in a graphical architecture. In *Proceedings of the 10th International Conference on Cognitive Modeling*
- Rosenbloom, P. S. 2011b. From memory to problem solving: Mechanism reuse in a graphical cognitive architecture. In preparation.