Hybrid ML/KB Systems Learning through NL Dialog with DL Models

Sergei Nirenburg[1,*,†], Nikhil Krishnaswamy[2,†] and Marjorie McShane[1,†]

[1]Rensselaer Polytechnic Institute, 110 8th St. Troy, NY 12180, USA
[2]Colorado State University, 1100 Center Avenue Mall, Fort Collins, CO 80523, USA

Keywords
Language-endowed intelligent agents, Learning through language understanding

Neurosymbolic approaches to artificial intelligence have so far been understood primarily as "a method of reincorporating some of the structure and speed of symbolic reasoning into the flexible representations provided by deep learning" [1, 2, 3]. Our approach is the inverse of the above: we use flexible representations, big data orientation and analogical reasoning of ML (and older-style data analytics) to a) support knowledge-based processing (for example, taking care of pre-semantic stages of deep language understanding []; and b) to support efforts on overcoming the notorious "AI knowledge bottleneck." In this abstract we (very) briefly describe our team's past, current and planned future work on the above topics.

We develop intelligent agent systems (language-endowed intelligent agents, LEIAs []) that include both ML-based and knowledge-based processing modules (see Figure 1). LEIAs are intended to serve as members of human-AI teams in critical applications where humans must fully trust the AI agents' conclusions and recommendations. Knowledge acquisition is a central concern in LEIA development [].

While originally the efforts concentrated on analytical support and ergonomics of manual acquisition, the availability of the LEIA system infrastructure offered a novel opportunity – knowledge acquisition as a side effect of the system's regular operation ("opportunistic learning"). This capability required several prerequisites:

- an NL understanding system that takes text as input and interprets it in terms of the underlying ontological model of the world, language and human behavior (in other words, extracts and represents the meaning of language inputs) [];
- an image interpretation system that uses same underlying knowledge base to interprets results of image recognition [];
- a conceptual learning module that extracts the information to be learned from the outputs of the above interpretation systems and updates the LEIA's knowledge infrastructure – the ontology, the lexicon, the opticon and the episodic memory of concept exemplars [].

In A. Martin, K. Hinkelmann, H.-G. Fill, A. Gerber, D. Lenat, R. Stolle, F. van Harmelen (Eds.), Proceedings of the AAAI 2023 Spring Symposium on Challenges Requiring the Combination of Machine Learning and Knowledge Engineering (AAAI-MAKE 2023), Hyatt Regency, San Francisco Airport, California, USA, March 27-29, 2023.
[*]Corresponding author.
[†]These authors contributed equally.
nirens@rpi.edu (S. Nirenburg); nkrishna@colostate.edu (N. Krishnaswamy); mcsham2@rpi.edu (M. McShane)

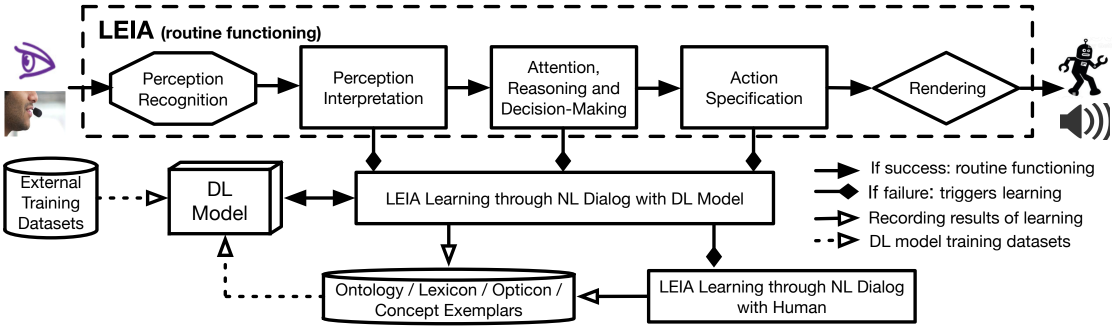

Figure 1: Integration of routine operation and two kinds of conceptual learning by LEIAs with DL model training involving LEIA knowledge resources. LEIAs are already hybrid systems, as perception recognition and much of rendering tasks are implemented using ML-based methods.

A LEIA with the above capabilities is first bootstrapped by a “seed” ontology, lexicon and opticon – and uses these resources to continuously learn new content. At present, we bootstrap the system with an ontology of $\sim$160K RDF triples, an English lexicon of $\sim$30K word senses and a small opticon [4].

At present, learning-through-NL understanding still requires large amounts of expensive human time. In order to gradually alleviate this inefficiency, we have embarked on a program of work on substituting generative DL models for humans in the dialog. These models are capable of generating coherent text on multiple topics. While this breakthrough in generative AI has potential for knowledge acquisition, current DL models may still generate incomplete or fallacious outputs. Until “autonomous” LEIA-DL model dialog has been demonstrated to be viable, learning through LEIA-human dialog will remain a part of the overall learning environment (Figure 1), but the expectation is that in the suggested configuration humans will have less to do.

Learning is *triggered* if LEIA inference fails — for instance if a new entity type is encountered, as we have previously demonstrated in [5]. Once this happens, the system generates a request to a DL model to generate text and images as appropriate for the particular prompt. Model-generated text and images are interpreted as a TMR that is processed by the conceptual learning module to extract the precise content that will be used to augment the LEIA’s ontology, lexicon or opticon. Images (in text or other context) are used to acquire the opticon. The human is "lurking" in the background to resolve incongruities, but is *not* asked to actively provide inputs.

The above interaction between a DL model and a LEIA can take a push-me-pull-you turn: the DL model output supports automation of ontology, lexicon and opticon acquisition, and the content of LEIA knowledge resources in turn augments the training datasets of DL models.

Retraining the models with LEIAs’ knowledge resources will enhance success of the LEIA language interpreter in at least a couple of ways: model outputs can be fomulated using only items in the LEIA lexicon; and model outputs can be at least partially formulated in the LEIAs’ ontological metalanguage.

We present both a methodological position and a challenge to AI researchers interested in both ML and knowledge-based approaches. Therefore the final paper (which will describe the above material in detail) can be presented as either a challenge or a position paper (it has features of both).

Nikhil's contributions on using VoxML [6, 7, 8, 9]

References

[1] J. Mao, C. Gan, P. Kohli, J. B. Tenenbaum, J. Wu, The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision, in: International Conference on Learning Representations, 2018.

[2] A. d. Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, S. N. Tran, Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning, arXiv preprint arXiv:1905.06088 (2019).

[3] A. d. Garcez, S. Bader, H. Bowman, L. C. Lamb, L. de Penning, B. Illuminoo, H. Poon, C. Gerson Zaverucha, Neural-symbolic learning and reasoning: A survey and interpretation, Neuro-Symbolic Artificial Intelligence: The State of the Art 342 (2022) 1.

[4] J. Pustejovsky, N. Krishnaswamy, Voxml: A visualization modeling language, in: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16), 2016, pp. 4606–4613.

[5] S. Ghaffari, N. Krishnaswamy, Detecting and accommodating novel types and concepts in an embodied simulation environment, arXiv preprint arXiv:2211.04555 (2022).

[6] N. Krishnaswamy, P. Narayana, I. Wang, K. Rim, R. Bangar, D. Patil, G. Mulay, R. Beveridge, J. Ruiz, B. Draper, et al., Communicating and acting: Understanding gesture in simulation semantics, in: IWCS 2017—12th International Conference on Computational Semantics—Short papers, 2017.

[7] N. Krishnaswamy, S. Friedman, J. Pustejovsky, Combining deep learning and qualitative spatial reasoning to learn complex structures from sparse examples with noise, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, 2019, pp. 2911–2918.

[8] J. Pustejovsky, N. Krishnaswamy, Embodied human computer interaction, KI-Künstliche Intelligenz 35 (2021) 307–327.

[9] N. Krishnaswamy, J. Pustejovsky, Affordance embeddings for situated language understanding, Frontiers in Artificial Intelligence 5 (2022).