Knowledge Representation in Knowledge-Enhanced Machine Learning: How? Where?

Fabio Gagliardi Cozman

Escola Politécnica, Universidade de São Paulo
fgcozman@usp.br

Amazing success has been attained by artificial intelligences that resort to data intensive machine learning, for instance in natural language processing and in recommendation systems. Can we build an artificial intelligence endowed with full logical and commonsense reasoning just out of pattern extraction from ever increasing datasets? Possibly. But it seems reasonable to assume that tasks at higher abstraction levels demand at least bits of knowledge representation mixed with machine learning. In any event, several questions must be answered before we can have *knowledge-enhanced* machine learning at our fingertips.

How can we bring theoretical insights and practical tools from knowledge representation into machine learning tasks? Where is it worthwhile to add the power (and the cost) of knowledge representation to available datasets? How to evaluate the resulting combination of formalisms? This invited talk discusses these questions, necessarily focusing on a small subset of possible answers. Overall we emphasize the knowledge representation aspects of knowledge-enhanced machine learning, optimistically assuming that optimization and estimation methods will be available whenever needed.

We start by examining languages that combine logical formulas/rules with probabilities, as such languages must be key tools in our intended mix. The combination of logic and probability has an old and rich history; connections have been rediscovered more than once in artificial intelligence research [7]. In particular, during the past two decades there has been steady interest in languages that mix probabilistic graphical models, such as Bayesian networks, and relational logic [4]. Another line of research under investigation for more than twenty years has focused on probabilistic logic programming [10]. There are now solid techniques, often imported from finite model theory, that support us in studying these languages; results discussed in the talk are extracted from Refs. [1–3]. We compare the various languages, arguing that several ideas behind probabilistic logic programming are particularly valuable.

However, given the often "unreasonable" effectiveness of data in producing ostensibly intelligent behavior [6], it seems that we should *not* try to force knowledge representation into any machine learning task. Rather, we should carefully look for those tasks where knowledge-enhanced techniques will really make a difference. In this talk we discuss the task of explaining a link prediction in a knowledge base. In such a

Partially supported by CNPq grant 312180/2018-7 and by FAPESP grant 2016/18841-0.

task we do have knowledge, and the state of art methods resort to embeddings that are very difficult to interpret (that is, all entities and relations are mapped into vectors, and relationships are then expressed by relatively simple mathematical operations such as addition) [9]. The difficulty with embeddings is that decisions depend on numerical values that are apparently disconnected from semantic meaning. We discuss how explanations for link predictions can be extracted from embeddings, explanations that for instance resort to Horn clauses and similar formalisms [5, 11].

But even explanations can be learned from data: one can learn how to explain the behavior of another learner… and so on. Thus one might just argue that we can keep improving our pattern extraction methods so that they learn both to decide and to explain decisions, leaving aside any need for knowledge representation. To investigate the limits of knowledge-free learning, we propose a test inspired by the Winograd challenge [8] that can exercise the connection between commonsense reasoning and data intensive language processing. We suggest that such a Winograd Explaining Challenge, where the goal is to explain the answer to a Winograd scheme, can help focus our attention on problems that can only be solved by a mix of machine learning and commonsense reasoning. We discuss how we might go about facing such a test, and which research directions it opens.

References

1. Cozman, F.G., Mauá, D.D.: On the semantics and complexity of probabilistic logic programs. J. Artif. Intell. Res. **60**, 221–262 (2017)
2. Cozman, F.G., Mauá, D.D.: The complexity of Bayesian networks specified by propositional and relational languages. Artif. Intell. **262**, 96–141 (2018)
3. Cozman, F.G., Mauá, D.D.: The finite model theory of Bayesian networks: descriptive complexity. In: International Joint Conferences on Artificial Intelligence, pp. 5229–5233 (2018)
4. Getoor, L., Taskar, B.: Introduction to Statistical Relational Learning. MIT Press (2007)
5. Gusmao, A.C., Correia, A.H.C., Bona, G.D., Cozman, F.G.: Interpreting embedding models of knowledge bases: a pedagogical approach. In: ICML Workshop on Human Interpretability in Machine Learning, pp. 79–86 (2018)
6. Halevy, A., Norvig, P., Pereira, F.: The unreasonable effectiveness of data. IEEE Intell. Syst. 8–12, March/April 2009
7. Halpern, J.Y.: Reasoning About Uncertainty. MIT Press (2003)
8. Levesque, H.: Common Sense, the Turing Test, and the Quest for Real AI. MIT Press (2019)
9. Nickel, M., Murphy, K., Tresp, V., Gabrilovich, E.: A review of relational machine learning for knowledge graphs. Proc. IEEE **104**(1), 11–33 (2016)
10. Riguzzi, F.: Foundations of probabilistic logic programming (2018)
11. Ruschel, A., Gusmão, A.C., Polleti, G.P., Cozman, F,G.: Explaining completions produced by embeddings of knowledge graphs. In: Kern-Isberner, G., Ognjanović, Z. (eds.) ECSQARU 2019. LNCS(LNAI), vol. 11726, pp. 324–335. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29765-7_27