

Knowledge Acquisition with Transferable and Robust Representation Learning

1 Introduction and Background

My research focuses on promoting the advancement of intelligent computational systems with better awareness of commonsense and expert knowledge, which leads to more efficient information exchange of the system with people and the world. My goal, in the long term, is to leverage unified methodologies to help machines understand the relations of lexemes, entities and concepts in human languages as well as the interactions of objects in nature (such as molecules and biomolecules). In the near term, I am motivated by the objective of developing new technologies in representation learning and information extraction, and extending their use in various tasks for knowledge base construction, natural language understanding, computational biology and medicine.

The challenges span over a range of fields, from the fundamental questions in the acquisition, representation and inference of knowledge, to systematic paradigms for scalable data management, mining and retrieval. In this decade, AI systems in various application domains are empowered by representation learning technologies for automatically discovering and acquiring relations, patterns and properties of objects from large-scale data. In particular, such technologies involve relational embedding, language modeling and constrained learning. My work seeks to trigger the advancement of these technologies, with focus on knowledge acquisition from data in different modalities, and in scenarios with or without plausible supervision signals. My investigation in these research directions has led to over 30 published papers, and have benefited real-world applications in various computational and interdisciplinary areas. The following of this statement presents parts of my investigation into these directions, followed by some exciting future directions that I am planning to explore.

Transferable Representation Learning for Multi-relational Data. Multi-relational data, such as knowledge graphs (KGs), lexical graphs, product graphs and biological networks, serve as the essential source of both commonsense and domain knowledge for various knowledge-driven AI systems. Such data can model semantic relations of various entities, concepts and language components as well as interactions between molecules and biomolecules in the nature.

My research has focused on multi-relational representation learning, which seeks to support the characterization and inference of relational knowledge in embedding spaces. This is the requisite for incorporating the symbolic knowledge into deep learning models. Specifically, a key contribution I have made to this community is on the *transferability* of such representations. Different sources of multi-relational data, often provide *interchangeable* and *complementary* domain knowledge. Hence, it is particularly important to develop a general representation learning method that is able to capture the association of knowledge across multiple data sources with minimal supervision, and further support the transfer of knowledge across different domains. I started this line of research by proposing the first embedding framework that bridges multiple language-specific KGs [10, 16], by joint learning multiple relational embedding models with a semi-supervised alignment learning model. To more precisely capture the knowledge association with minimal supervision, I have extensively extended the alignment learning process based on iterative co-training [8], multi-view learning [33], incidental supervision from free text [6] and unsupervised visual pivoting [21]. I have also contributed with relational embedding techniques that are robust against scarcity and structural heterogeneity of data, particularly based on hyperbolic metric learning [26], concept contextualization [27] and attentive neighborhood aggregation [28]. For knowledge transfer from multiple sources of (inconsistent) learning resources, my work addresses the problem of inducing trustworthiness of inconsistent inference results with meta-learning [17].

This line of research has received a wide recognition by the community for its importance, and has received over two hundred citations in the past two years. A wide spectrum of applications have also been benefited from the techniques proposed in my papers and follow-up works. Besides knowledge alignment and KG completion, related techniques reached SOTA on few-shot entity typing [19], and improved protein-protein interaction prediction and disease target prediction with knowledge transfer from gene ontologies [20]. The advancement in this research topic has been systematically summarized in a part of our tutorial at AAAI-2020 [2] and our recent survey paper [29].

Knowledge Acquisition from Unstructured Data. While multi-relational data crucially provide structural and actionable knowledge representations to many AI applications, acquiring such data is costly and has often relied on extensive human efforts. Hence, there is a pressing need to develop methods for automated knowledge acquisition from unstructured data.

My research has developed data-driven relation induction methods from massive data of other modalities. Under this topic, I have studied the problem *large-scale document relation extraction*. My collaboration with Google Knowledge Graph has delivered a scalable system for document relation induction [4]. The system has received a near-perfect F1 score of over 98% for *subtopic* relation extraction on a ten-thousand-scale annotated data. Based on distributed inference with MapReduce, the system was able to efficiently extract relations from 24 trillion Wikipedia article pairs within several hours. My study has also delivered competitive end-to-end systems for *interaction prediction of biomolecules*, which sought to populate biological data banks and provide effective *in silico* techniques seeking to alleviate the expensive *in vivo* and *in vitro* efforts from the laboratory. These include the first sequence-based end-to-end system for multifaceted protein-protein interaction (PPI) prediction [3], and the point mutation effect estimation method based on language modeling technologies on amino acids [15]. In this line of research, I have also studied constrained learning for event-event relation extraction [31], retrofitting of language models for robust discourse relation detection [7], indirect supervision for ultra fine-grained intention inference for event processes [13], and joint learning of lexical and sentential semantics for semantic search [9, 24].

Learning and Inference with Complex Constraints and Properties. Data can possess highly complex properties and structures. In such situations, representation learning and inference of knowledge is often intractable. My work accordingly suggests solutions to such problems from several perspectives. In [18], I have proposed a novel uncertainty-aware representation learning method, UKGE, to capture both uncertainty and logical constraints with embeddings, which Probabistic Soft Logic rules incorporated in the learning phase. UKGE has been successfully applied to drug discovery in recent study [25]. A relevant technology has also been used to preserve the logical consistency of temporal, coreference and membership relations extracted for events from natural language text. I have also designed an approach to preserve hierarchical relational structures with rich profile information of relations in hyperbolic spaces [5], and extended the aforementioned approach to achieve SOTA performance on knowledge alignment and entity typing tasks [26]. My work has also demonstrated the potential of these representation learning approaches to improve hierarchical classification and extreme classification problems with better inference ability on the label space [19]. In this line of research, I also studied approaches for enforcing logical consistency for relation extraction from text [31], capturing logical properties of relations in linear embeddings [11] and supporting multi-label learning on multigraphs [1].

2 Research Agenda: Directions for Future Work

My research goal is to help machines acquire commonsense and expert knowledge through unified methodologies of representation learning, and leverage transferable representations to solve problems in various domains and interdisciplinary areas. I have already made significant strides toward this goal, and intend to focus on further extension of the methodologies described above over the next few years. At the CS department of USC, I am excited to further investigate in the following directions.

Learning with Minimal and Incidental Supervision. Training data-driven machine learning models require extensive collection of learning resources. However, annotating task-specific outputs is difficult and often requires expertise and domain knowledge, making it costly to obtain annotations that are with high quality and fairness. On the other hand, the Internet has produced enormous amount of textual and multimodal data. Although these data may not directly come with annotations for specific tasks, they involve related information that can be transformed into incidental supervision signals. However, the massive volume and noisy nature of unstructured and multimodal data present challenges to existing learning algorithms. Based on my pilot study of knowledge alignment and ultra fine-grained event process typing with indirect supervision [6, 13], I seek to further tackle the technical challenges of exploiting implicit and analogous supervision signals [32] from

heterogeneous learning resources by designing effective self-supervised learning algorithms. Towards this end, I am also interested in incorporating transferable representation learning of symbolic knowledge into learning systems, so as to support the systems with better inference ability on tasks with insufficient, biased training data or extreme label spaces.

Robust Machine Learning with Generalizability. Although data-driven machine learning benefits from automatic feature extraction, end-to-end training and inference, the produced models often suffer from generalization issues. Hence, I am interested in developing robust learning systems for various domain tasks. The methods I seek to explore include massively pre-training representation learning models and supporting domain invariant feature extraction. Through the investigation, the proposed techniques will be applied to a broad range of cross-lingual and low-resource language understanding tasks. Meanwhile, related techniques are also urged by the computational biology community to support tractable cross-species prediction tasks. A recent study [22] has shown that a machine learning model like PIPR [3] can effectively infer protein functions of *tomato* when being trained on *Arabidopsis*, but found the transferability from *yeast* to *Arabidopsis* to be very limited (the later two have a $10\times$ evolution distance in comparison to the former two). Given the intractable cross-species transferability of current methods, my future research would seek to deliver multi-species M-BERT-like models and species-invariant adversarial learning models. This is with the purpose of generally benefitting the computational biology community that studies on genomic and proteomic data. On the other hand, I also seek to better generalize the inference of models by addressing the “null prediction” problem, by teaching the model to understand the cases where there is no correct answers. I would like to tackle this challenge by developing a general adversarial learning framework to help a task-specific model to discriminate between answerable and unanswerable cases, and explore with advanced ordinal metric learning techniques to improve the characterization of label spaces.

Non-linear Characterization of Complex Structures. Many applications involve highly complex structures, which are particularly difficult for linear representation learning models to capture. I am interested in investigating new paradigms for representation learning in non-linear spaces. Based on my previous study on hyperbolic representation learning for multi-relational data [5, 26], I plan to enable learning systems with non-linear characterization of complex objects that form hierarchies, such as the evolution of claims and programs in online social media and software repositories. I am also interested in developing efficient *set learning* methods to help systems seizing order-invariant features, such as concurrent clinical events in EHR data.

Knowledge-aware Language Modeling. Language models are the backbone of many NLP systems. They critically capture the statistical relations of language components, and provide contextualized semantic representations for lexemes and sentences. While language models are mostly trained from scratch on raw text corpora, a line of research that has always been in my mind is to incorporate language models with both linguistic and commonsense knowledge [12]. My previous studies have achieved robust contextualized representations of lexemes and sentences by making a language model aware of contextual paraphrasing [7], and demonstrated that incorporating gloss knowledge helped a language model making prediction under an extreme multi-label and few-shot learning scenario [13]. Specifically, I believe a more reliable semantic representation method should also consider other linguistic properties of contexts, including different discourse relations and semantic classification of the linguistic contexts. On the other hand, various knowledge-driven NLP tasks would tremendously benefit from a language model that is aware of commonsense knowledge. Through this line of research, I intent to study how we can probe language models with multiple aspects of relational knowledge, gloss knowledge [9] and temporal knowledge.

Cross-domain and Interdisciplinary Research. I always believe that a useful technology should address problems in several related research areas rather than a single one. Therefore, my research for multi-relational representation learning and relation induction has contributed to a broad range of tasks in knowledge base construction [10, 4], natural language processing [9, 7], computational biology [3, 15], computer networks [14] and social media analysis [30, 23]. Some of my ongoing or planned works are also extending the proposed

methods in my Ph.D. study to new application scenarios, including diagnostic prediction on EHR data [34] with knowledge transfer from disease ontologies, and projecting protein and disease knowledge to predict interactions of drugs. Given the previous success in transferring technologies to different areas, I am enthusiastic about developing open-source libraries and software, and facilitating collaborations with people outside my areas. I am excited about any opportunities to apply my expertise in representation learning and NLP to solve important problems in other areas and disciplines.

References

- [1] CHEN, H., TIAN, Y., CHEN, M., PEROZZI, B., AND SKIENA, S. Enhanced network embeddings via exploiting edge labels. In *CIKM* (2018).
- [2] CHEN, M., CHANG, K.-W., AND ROTH, D. Recent advances in transferable representation learning. In *AAAI Tutorials* (2020).
- [3] CHEN, M., JU, C., ZHOU, G., CHEN, X., ZHANG, T., CHANG, K.-W., ZANIOLO, C., AND WANG, W. Multifaceted protein-protein interaction prediction based on siamese residual rcnn. *Bioinformatics* 35, 14 (07 2019), i305–i314 (Procs of ISMB 2019).
- [4] CHEN, M., MENG, C. P., HUANG, G., AND ZANIOLO, C. Neural article pair modeling for wikipedia sub-article machine. In *ECML* (2018).
- [5] CHEN, M., AND QUIRK, C. Embedding edge-attributed relational hierarchies. In *SIGIR* (2019).
- [6] CHEN, M., SHI, W., ZHOU, B., AND ROTH, D. Cross-lingual entity alignment for knowledge graphs with incidental supervision from free text. In *EACL (In review)* (2020).
- [7] CHEN, M., SHI, W., ZHOU, P., AND CHANG, K.-W. Retrofitting contextualized word embeddings with paraphrases. In *EMNLP* (2019).
- [8] CHEN, M., TIAN, Y., CHANG, K.-W., SKIENA, S., AND ZANIOLO, C. Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In *IJCAI* (2018).
- [9] CHEN, M., TIAN, Y., CHEN, H., CHANG, K.-W., SKIENA, S., AND ZANIOLO, C. Learning to represent bilingual dictionaries. In *CoNLL* (2019).
- [10] CHEN, M., TIAN, Y., ET AL. Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In *IJCAI* (2017).
- [11] CHEN, M., TIAN, Y., ET AL. On2vec: Embedding-based relation prediction for ontology population. In *SDM* (2018).
- [12] CHEN, M., AND ZANIOLO, C. Learning multi-faceted knowledge graph embeddings for natural language processing. In *IJCAI* (2017).
- [13] CHEN, M., ZHANG, H., WANG, H., AND ROTH, D. “what are you trying to do?” semantic typing of event processes. In *CoNLL* (2020).
- [14] CHEN, M., ZHAO, Q., DU, P., ZANIOLO, C., AND GERLA, M. Demand-driven cache allocation based on context-aware collaborative filtering. In *MobiHoc* (2018).
- [15] CHEN, M., ZHOU, G., JU, C., WANG, Z., JIANG, J.-Y., AND WANG, W. Mutation effect estimation on protein-protein interactions using deep contextualized representation learning. *NAR Genom Bioinform* (2020).
- [16] CHEN, M., ZHOU, T., ET AL. Multi-graph affinity embeddings for multilingual knowledge graphs. In *AKBC* (2017).
- [17] CHEN, X., CHEN, M., FAN, C., UPPUNDA, A., AND ZANIOLO, C. Cross-lingual knowledge graph completion via ensemble knowledge transfer. In *EMNLP - Findings* (2020).

- [18] CHEN, X., CHEN, M., SHI, W., SUN, Y., AND ZANIOLO, C. Embedding uncertain knowledge graph. In *AAAI* (2019).
- [19] HAO, J., CHEN, M., YU, W., SUN, Y., AND WANG, W. Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts. In *KDD* (2019).
- [20] HAO, J., JU, C., CHEN, M., SUN, Y., ZANIOLO, C., AND WANG, W. Bio-joie: Joint representation learning of biological knowledge bases. In *ACM BCB* (2020).
- [21] LIU, F., CHEN, M., ROTH, D., AND COLLIER, N. Visual pivoting for (unsupervised) entity alignment. In *AAAI* (2021).
- [22] MAKRODIMITRIS, S., VAN HAM, R., AND REINDERS, M. Sparsity of protein-protein interaction networks hinders function prediction in non-model species. *bioRxiv* (2019), 832253.
- [23] MENG, C., CHEN, M., MAO, J., AND NEVILLE, J. Readnet: A hierarchical transformer framework for readability analysis. In *ECIR* (2020).
- [24] SHI, W., CHEN, M., TIAN, Y., AND CHANG, K.-W. Learning bilingual word embeddings using lexical definitions. In *RepLANLP* (2019).
- [25] SOSA, D. N., DERRY, A., GUO, M., WEI, E., BRINTON, C., AND ALTMAN, R. B. A literature-based knowledge graph embedding method for identifying drug repurposing opportunities in rare diseases. In *PSB* (2020).
- [26] SUN, Z., CHEN, M., HU, W., WANG, C., DAI, J., AND ZHANG, W. Knowledge association with hyperbolic knowledge graph embeddings. In *EMNLP* (2020).
- [27] SUN, Z., HUANG, J., HU, W., CHEN, M., AND QU, Y. Transedge: Translating relation-contextualized embeddings for knowledge graphs. In *ISWC* (2019).
- [28] SUN, Z., WANG, C., HU, W., CHEN, M., DAI, J., ZHANG, W., AND QU, Y. Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In *AAAI* (2020).
- [29] SUN, Z., ZHANG, Q., HU, W., WANG, C., CHEN, M., AKRAMI, F., AND LI, C. A benchmarking study of embedding-based entity alignment for knowledge graphs. *PVLDB 13* (2020).
- [30] TIAN, Y., CHEN, H., CHEN, M., PEROZZI, B., AND SKIENA, S. Social relation inference via label propagation. In *ECIR* (2019).
- [31] WANG, H., CHEN, M., ZHANG, H., AND ROTH, D. Joint constrained learning for event-event relation extraction. In *EMNLP* (2020).
- [32] ZHANG, H., CHEN, M., WANG, H., AND ROTH, D. Open-domain process structure induction. In *EMNLP* (2020).
- [33] ZHANG, Q., SUN, Z., CHEN, M., GUO, L., AND QU, Y. Multi-view knowledge graph embedding for entity alignment. In *IJCAI* (2019).
- [34] ZHANG, T., CHEN, M., AND BUI, A. Diagnostic prediction with sequence-of-sets representation learning for clinical event. In *AI in Med (AIME)* (2020).