Although the intelligent systems of today can easily solve a multitude of tasks, their lack of world knowledge makes them uninterpretable and incapable of complex reasoning. We teach robots to cut oranges by copying the actions of humans instead of imparting them with the knowledge of an orange's physical properties. We teach machines to read encyclopedias by pointing them to texts written by humans rather than the factual knowledge underlying the texts. Without the understanding of such knowledge, these large neural models are trained to be unreliable mimics, posing risks including amplifying pre-existing biases in the data. I am thus interested in integrating knowledge into our powerful yet opaque neural models so their complicated behaviors can be grounded to human-understandable world knowledge.

My past research explored the interplay between language and knowledge, focusing on two research directions:

- **Knowledge-enhanced NLP**: How can knowledge help natural language understanding?
- **Language-inspired KR**: How can language inspire the design of knowledge representation?

Knowledge-enhanced NLP Current NLP models lack various kinds of knowledge including encyclopedic knowledge about entities of interests, commonsense knowledge about the physical world, and cultural knowledge that rationalizes human social and ethical behaviors. I am fascinated with the possibility of enabling NLP models to reason with such knowledge.

To investigate deeper into the integration of knowledge and language, I have been working with Professor Yangfeng Ji studying language tasks that can be contextualized by existing declarative knowledge within graph-structured data. Specifically, I focused on the challenging problem of introducing meaningful interaction between language and knowledge representation. The first task I looked at was Text-to-SQL semantic parsing guided by relational knowledge in database schemas. In this task, a graph neural network module can explicitly encode the domain knowledge in the database schema, allowing us to build a semantic parser that can work across different domains. In our paper at COLING (Chen et al., 2020), I showed that different types of interaction between language and knowledge can be formulated at different abstraction levels of their representation, and the most important type of interaction varies based on which SQL component is currently being generated.

Along this direction, I intended to explore whether such multi-level interactions can be introduced between state-of-the-art language models (e.g., BERT) and more universal knowledge graph representation. In my master's thesis (Chen, 2020), I demonstrated that combining a BERT model with a multilayer knowledge graph representation model using a crossmodal attention mechanism can substantially enhance the model's ability to answer factoid questions. More interestingly, I found that different connection strategies between layers of BERT and the knowledge graph representation can lead to very different results, suggesting that a deeper understanding of the knowledge stored in each of their representation layers is important to build meaningful interactions between them.

While such domain-specific knowledge and encyclopedic knowledge can directly contextualize language understanding, the benefits they bring are not just additive. Humans do not learn language from textual symbols alone, instead, we ground the meaning of language and heterogeneous multimodal input (e.g., images and video) to our existing world knowledge. This alignment between language and other modalities, associated with a shared body of knowledge, helps us to learn language efficiently. As current neural models are extremely overparameterized and trained on billions of words, I would like to explore this potential of greatly improving sample efficiency to help us make further progress.

Language-inspired KR Although current natural language understanding models seem to pick up a large amount of knowledge, their representation of knowledge is implicit. We do not understand what knowledge is contained in the representation or how it is organized and utilized, making it unreliable and uncontrollable. In an attempt to address this issue, I intend to draw insights from language, a complex but efficient symbolic system we have, to build representations that organize knowledge in an explicit way.

As the main conduit for humans to pass knowledge across generations, language embodies the organization of knowledge in our mind, reflecting many important properties of the underlying knowledge it is used to represent. My past research explored building knowledge representation using learning methods that are designed for common properties of language and knowledge. In our EMNLP paper (Chen et al., 2021), inspired by recent findings on the emergent linguistic structures from attention weights of Transformer-based language models, I adapted the Transformer to learn the structures of noisy knowledge graphs. Using flattened subgraphs as input sequences, the model learns multilayer representation that captures the hierarchical structures of knowledge, achieving new state-of-the-art results in knowledge graph completion tasks.

The compositionality of language shows us that complex meaning can be composed through a small set of primitives. In an ongoing project, I attempted to represent entities in knowledge graphs using subwords from their language descriptions. This combination of structured and unstructured knowledge offers us the best of both worlds, achieving extensive knowledge coverage while avoiding excessive memory usage caused by the formidable number of entities in large-scale knowledge graphs. Following this direction, I plan to further explore a unified framework for such general properties of knowledge. Within it, we can define a similar set of atomic units for knowledge, on top of which levels of abstraction can be formed to manage the mutual interaction and information sharing among different types of knowledge.

Future plan After completing my Ph.D. study, I plan to be a professor to continue performing research and education in this field. In the past, I was fortunate to be mentored by many people when I was lost in doubt about my academic career. On the other hand, teaching and mentoring students have been a central part of my undergraduate life—I used to lead our university's programming competition training team consisting of nearly 40 members and serve as the (head) TA for several programming classes. Education lies in the heart of scientific progress. The knowledge accumulated by one generation is limited, so I would like to pass my little understanding of artificial intelligence research to the next generation.

Duke University attracts me to its wonderful NLP environment and broader interdisciplinary artificial intelligence research. Professor Bhuwan Dhingra's interests in knowledge representation and their applications on language understanding and Professor Samuel Wiseman's interests in structured models that track entities and other world states underlying texts are perfectly aligned with my research interests, and my past experience in modeling graph-structured knowledge with language will directly contribute to these directions during my Ph.D. study at Duke.

Sanxing Chen. 2020. Contextualizing language understanding with graph-based knowledge representations. Master's thesis, The University of Virginia.

Sanxing Chen, Xiaodong Liu, Jianfeng Gao, Jian Jiao, Ruofei Zhang, and Yangfeng Ji. 2021. HittER: Hierarchical transformers for knowledge graph embeddings. In *Proc. of EMNLP*.

Sanxing Chen, Aidan San, Xiaodong Liu, and Yangfeng Ji. 2020. A tale of two linkings: Dynamically gating between schema linking and structural linking for text-to-sql parsing. In *Proc. of COLING*.