Cognitive Computation Group Research Overview

2019-2020

Incidental supervision

How should we understand, acquire, and use signals that were not put there to help a specific target task?

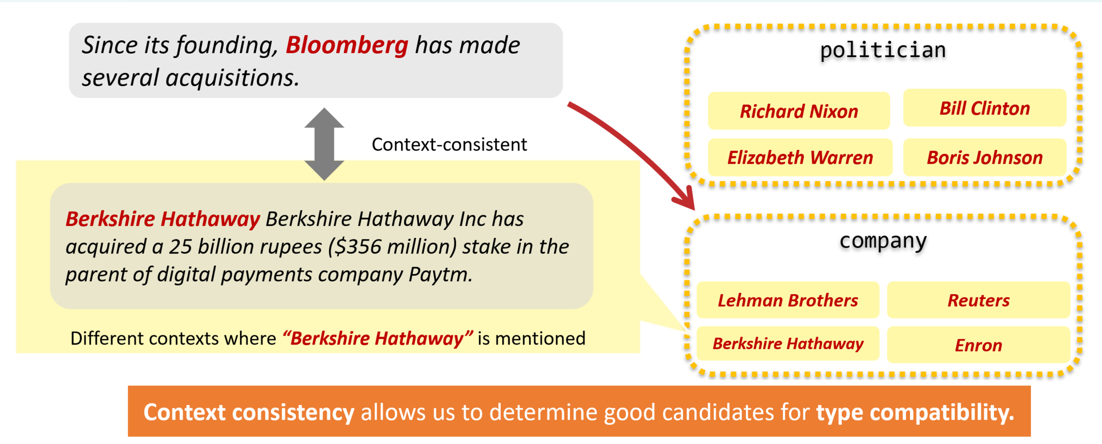

In Zero-Shot Open Entity Typing (ZOE), *type-compatibility* is used as a supervision signal.

Machine Learning and Inference methods have become ubiquitous in our attempts to induce more abstract representations of natural language text, visual scenes, and other messy, naturally occurring data, and to support decisions that depend on it. However, learning models for these tasks is difficult, partly because generating the necessary supervision signals for it is costly and does not scale.

We study several learning paradigms designed to alleviate the supervision bottleneck, from zero-Shot (Dataless) learning to Response Driven Learning – a learning protocol that supports inducing representations simply by observing the model's behavior in its environment – to learning from definitions and available text. We develop theoretical understanding for these paradigms and make use of them in a range of NLP applications, from semantic typing to (cross-lingual) text classification to temporal relations.

Learning and Reasoning

Humans engage in reasoning – we make decisions that involve (i) assigning values to multiple interrelated variables, (ii) making multiple, interdependent, inference steps, and (iii) using discrete computations (logical or other) over inferred variables. These computations often require incorporating background knowledge to facilitate robust behavior in new situations. Our earlier work on **Learning to Reason** suggested that Reasoning should be studied together with Learning and the Representation it produces.

CCMs (a.k.a. **Integer Linear Programming** formulations for NLP) provide an abductive framework addressing some aspects of this view in a learning and

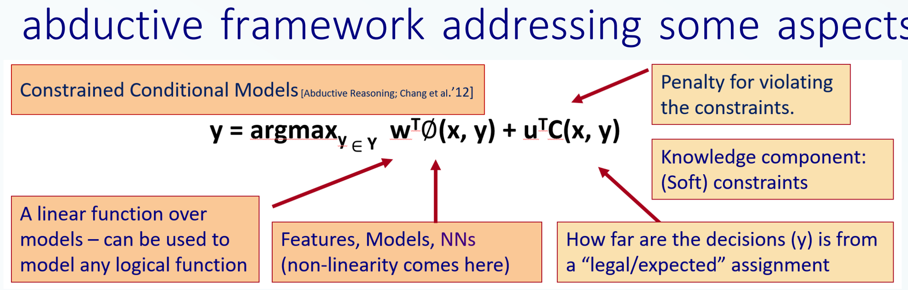

Reasoning approach that augments learning of models with declarative constraints (background knowledge) to support assigning values to multiple interrelated variables.

But our models still don't know how to respond to surprising questions such as "Did Aristotle have a laptop?", read a football game recap and reason about scoring scenarios, or reliably solve algebra word problems

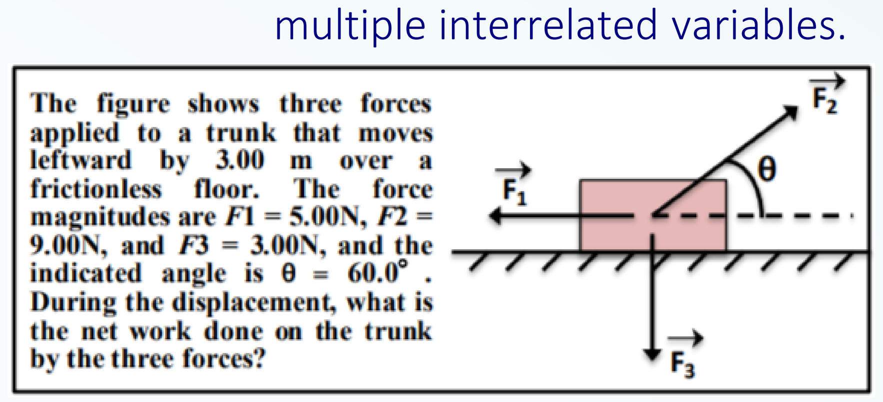

that require both text understanding and some "reasoning" capabilities. We study representations, reasoning paradigms, and learning approaches to address these questions.

Research Focus

Our research focuses on the computational foundations of intelligent behavior. We develop theories and systems pertaining to intelligent behavior using a unified methodology, at the heart of which is the idea that learning and reasoning have a central role in intelligence. Our work centers around studying machine learning and inference methods that facilitate **Natural Language Understanding (NLU) –** developing programs that support multiple aspects of machine reading and that will eventually communicate with humans the way humans do. Such systems must acquire the bulk of their knowledge from real world data, and behave robustly when presented with new, previously unseen situations. Therefore, our technical focus has been on paradigms for incidental supervision, and for inference that makes use of knowledge learned, read, and given. The foundational work is driven by a range of need-to-be-solved NLU tasks, and by applications such as English as a Second Language (ESL), NL acquisition, multilingual NLP, medical NLP, and navigating Information Pollution.

Applications and Driving Forces

The foundational work described on the left is driven by and studied in the context of Natural Language Understanding (NLU) tasks that we deem important. Some key representatives are described below.

Events and Situations (not sentence processing) are the backbone of NLU. This perspective drives a lot of the NLU research we have done in the last few years and will continue doing so. There is a need to identify events at multiple granularities, and understand their logical and temporal structure, components and participants, as well as relations between them. We have worked on several aspects of this level of understanding – identifying events, understanding **time** and temporal relations between events, event-level **(semantic) language models**, and more. We have also developed tools that support language understanding at the primitive event level (e.g., **Semantic Role Labeling** with respect to multiple predicate types). We work on a several important **Information Extraction** tasks, from understanding **Quantities** (and solving **algebra word problems**) to **semantic typing** and **NER** to Entity **Linking** (Wikification) and **coreference**. While most of the work in NLP has been done in English, thousands of other languages are being used daily, many of which are **low-resource languages**, making most of the current NLP technology useless. We study approaches that provide access to low-resource languages by English speakers, even when translation isn't available – **cross-lingual representations**, **text classification**, NER, entity-linking, etc. We also work on **English as a Second Language (ESL),** developing methods to improve and correct the writing of non-native speakers. Our **BabySRL** project reflects the view that **investigating models of language acquisition by children** could enrich our NLP work, while our machine learning expertise can help guide the work on Psycholinguists. Our joint work with psycholinguists focuses on predicate-argument acquisition.

Communication: Language in Context

The study of the learning, inference, and knowledge representations mechanisms that facilitate NLU requires that we study understanding human language in context. We need to study how systems interact with data, with knowledge, and with humans. This involves thinking about grounding, learning from the environment's response, and learning in

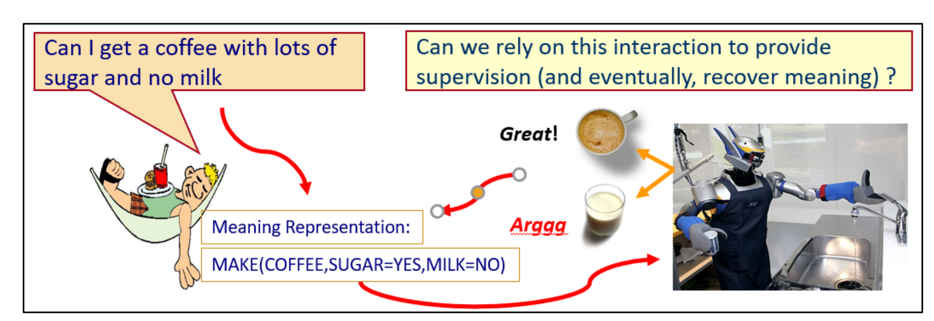

context, while accounting for the domain, the task at hand, and the human-machine shared knowledge.

Head: Dan Roth
Research Engineer: Hegler Tissot
Post Docs: Muhao Chen, Elior Sulem
Undergrad Students: Jamaal Hay, Nicole (Xinran) Han, Celine Lee, Francesca Marini, Tatiana Tsygankova.

Program Coordinator: Jennifer Sheffield
PhD Students: Sihao Chen, Soham Dan, Dan Deutsch, Hangfeng He, Nitish Gupta, Qing Lyu, Kaifu Wang, Xiaodong Yu, Yi Zhang, Ben Zhou.
MS Students: Disha Jindal, Sanjna Kashyap, Krunal Shah, Haoyu Wang

Navigating Information Pollution

In an era where generating content and publishing it is so easy, we are bombarded with information and are exposed to all kinds of claims – in news, the medical domain, education, and commerce – some of which do not rank high on the truth scale.

This Information Pollution – the contamination of the information supply with irrelevant, redundant, unsolicited, incorrect, and otherwise low-value information, is the subject of this line of work. Our goal is to define and address some of the key research questions raised by the need to navigate our way through it: from key natural language processing problems that arise when attempting to identify and present the multiple perspectives a claim might have, along with its supporting evidence, to understanding information

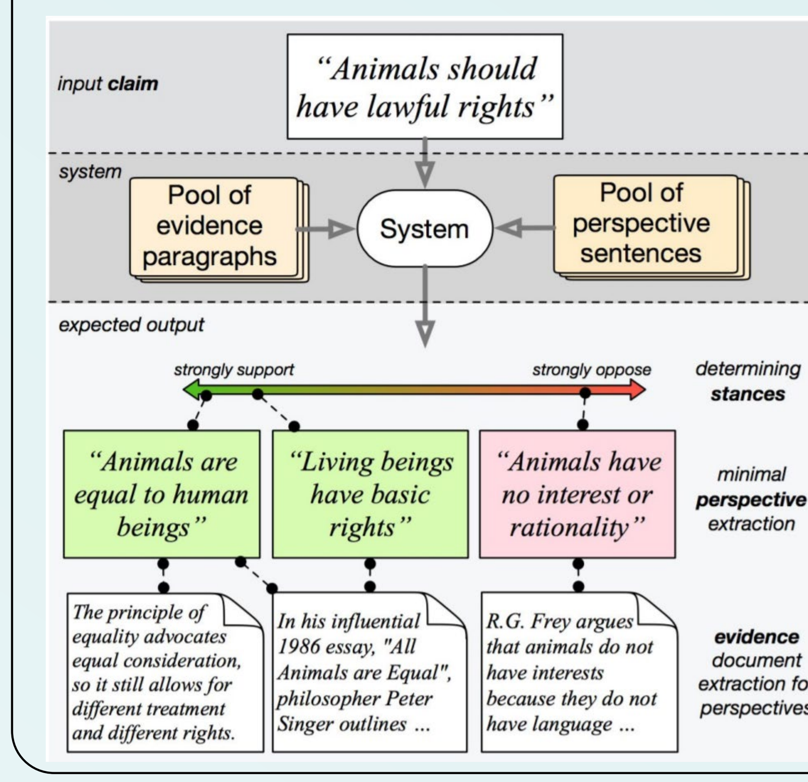

sources, the claims they make, and evidence they provide, to an algorithmic inference framework for trustworthiness. We define novel learning and inference tasks that would provide important building blocks for addressing information pollution, and novel NLU tasks to characterize similarities and differences among claims, the intent behind them, perspectives they express, and their implications.

Commonsense Reasoning

Humans have a store of commonsense knowledge that we can quickly access and reason with, to make sense of new situations and make inferences about the world around us. Automating natural language understanding requires models that are informed by commonsense

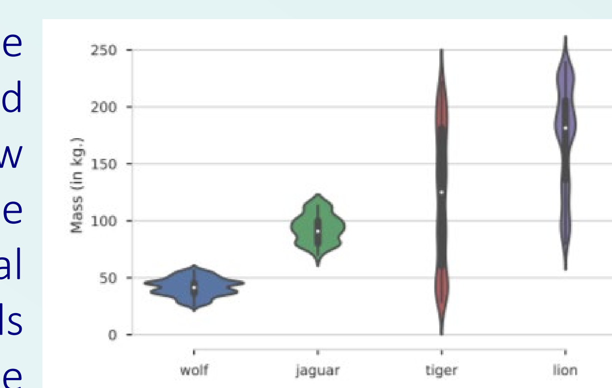

knowledge and the ability to reason with it in both common and unexpected situations. The success of statistical and deep learning methods has supported advances in some aspects of AI, but our models still do not know that "get me a piece of cake" requires first getting utensils, then cutting the cake, and placing it on a plate, and that it typically takes minutes (as opposed to "baking a cake"); and they don't know that NYC is *always* on the East Coast, but Paul Simon is *sometimes* there. We study an encompassing approach to commonsense reasoning and AI that avoids nonsensical decisions. Our approach builds on a knowledge acquisition effort – we have worked on Quantities and Time (shown on the right) already – along with a reasoning effort inspired by the observation that "reasoning is common sense".