Knowledge Representation and Reasoning

- **Initial Assumption:** Manually engineered systems for extracting and representing the meanings of sentences will not be good enough. Learning must be incorporated.
- **Key Challenge:** How to make learned representations effective enough so that they can support representation of meaning, and reasoning with it, and can do it across domains and in both known and new situations.

- **Goals**: Already talked about it.

    - Warning: very easy to believe the new datasets represent the goals of reasoning. Always think about the challenging scenarios we discussed, and just look at text and think about what's needed.

- **Background:** We'll talk a bit about history, from the NLU perspective
    - The Classical papers gives you a (very partial) perspective into it.

- Talk about **learning** (mostly: learning protocols)

Representations:

To facilitate answering questions there is probably a need to map a surface form to some abstract, executable representation.

What is Meaning?

> **[Lewis, David. 1970. General semantics.]** In order to say what a meaning is, we may first ask what a meaning does, and then find something that does that.

Classical view: The semantic meaning of the sentence is taken to be some expression of formal logic. This is an instance of the "**truth-Theoretic**" view of formal semantics:

The meaning of the sentence as the set of possible situations in which that sentence is true. (**Possible Worlds** semantics).

The key representational tool used under this approach is that of **Logical forms:** a functional representation of meaning that can be evaluated alongside a **model** to yield a truth value (T, F).

A model is an assignment of variables to the sets (man, woman, love, in this case), in such a way that predicates in the logical forms takes truth values.

(i) A woman loves every man.
(ii) $\forall x\big[man(x) \rightarrow \exists y[woman(y) \wedge love(y,x)]\big]$

The truth value of the sentence is thus determined **compositionally.** The (natural) assumptions is that the meaning of an expression is a function of the meanings of its parts and the way they are syntactically combined.

That requires a more refined **representational theory**, along with an **algorithmic** (proof theoretic, or model theoretic) way of computing the truth value of a statement (sentence).

Specifically, the process of interpretation in this view of language understanding is generally broken into three components: **lexical semantics**, c**ompositional semantics**, and **pragmatics**.

- **Lexical semantics** characterizes the meanings of the smallest meaningful units of language.

- **Compositional semantics** describes the formal procedures by which these minimal units are assembled together into sentence meanings, yielding logical forms like the one above.

- **Pragmatics** describes the systematic reasoning processes by which listeners infer aspects of sentence meaning that aren't supported by the literal content of a sentence.

A lot of work went into specifying a truth-conditional semantics; this involves constructing complex functional forms for words (and other basic units) within the lexical semantics such that those words can combine with other words in a way that yields correct sentence meanings.

A lot of formalisms have been developed – the most important one studied today is that of **Combinatory categorial grammar (CCG)**, a formalism developed by **Mark Steedman** that is interface between surface syntax and underlying semantic representation. This has been quite successful as a framework for semantic parsing, especially when the goal is to represent single sentences well enough to access a database.

There are other versions of CCG, e.g., **Dependency-Based Compositional Semantics** (DCS, Percy Liang) that are attempting to achieve about the same goal: representing the meaning of a single sentence well enough, so that it facilitates access to knowledge bases.

The bottom line is that a complete truth-conditional semantics has not yet been completed for any language. However, almost all other representations and reasoning formalisms take a lot from the **truth-theoretic** approach to representation.

In one way or another, many other **semi-logical formalisms fall into this category**, and extend it in different ways.

- Semantic networks
- Descriptions logics
- Probabilistic logics

Relational view of meaning:

This view has two components, **representational and inferential**, although this has not necessarily been viewed this way by any of the people pursuing this direction.

The key conceptual difference is the view that instead of "mapping" text (and our knowledge) to a **"structured" representation** of some sort, we'll continue to represent it directly in text. Of course, we need to enrich the text somehow (some syntactic annotation; some shallow semantic annotation, etc.) and then we can use this (enriched) textual representation to **directly reason** about truth and false.

It's relational, since we are reasoning about **relations between sentences** (more generally, pieces of text), but also because we **enrich the text with relational information.**

A natural inference formulation here then becomes that of **textual entailment**: determining whether one snippet of text (Text; Premise) entails another (Hypothesis).

Text: Currently, there is no specific treatment available against dengue fever, which is the most widespread tropical disease after malaria. Sanofi Pasteur is collaborating with the Communicable Disease Center in Singapore and the Pasteur Institute in Vietnam to conduct these clinical studies in children and adults. "Controlling the mosquitoes that transmit dengue is necessary but not sufficient to fight against the disease," he says, ….

Hypothesis 1: Malaria is the most widespread disease transmitted by mosquitoes.

Entailed? Probably not; why, and what NLU phenomena participate?

Hypothesis 2: Dengue fever is the most widespread tropical disease after malaria.

Entailed? Yes

Hypothesis 3: Malaria is the most widespread tropical disease

Entailed, Yes.

Most importantly:

1. It is clear that we can model any NLU phenomenon as textual entailment.
    a. **Text:** Obama and his wife Michelle travelled to Paris to meet with the….

2. The definition is ill defined – it's not logical entailment; it's "textual entailment" and requires NLU.
    a. **Hypothesis:** Obama travelled to France

3. It's just a framework for thinking about/doing NLU. It does not commit you to any representational nor inferential framework.
    a. Indeed, everything has been tried.

However:

One interesting formalism is trying to take this approach seriously as a way to advance NLU. (Historically, it was developed before Textual Entailment).

Natural logic is a proof theory over the syntax of natural language. It eliminates the need for semantic parsing and domain-specific meaning representations by providing **inference rules** that apply **directly at the level of the text.** It claims that the inferences warranted by the logic tend to be the same inferences that are cognitively easy for humans – that is, the inferences humans assume a reader will effortlessly make.

T: the cat ate a mouse
H: no carnivores eat animals
Entailment: False

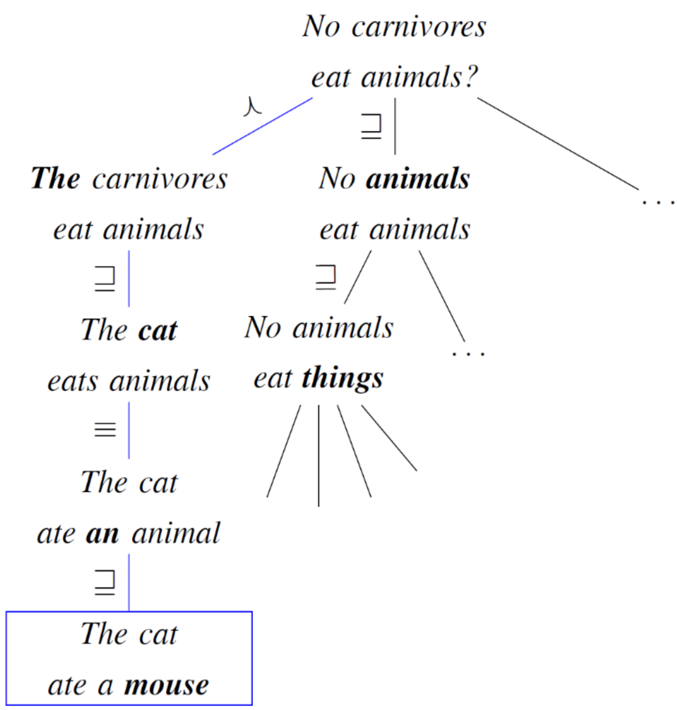

The system constructs a search problem for each queried hypothesis over relaxed natural logic inferences: **the surface form of the hypothesis is allowed to mutate** until it matches one of the facts in the knowledge base. These mutations correspond to steps in a natural logic proof; a learned cost for each mutation corresponds to the system’s confidence that the mutation is indeed logically valid (e.g., mutating to a hypernym has low cost, whereas nearest neighbors in vector space has high cost).
This really become a “fuzzy theorem prover”; **we still have the problem of Knowledge (facts, and inference rules).**

Not the only formalism that considers transformations – from Hypothesis to Premise or the other way around. But, how do we acquire transformation? how to we decide which to take? How do we weigh them?

Leaves the question of representation almost intact:

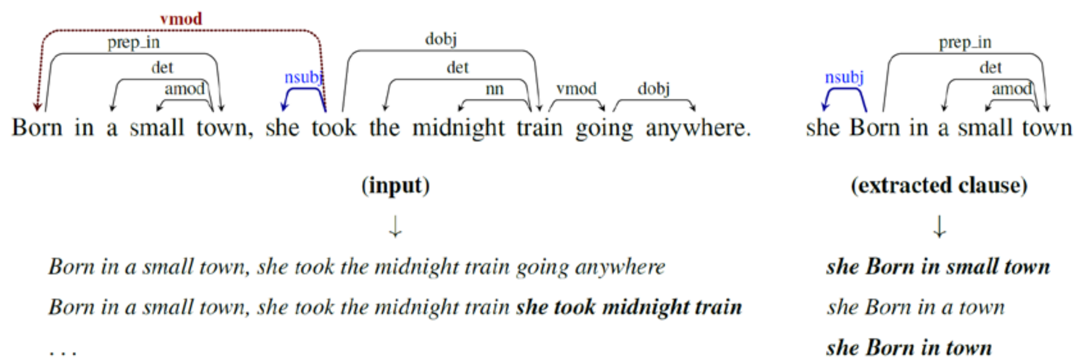

Can we do something without these kinds of representations?

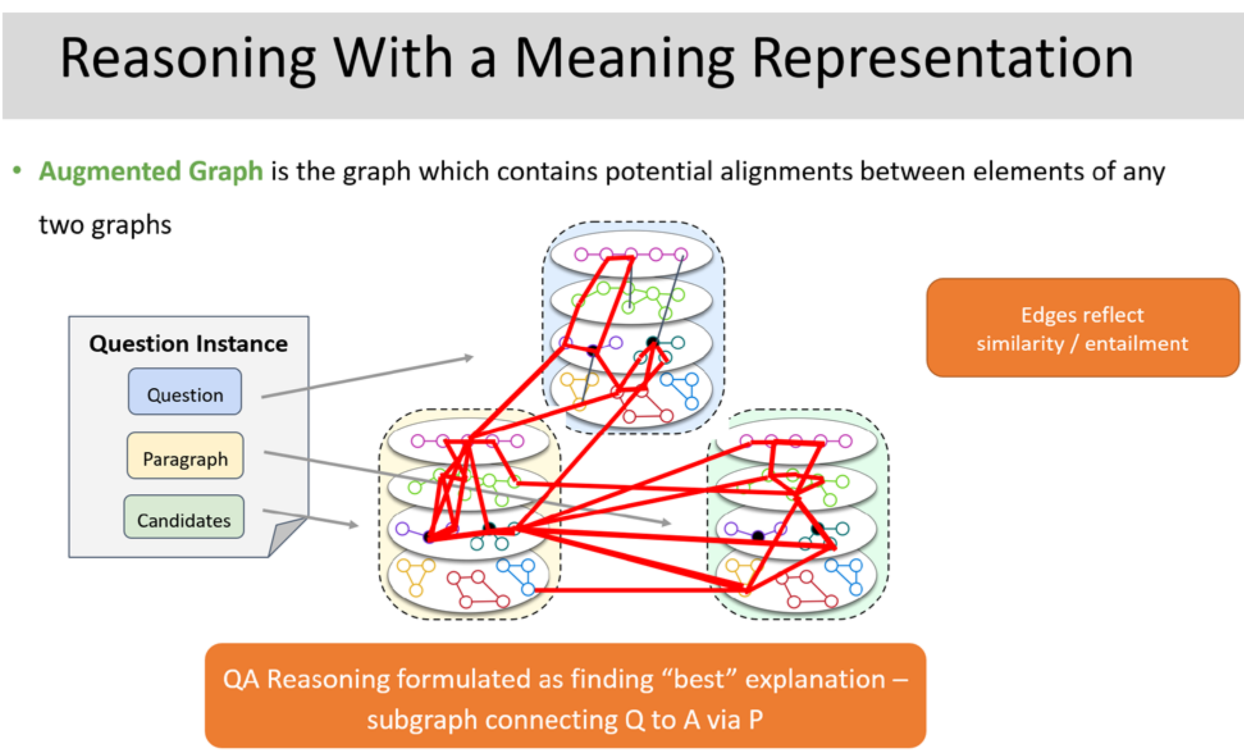

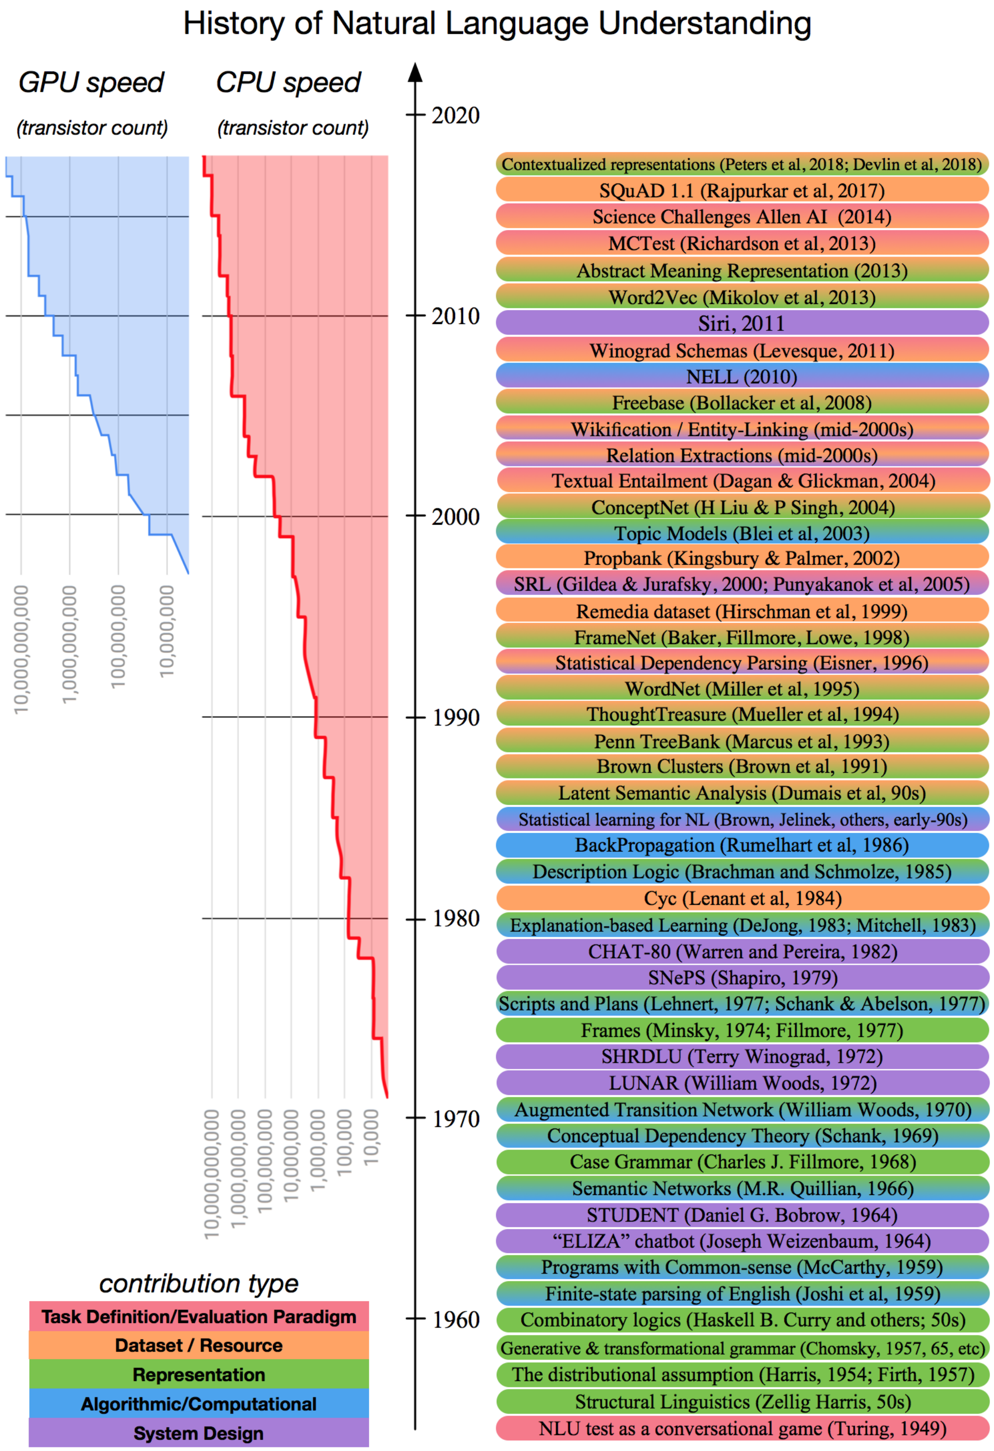

Figure 5: Major highlights of NLU in the past 50 years (within AI community). For each work, their contribution-type is color-coded. To provide perspective about the role of the computational resources available at each period, we show the progress of CPU/GPU hardware over time.