

MOTIVATION

Texts, such as the news, are driven by the entities they discuss and these entities are referred to using different words throughout a text.

Can we automatically resolve which nouns in a text refer to the same real world entity? If we are able to do this, computers become more similar to human readers, able to extract the quotes, facts, and other information from the articles in which an entity appears

How do we get a computer to understand that Kevin Bacon is the “resident pet” at The Grounds? Humans resolve these kinds of references easily while reading a text, but the difficulty of encoding the real world knowledge shared between readers and authors has so far inhibited the automatic resolution of such common noun references.

COREFERENCE RESOLUTION

Coreference resolution is the task of extracting markable noun phrases (mentions) in a text and clustering these into groups which correspond to a real world entity. Some of the surface level linguistic cues currently used are:

- canonical name
- semantic class
- animacy
- gender
- grammatical and contextual information
- *! real world knowledge

Kevin Bacon has been pig-napped. The much loved resident pet at The Grounds of Alexandria, who has a solid following on Facebook, was stolen from his pen overnight, along with ...

Example showing resolution of the pronoun “who” with “Kevin Bacon” using syntactic cues of relativisation and “his” with “Kevin Bacon” based on both likely being references to a male entity. Note that there is no straightforward way to resolve “resident pet”.

RESEARCH QUESTIONS

1. Current state-of-the-art models are “knowledge poor”.

How can we better make use of the web scale resources to learn semantics? Discourse patterns? Inference?

2. If clustering is performed post-process, any information they implicitly encode can't be accessed during learning.

How can we best include cluster level consistency information during learning?

3. Simple pairwise comparison compares each mention against (potentially all) those preceding it.

Can we model the badness of fit between mentions, as well as traditional goodness of fit features?

What practical methods can we use to improve this efficiency to enhance scalability?



The Grounds: missing a pig and a sheep. Photo: Supplied

Kevin Bacon has been pig-napped.

The much loved resident pet at The Grounds of Alexandria, who has a solid following on Facebook, was stolen from his pen overnight, along with another children's favourite, Bradley the lamb.

Two women and a man were captured on CCTV camera entering the cafe's grounds at 12.57am on Monday and stealing the animals, which were popular among parents and children who visited the cafe.

All the kids this morning are running up to the pen, asking “Where are they?”, and we've just been saying that they're on holidays.

The thieves, who were filmed walking around The Grounds with torches, were believed to have led Kevin and Bradley through the carpark and into a car on Huntley Street before driving away.

General manager Michael Gebran said he was desperate to find the cafe's two beloved “family members”. ...



EVALUATION

We evaluate our experiment on the standard dataset for coreference resolution, namely the English portion of OntoNotes: a 1G+ word collection of text across genres (including newswire, broadcast news and conversation, telephone conversation, magazine and web text), annotated for coreference, as well as syntactic parse and semantic word sense and class information. Our experiments so far have been exploratory and based on these, we propose the following system architecture to improve the resolution of common nouns.

