

Animism in Scientific Writing

Supervisor: Prof Simone Teufel (sht25@cam.ac.uk)

Animism is the process of referring to inanimate entities as if they had a mind (that is, the ability for volition). This unfortunate thing happened in the following sentences:

A *This thesis embarks on an unprecedented exploration into the complex field of a fortiori arguments, a critical form of deductive reasoning that draws conclusions upon comparison.*

B *The Groningen Meaning Bank (Basile et al, 2012) project annotates text by manually correcting the output of existing semantic parsers.*

We can discuss whether these should be "embarking" on anything or not, but it is, of course, not never the reasoning that draws the conclusions, but the people reasoning. (In a way, the drawing of conclusions is itself the reasoning.) And of course it is also not the project that annotates, but the project *is* the annotation (and it's instead people who perform the annotation work). One might argue that the above examples are simply examples of metonymy. There are certainly commonly accepted metonymies in science (statements such as "this paper presents" and "the results show"), which are conventionally used. However, metonymy is used unsystematically in the above examples (we could call it "metonymy overstretch") and it is likely that better formulations could be found. Such usage should therefore be flagged up to a writer. The aim of this project is to work towards a semantically-aware writing assistant that spots and corrects cases of authors' involuntary animism.

The task is to a) detect cases of animism and b) reformulate the sentence in such a way that only volitional entities (such as researchers) do any actions that require volitionality.

Research that is relevant to our task comes from several avenues. There is a large literature concerned with the question of modelling selectional preferences which is relevant here. There is also a literature about metaphors. There are lexical resources recording the semantics of SOME verbs with respect to volitionality, such as FrameNet and Levin's verb classes, which one could expand to cover more cases.

There is also a connection with the syntax of English. This is because if an author chooses a subjectless construction, this can pose a particular threat of introducing animism in otherwise harmless-looking text.

C *The article creation process on Wikipedia relies on the integration of different perspectives to develop a fair and accurate knowledge base.*

The purpose clause in C (starting with "to") is a subjectless construction that requires an implicit subject. That subject needs to be capable of volition. The only syntactically available subject, however, the process, has no volition, causing a semantic clash. If the sentence is rewritten, the infelicity disappears:

Our goal is to develop a fair and accurate knowledge base. We therefore developed an article creation process on Wikipedia that relies on the integration of different perspectives.

Case D presents a method clause combined with a passive.

D *By keeping track of the addition and removal of dispute tags over various iterations of an article, WikiDisputes was created in Chapter 4.*

This results in animism because WikiDisputes (a corpus) isn't the one keeping track. This case can be fixed by simply creating an active voice version of the sentence, putting the semantic agent "us" (the default agent in a scientific paper) back in the driver seat. *By keeping track of the addition and removal of dispute tags over various iterations of an article, we created WikiDisputes in Chapter 4.* (There is still something not quite right about this sentence, but this project won't touch that problem.)

This is a phenomenon that nobody has worked on before, to my knowledge. There is no benchmark and no training data. Only an evaluation corpus of around 250 sentences exists, along with fixed versions of the sentences. The work will therefore involve unsupervised and symbolic methods for the main workflow. A minimalist type of evaluation would use only on the test corpus, but if the project progresses well, its scale and impact could be increased by performing a human evaluation of the detected cases and their reformulations.

This is not a straight Machine Learning project where advanced ML techniques are applied to a standard NLP problem. Even if the lack of training data wasn't a problem, an end-to-end system wouldn't be desirable in this setting. This project is a first exploration of a hitherto unstudied problem. The student should study the phenomenon and possibility of correction, both linguistically and computationally.

This does not mean that no ML is involved in any of the component steps (it is very likely to be). For instance, it is OK to have some components that provide non-explanatory output. It would be extremely useful to develop a neural model of the volitionality of verbs, but this would be a subsystem.

There is also a version of this project that involves syntactically coordination -- talk to me if that one interests you more. I am also always interested in self-proposed project on the pattern of my previous projects, particularly if they involve the scientific literature, metaphor, discourse effects or argumentation.

Literature

- Haagsma et al. (2017) Detecting novel metaphor using selectional preference information.
- Levin (2003) English verb classes and alternations: A preliminary investigation
- Baker et al. (2003) The structure of the FrameNet database
- Quirk et al (1995) A Comprehensive Grammar of the English Language