



Alhub.org

## Interview with Fiona Anting Tan: Researching causal relations in text

by [Lucy Smith](https://aihub.org/author/lucysmith/)

01 February 2024



The [AAAI/SIGAI Doctoral Consortium](https://aaai.org/aaai-conference/dc-24-program/) provides an opportunity for a group of PhD students to discuss and explore their research interests and career objectives in an interdisciplinary workshop together with a panel of established researchers. This year, 30 students have been selected for this programme, and we'll be hearing from them over the course of the next few months. In this interview, we meet [Fiona Anting Tan](https://tanfiona.github.io/) and find out about her work on casual relations in text.

### Where are you studying for your PhD and what's the topic of your research?

I'm at the National University of Singapore and I'm in the fourth year of my PhD. The specific topic of my research is text mining for causal relations. That falls under the area of natural language processing and reasoning.

### Could you give us an overview of the research you've carried out so far during your PhD?


I focus on the causal relations that are expressed in text. My work is split into three parts: extraction, representation, and application.

In my earlier work I focused on extracting, and I created datasets for areas that were lacking, consolidated benchmarks, and trained and designed models that are well suited for the task of extracting causal relations. Causal relations themselves are more tricky than relation extraction because they tend to co-exist with temporal relations. When you are creating your datasets you need to ensure that you include tricky edge cases, and that you test them properly.

When it comes to representation, there are a lot of techniques. The most common one, if you have very clean relationships, are knowledge graphs. These enable you to extract relationships that you didn't observe in the original text. In other words, you can expose causal chains that are new to you. This helps in applications such as prediction, where in the original text they didn't think that that was a downstream effect, but then, because in another piece of news or document someone says that event B has led to event C, you can make the connection from A to B to C. That's one method. Another method is to work directly with embeddings.

In the last bucket of my research, I have been finding applications that could use this knowledge, specifically with summarization and prediction. Working on this aspect, I have done internships. One of them was with Panasonic, and [that work \[https://ojs.aaai.org/index.php/AAAI-SS/article/view/27476\]](https://ojs.aaai.org/index.php/AAAI-SS/article/view/27476) featured in the AAAI Summer Symposium. The whole idea there was that we wanted to extract causal knowledge from news and then let that inform the supply and demand planners in terms of the trends in the electronics supply chain industry. For them, the use case was summarization and prediction based on current events, recent trends, and predictions of future events. For example, what led to the fall in demand for EV batteries?

I'm still building on this idea of using causal knowledge in various applications. At the moment, I'm doing an internship at Amazon, based in Seattle, and I'm using a very similar framework: extracting causal relations, then finding applications for them.



Overview of methodology for constructing causal knowledge graphs from news. From [Constructing and Interpreting Causal Knowledge Graphs \[https://ojs.aaai.org/index.php/AAAI-SS/article/view/27476/27249\]](https://ojs.aaai.org/index.php/AAAI-SS/article/view/27476/27249).

**Have these internships helped you with your PhD?**

The internships have helped solidify my PhD thesis by giving me more applications to work on. Prior to that, the kind of problems I was working with were restricted to NLP benchmarks and open-source datasets. Now I can point to a particular use case, which makes my thesis a bit more concrete.

**Is there a result from your research that you are particularly proud of?**

I feel like my most important work was [creating the dataset called "Causal News Corpus" \[https://aclanthology.org/2022.lrec-1.246/\]](https://aclanthology.org/2022.lrec-1.246/), featured at LREC 2022, and the subsequent resource called RECESS, featured at IJCNLP-AAACL 2023. I got to work with so many people that I really respected and that was the first big project that I had to lead. That was the most significant point of growth for me as a PhD student.

**Would you like to build on your research after you have finished your PhD?**

One reason why we focus on causal relations is because they are very tricky to build into a model. It's a very important facet that you should test to see whether your model is "understanding". I think with the onset of large language models there is still scope for doing this kind of research, in terms of checking the robustness, and in terms of understanding real scenarios. I'm thinking about things like complex reports where there are certain causal connections that need to be made in order to draw a conclusion, say about a probable cause. I think that research so far has shown that large language models still can't really do planning which I think relates to the lack of ability to look forward in terms of events. Hence, I see the relevance of my work in understanding how language models reason and understand causal events even after I graduate.

**What was it that made you want to study AI?**

I was already doing data science during my undergraduate, which was actually in Econometrics. I was in London for three years. I really enjoyed the whole research process, particularly the experimental side of it. It felt natural to pursue more data-focused research, working with the data and describing what I see. I felt that I had very little experience with text specifically and I thought if I took a masters or PhD, I could build on this skill. I decided to do a PhD – I wanted the time to go deeper in text mining and natural language processing, which I saw had a huge potential for growth back in 2020. At that time, there were so many methods for structured data, with all the tools that were available out there, and I thought that the NLP field was relatively lacking at that point in time. I was interested in working on some ideas that I had.

Originally, when I first applied for the PhD, the ideas that I had were very simple, like "I just want clean text", or "I just want to make text more structured for a business use case". I guess, as I dived deeper, I found the causal relations problem very interesting, so I focused on that.

**What advice would you give to someone who's thinking of doing a PhD in the field?**


I would say that prior to applying you should read up on the field. Initially, it might be hard to understand the papers, but definitely browse through to see if that's something that you want to do. Reading papers also gives you some idea of the kind of things you should be achieving during a PhD.

I'd also recommend talking to people who are working in the field. I know in some countries you have to know who your supervisor is going to be before you apply. I didn't have to do that in Singapore, but I still reached out to some professors. That was helpful because I found out the areas that I needed to work on, and the areas in which I already had expertise, which gave me some confidence to apply.

**Could you tell us an interesting non-AI related fact about you?**

I got married during my PhD. I'm quite pleased that I managed to organize a wedding whilst doing my research at the same time.

In terms of hobbies, I enjoy netball. I've also gotten into yoga during my PhD. It definitely helped me to find mental space and quietness.

**About Fiona**

Fiona Anting Tan is a PhD student at the National University of Singapore and a recipient of the President's Graduate Fellowship (PGF). Her research interests span across natural language processing and reasoning, with a focus on extracting causal relationships from text for various applications. She has had multiple research internships and attachments with companies like Amazon, American Express, Panasonic and Changi Airport Group. She graduated from the London School of Economics and Political Science with first-class honors for BSc in Econometrics and Mathematical Economics as a Singapore-Industry Scholarship holder.

tags: [AAAI](https://aihub.org/tag/aaai/) [https://aihub.org/tag/aaai/], [AAAI Doctoral Consortium](https://aihub.org/tag/aaai-doctoral-consortium/) [https://aihub.org/tag/aaai-doctoral-consortium/], [AAAI2024](https://aihub.org/tag/aaai2024/) [https://aihub.org/tag/aaai2024/], [ACM SIGAI](https://aihub.org/tag/acm-sigai/) [https://aihub.org/tag/acm-sigai/]



Lucy Smith , Managing Editor for Alhub.