

Rules

- ▶ No participant may record the session
- ▶ Abide by the ACL Anti-Harassment Policy
https://www.aclweb.org/adminwiki/index.php?title=Anti-Harassment_Policy
- ▶ Questions: Feel free to interrupt at any time!

Introductions

Please introduce yourself

- ▶ Name and location
- ▶ Affiliation and role
- ▶ Research area/interests

Topic: Navigating research problems/directions in NLP

- ▶ I would be happy to help with other topics/questions, either at the end of today's session, or get in touch: robvanderger@live.nl
- ▶ Note that this is a process that is subjective, my approach might not work universally, but I hope to give some valuable suggestions
- ▶ You can ask others (i.e. at posters), how did you find this topic/idea? every paper has a story behind it!

Suggested questions:


- ▶ How can I evaluate the long-term impact of an idea/project?
- ▶ What do you think are the most interesting problems that NLP researchers should be focusing on right now?
- ▶ How did you choose your current research direction? What motivates you to work on this problem?
- ▶ As a PhD student, how did you balance short-term (publications/funding obligations) and long-term (building a research identity) research goals?


What I was afraid of:



Reality:







Starting directions:

- ▶ Interest in task/benchmark X
- ▶ Interest in a setup with source data X and target data Y
- ▶ Interest in a specific model/have an idea for an improved model
- ▶ I am interested in X and my colleague in Y
- ▶ How/why/when does model X perform well?
- ▶ Reading interesting papers

I am interested in task/benchmark X?

- ▶ Are you targeting a very popular benchmark (WSJ/Glue)?, or a less popular one (i.e. constituency parsing of tweets)?
- ▶ Read the previous work on this task, and try to find what they leave out
- ▶ Qualitative error analysis!
- ▶ Can we exploit more information (or less?)

I am interested in a setup with source data X and target data Y

- ▶ Is the setup unique? make sure you have a strong baseline!
- ▶ What are existing models doing for similar setups
- ▶ How can you improve?, why do you think this is viable?, can you do preliminary testing?

I am interested in a specific model/have an idea for an improved model

- ▶ What would the ideal setup/data be for testing this model?
- ▶ Does this data exist?
- ▶ Is this an interesting extension?

I am interested in X and my colleague in Y

- ▶ Great!, can these be combined?, could this lead to interesting findings?
- ▶ Could be multi-task, pipeline, or ensembling

How/why/when does model X perform well?

- ▶ Begin with some intuition/hypothesis
- ▶ Make sure you can answer your research question!
- ▶ Is this relevant for other people?

Reading interesting papers

- ▶ Try to figure out what is missing
 - ▶ Besides answering research questions, papers often raise them!
- ▶ **Question assumptions!**
- ▶ What do the takeaways mean for the larger (sub)field?
- ▶ The most interesting follow-ups are often not in the future work description!

Some more general advice if you have problems with finding a unique direction:

- ▶ Try to find a niche, change the setup (domain/languages, training size, etc.)
- ▶ Try to relax assumptions of previous work
- ▶ Try to find unanswered aspects in previous work
- ▶ Error analysis!
- ▶ Don't be afraid to reach out to colleagues!

How to know if a research direction is worth following:

- ▶ It interests YOU
- ▶ Start out with baseline/upperbound wherever possible!, so that you can do preliminary experiments
- ▶ You can bring something novel


”Case study”: Lexical normalization (2015)

social	ppl	r	troublesome	.
IV	OOV	OOV	IV	IV
	people	are		

- ▶ Default benchmark: LexNorm 519 annotated sentences (Larger training set soon thereafter)
- ▶ Most models first generate candidates and then rank them (based on some distance metric)
- ▶ What are potential follow ups?

- ▶ Detect which words need normalization
- ▶ Include tokenization
- ▶ Fix capitalization
- ▶ Think about scope of task more general, what is included (lol, hahaha), and can we make custom models?, yes by annotating categories
- ▶ Go multi-lingual
- ▶ Go cross-domain
- ▶ What can we use this for?
- ▶ How can we evaluate performance in depth?

Viability Dependency Parsing as Sequence Labeling (Strzyz et al, 2019):



- ▶ We can try different "encoding" strategies, or improve the sequence labeler (less interesting)
- ▶ Can this be used for other tasks?:
 - ▶ (Biological) event extraction
 - ▶ Coreference resolution
 - ▶ Enhanced dependency parsing
 - ▶ Constituency parsing
 - ▶ Semantic parsing?
 - ▶ ...

Do you have any other questions about your own research directions?

- ▶ Concrete examples are welcome!

Any other questions? (robvanderger@live.nl)