Mentoring Session at ACL 2020 on Effective Communication, Long Term Career Planning in NLP, and Changing Career to NLP

Mentors: [Ahmed Awadallah](), [Pradeep Dasigi]()
Original doc:
https://docs.google.com/document/d/1UM1-UZ9OfwhvqoNZvetoJB7K6XUXc0272wvIKYKj92w/edit
[disclaimer: this pdf version is only for archive purposes.]

Q. How do I reach out to other NLP researchers when I am the only person at my company doing NLP?
      Follow up Q. Should I be worried about sharing ideas that are not well-formed yet?
Ahmed: You can invite students at universities for internships so that you can learn from them. Also, look for universities and strong research groups close to you and reach out to them. Try to establish a good connection before you share ideas and seek feedback. Don't worry about sharing ideas. Ideas are usually overrated. Only 10 to 20% of a successful project depends on a good idea, the remaining being how it is executed, and quite often the idea itself evolves over the course of a project.

Pradeep: Agree with Ahmed. If you're worried about other people stealing your ideas, I don't think that's a huge concern. Most researchers have their own agendas and are too busy already.

Q. How do I choose which topics to learn? Easy to learn vs. more important
      Follow up Q. How do I balance breadth-first vs depth-first learning?
Pradeep: It might be useful to think of what problems you're trying to solve and identify the topics you need to learn for solving them. Regarding breadth-first vs depth-first, it's useful to do an in-depth literature review when you're starting on a new project, and be informed of the advances in the field by reading a few papers outside your core research topic on a regular basis.

Ahmed: It's important to define what your area of interest is, and when you see a new paper, be able to identify how close it is to your interests. Spend enough time to learn the topics relevant to your area. For the other topics, allot a limited amount of time and don't waste too much of your time there.

Q. As a newbie, which deep learning platform should I learn, Pytorch or Tensorflow?
      Follow up Q. I heard that Tensorflow makes it easier to deploy code in production, compared to Pytorch. Is that true?

Pradeep: Most projects typically build on top of others. So the most important factor here is which framework the projects relevant to you are implemented in. I believe Pytorch has an edge over Tensorflow at least in writing research code.

Ahmed: I agree that many more NLP projects use Pytorch than Tensorflow. In terms of deployment in production, it used to be the case that Tensrflow had a strong advantage, but it is changing, thanks to frameworks like [Onnx](#).

Q. My current job requires using basic NLP tools, but not necessarily ones that require deep learning expertise. How do I make myself a competitive candidate for my next job search?
Ahmed:

Pradeep: Some companies now have open-source tools or libraries. Consider contributing to them to familiarize yourself with those technologies and also with the development teams at those companies.

Q. How do I communicate what NLP can do to my clients who are not familiar with the technology?
Ahmed: This is a serious problem, and quite often you do not have people who can both understand the client's needs and what NLP/ML can do. It is important to hire such people who are "bilingual" in that sense so that they can bridge the gap.

Pradeep: If possible, it might help to build demos of the services you offer so that your clients can get a sense of what the strengths and weaknesses of the technologies are.

Q. What are your favorite papers either at this conference or in general?
Ahmed: It's hard to pick one paper. But at this conference, I really enjoyed Kathy McKeown's keynote talk. She provided a great overview, and very useful advice on how to choose problems to work on, etc.

Pradeep: I agree with Ahmed about Kathy's talk. In terms of papers, there isn't a single one, but there are a couple of recent trends that I think are quite exciting: The community has been focusing on building newer datasets and thinking about testing the phenomena that have not been tested earlier. That's great. We've also been thinking about making NLP useful by building practical benchmarks and not just solving academic research problems. One example of such a benchmark is the [TechQA dataset](#) from IBM.

Q. Is it better to choose statistical NLP research path or Deep Learning research path for future career?