How to write "statements of purpose"

Please stop me and ask questions throughout the talk!!!

slido.com #466164

https://app.sli.do/event/jbyxrcsa

Boaz Barak

Harvard John A. Paulson School of Engineering and Applied Sciences

LEARNING THEORY ALLIANCE

COLT 2021 Virtual Mentoring Workshop

Who am I

- Professor, Princeton University 2005-2011
- Researcher, Microsoft Research New England, 2011-2016
- Professor, Harvard University, 2016-

$\Rightarrow$ read a lot of statements

General disclaimer

Advice always comes with

This talk

Statements of purpose / research statements for grad school applications and fellowships.

Graduate admission from faculty POV

1) How likely is student to succeed in my area of research?

- General qualities: drive/grit, curiosity, creativity…
- Specific skills/experience: math, programming, ML, background in X (e.g. physics/econ/biology/…)

2) How good of a fit is the student with me?

Are they really interested in working with me / in my area?

And are they interested in this for the right reasons?

Evaluate answers using reference letters, transcript, and statement of purpose.

Goal of SOP: Help answer above Q's, complementing or highlighting information from reference letters and transcript.

Writing a statement of purpose

Who is your audience?

Faculty at school X that work in area Y

What do you want to convey to them?

Why you want to study Y, and what makes you good fit for Y-group at X

General tips

- Keep it short (e.g. two pages).
- DON'T start with *"ever since I was 3, I loved robots"*
- Talk about what you did, showing you understand the science.
  Not "I implemented 7 features using 11 libraries" or "I proved Lemma 2.3" but "project was important because of X, and my role in it was Y"
- Sometimes: makes sense to talk about difficulties/challenges you've overcome.
- Don't be obnoxious but also not overly humble.

University-specific part of SOP

- Do your research: know who works there, and what they are working on now. DBLP page / google scholar
- Don't have to be super-specific but also not be super general (e.g., "I'd like to work with Boaz Barak, Eddie Kohler, or Elena Glassman")
- No need to go overboard with flattery ☺

It's not a zero sum game

Graduate school has huge opportunity cost for you.

⇒ you want to make sure the fit is good.

⇒ not in your interests to get in at all costs!

Great example: Naama Ben David

B.Sc, University of Toronto
Ph.D, CMU

One page SOP

Introduction: short & to the point

context

> I am an undergraduate student at the University of Toronto, about to graduate with a B.Sc. in Computer Science and a minor in Linguistics. I have been involved in research in both computational linguistics and software engineering, and each of my projects has produced a result for publication. However, I find that my interests lie in theoretical computer science; I am currently working as a research assistant in the theory of distributed computing, and will be continuing this work in the next few months.

- Successful at research
- Interests shifted (absolutely OK!)

Describing research project

During the summer of 2014, I worked with Professor Chechik and her team in the field of software engineering. For the past few years, Professor Chechik's team has been working on introducing a concept of uncertainty to early stages of software design. This allows for developers to postpone some design decisions while continuing to work. Over the course of the summer, I was involved in creating a tool that handles uncertainty in software models; it allows for uncertainty to be expressed more easily, and has the ability to visualize and refine uncertain models based on given properties. We have submitted our developments on the tool for publication.

Understands overall goal of project

Describes clearly her role and why it's important

Interest in area of research

Earlier this month, I started working as a research assistant for Professor Hadzilacos and Professor Toueg in the theory of distributed computing. My interest in the field began in an independent study course that I took in the fall semester. This experience offered me a peek into the fascinating world of research in computer science theory. In the upcoming semester, we will be examining a class of shared objects, called Deterministic Abortable objects, to see how they compare to their wait-free counterparts. Linearizable wait-free objects are heavily studied in distributed computing, as they display powerful and attractive behavior. However, they are often expensive to implement. Deterministic Abortable objects exhibit slightly weaker, though still useful properties, and recent work has shown that they might offer a cheaper alternative.

A little bit of "passion" statement

Demonstrates technical understanding of area

Extra useful information

During my undergraduate studies, I held teaching assistant positions in two different courses: Data Structures and Software Design, and I worked as a private tutor in an introductory Theory of Computing course. I find that I enjoy teaching – both the interaction with other students, and the opportunity to deepen my understanding of the material. In addition, I gained software development experience during an internship at BlackBerry, where I received an "Outstanding" ranking in my final worker evaluation – the highest possible ranking, only given out in special circumstances.

TA in relevant courses

Highlight achievement that may not "jump out" from transcripts / recs.

What's missing?

- More theory-specific experience (unavoidable)
- Add more details on particular faculty (could have been added)

Questions / discussions / other examples

Fellowship research statement

Junior fellowships (undergrad / beginning grad):

Share similarities SOP but many fellowship-dependent details

Senior fellowships (2+ year of grad school):

More in common with job-search research statement than grad admission statement of purpose.

At this point you should have clear research area, vision for its future direction, and already have done work in it.

Naama Ben-David research statement

Theoretical Foundations for Practical Concurrent Systems

Naama Ben-David
Computer Science Department, Carnegie Mellon University

The theory and practice of computer science are meant to go hand in hand. Interesting theoretical problems stem from the requirements and realities of practice, and well-developed theories can lead to the design of more efficient and reliable technology. Unfortunately, in many areas of computer science, including the study of concurrent systems, theory and practice are still far apart. Virtually all modern machines, from our laptops to our phones, have multiple cores. It is thus of the utmost importance to have a solid theoretical understanding of concurrent systems to guide our design of practical programs.

The behavior of concurrent shared memory is difficult to predict. There are many factors that may affect the performance of a given algorithm; from the number of processors available for use, to the structure of the memory hierarchy, and even to how busy the machine is. It is fair to assume that every time a concurrent algorithm runs, its behavior will be different. To reason about such a chaotic setting, theoreticians often abstract away these concerns by considering an all-powerful adversarial scheduler that determines the order of events in the system. A bound or a proof of correctness must then hold under any schedule that the adversary chooses. In this way, we can guarantee that any real execution will be at least as good as the theoretical bounds.

However, as it turns out, the theoretical guarantees that can be derived under such an adversarial model are too pessimistic, and often fail to accurately represent observed behavior. For example, a common class of concurrent algorithms, known as lock-free algorithms, provides system wide progress, but cannot guarantee that any specific process will succeed in its operations. Interestingly, these are often the algorithms of choice in practice; despite their weak theoretical progress guarantees, they perform quite well in the majority of systems, displaying quick progress for all participating processes.

My research aims to address this discrepancy, and more generally, to provide a solid theoretical foundation for the practice of concurrent systems. A good theoretical model should yield results that are relevant to practice, but must also be simple enough to work with, and general enough to be applicable to many situations. Ideally, a theoretical model can provide significant insight into an algorithm's behavior regardless of which machine it is run on. Furthermore, as technologies develop, the model should remain relevant even in the face of change. Throughout this statement, I will use examples from my work at Carnegie Mellon University, as well as the VMware Research Group, to discuss approaches to modeling concurrent systems that I believe are promising and can lead to important new discoveries.

Overview of area & vision, followed by concrete research directions.

For each one:

1) What I've done before
2) What are my next steps

Questions / discussions / other examples

Thanks!

Naama Ben-David, Kira Goldner, Sam Hopkins, Thodoris Lykouris, Siddharth Prasad, Colin White, Elen Vitercik, Xinyu Wu.

Harvard colleagues: Nada Amin, Finale Doshi-Velez, Krzystof Gajos, Eddie Kohler, Harry Lewis, Michael Mitzenmacher