CHICAGO BOOTH
The Center for Applied Artificial Intelligence

Applied Machine Learning Pre-Doctoral Fellows Program
Chicago Booth | Center for Applied AI
Professor Sendhil Mullainathan, Director

Are you interested in economic or social science applications of machine learning?

Are you contemplating a PhD in economics, social sciences, computer science, or a related field, but want to take two years to develop research skills?

You might be a good fit for our pre-doc program.

Like with any research assistantship job, you will get hands-on experience.

But our predoc is not just a job. We recognize that the time between degrees is invaluable, so we've created a learning experience to maximize that time. For example, we have a series of special seminars focused on learning the tricks-of-the-trade from faculty. We have dedicated advising to help in choosing the right courses. Our faculty will work 1:1 with you on interesting projects, and you'll learn from a cohort of peers in weekly lab meetings.

The Kind of Work We Do

Chicago Booth has faculty at the cutting edge of empirical science across many areas – economics, finance, social policy, behavioral science. Many of these faculty are incorporating machine learning, AI, and other tools into their projects. Here are a few of the problems we are tackling:

- So many people are needlessly in jail. Can we build a tool that helps cities radically de-populate their jails, without having any noticeable effect on public safety? See this paper for previous work in this area.
- Finance is all about predictability. No wonder machine learning is beginning to spread. But finance has a strong empirical and theoretical foundation – how do we integrate machine learning into that foundation? Does it help characterize asset prices better than factor models do? Can we create better models of investors' behavior?
- Firms now have more data than ever on consumers. So much so that they can give each person their own price. What are the economic and social implications of such radically targeted pricing? Is it good for consumers or will they be exploited? See this paper for previous work on this topic.
- There is growing fear that algorithms might be biased. But how are we supposed to find bias? And what do we do when we do find it? Our recent paper in Science talks about how we did just that – for an algorithm in health care that affected lives of over a hundred million people.

The projects are diverse in scope and methods – from observational data to quasi-experimental methods to large-scale field experiments.

You will choose and be hired by a Booth faculty member, with whom you will work closely throughout the length of the program. In addition, you'll be surrounded by other predocs working on other projects, with whom you'll interact at weekly lab meetings and other events. You'll get the benefit of working

closely with a faculty member on one particular set of projects while learning about many others in a community of researchers.

Who you are:
You have a bachelor's or master's degree. You have some interest in a PhD, enough to invest in or explore that interest, whether in economics, social science, data science, computer science, or other related fields.

Our experience has taught us that cross-cutting work requires "bilingual" data scientists. You need to excel at the latest computational tools. You also need to know the micro-econometric issues that arise with social science data. Our experience has also taught us that even the best universities are not producing bilingual people. So, you're a good fit if you know one language and are willing to commit to learn the other.

Are you *excellent* at one of these skills and *humble* enough to want to learn the others?

- Empirical micro-econometrics – do you know what a regression discontinuity is or what LATE is? More importantly have you had experience in class (or ideally on some project work) applying these ideas and seeing the subtleties of social data?
- Computer science / data science – do you know what LDA is or have you built a pipeline that tackled a big hairy dataset? Do you feel some expertise in either a single methodology (e.g. convolutional nets or generative models) or a single modality (such as images or language)?

What you will do:
The position will span the spectrum as any data science role does. In addition, we will view this as in part an apprenticeship to learn the trade-craft of research.

- Contribute to the design and implementation of an efficient and reproducible data processing pipeline.
- Build and rigorously evaluate statistical models using best practices of machine learning and statistical inference.
- Prepare project memos, summaries, presentations, reports, and other work products for dissemination to academic researchers, policymakers and other stakeholders, as needed.

What we will do for you:
As part of your development as a researcher, you will enjoy being part of a community of scholars learning and pursuing research together, as well as:

- Weekly joint lab meetings, so various research teams can collaborate and learn from each other.
- Interactive seminars with Booth and University faculty to investigate new research and possible challenges associated with the research trade.
- Individual academic advising to prepare you for your PhD career.
- Opportunities to connect with visiting experts and external collaborators.
- Provide opportunities for publication, including co-authorship with Booth faculty.

Some formalities

Competencies:

- Advanced knowledge of data science techniques OR applied micro-econometrics OR theoretical statistics, math, physics or related field.
- Strong initiative and a resourceful approach to problem solving and learning required.
- Ability to work independently and as part of a team in a fast-paced environment required.
- Sound critical thinking skills required.
- Strong attention to detail with superb analytical and organization skills required.
- Excellent written and verbal communication skills, with the ability to present data in a simple and straightforward way for non-technical audiences required.

Education, Experience, and Certifications:

- Bachelor's or master's degree in computer science, statistics, data science, economics or a closely related field required;

Technical Knowledge or Skills:

- Proficiency with statistical data analysis and machine learning using Python or R is required. The ability to work in both is preferred.

To apply:

- Upload your resume, cover letter, transcripts, and a sample of data work (class project, GitHub repo, blog, etc.) through Booth's online form: https://forms.gle/XmmtTKUHoq5PjejZ7
- Submit a formal application through the UChicago jobs portal: https://uchicago.wd5.myworkdayjobs.com/External/job/Hyde-Park-Campus/Research-Professional_JR07983

The University of Chicago is an Affirmative Action/Equal Opportunity/Disabled/Veterans Employer and does not discriminate on the basis of race, color, religion, sex, sexual orientation, gender identity, national or ethnic origin, age, status as an individual with a disability, protected veteran status, genetic information, or other protected classes under the law.