Panel: Trustworthy Data Science

Adam Doupé*
Arizona State University
doupe@asu.edu

ABSTRACT

Much of the research that our community publishes is based on data. However, an open question remains: are the results of data science trustworthy, and how can we increase our trust in data science? Accomplishing this goal is difficult, as we must trust the inputs, systems, and results of data science. This panel will discuss the current state of trustworthy data science, and explore possible technical, legal, and cultural solutions that can increase our trust in the input, systems, and results of data science.

CCS Concepts

•**General and reference** → **General conference proceedings;**

Keywords

Data Science; Panel; Trust

1. SUMMARY

Much of the research that we, as a community, publish is based on data. For this panel, we will call any research results based on data *data science*. An open question remains: are the results of data science trustworthy, and how can we increase our trust in data science? Accomplishing this goal is difficult, as we must trust the inputs, systems, and results of data.

Can we ensure that the input data sets that we are using are representative and unbiased, thus establishing trust in the inputs to our data science system? This problem is very difficult when we consider the large datasets that we work with, for instance the web, malware, or mobile applications. How can we ensure that the datasets that we use and base our research results on are unbiased and representative?

Many of our data science systems rely on machine learning to generate results, however recent work in adversarial machine learning has shown that machine learning classifiers can be evaded by manipulating the input data, and therefore the systems that we build can be untrustworthy. Can these systems be made resistant to an active adversary? What threat models are realistic and should be considered?

Finally, even if we fully trust the inputs and the systems, how can we trust the results of data science? Traditional academic disciplines rely on the scientific principle of reproducibility—however, this seems to be missing in our field, due to complex technical, legal, and cultural reasons. What are some techniques that we can use to share datasets and improve reproducibility of our results?

As more and more research is based on data, it is critical that we explore ways that we can increase the trust in our inputs, systems, and outputs of data science.

*Panel Moderator.

CODASPY'17 March 22-24, 2017, Scottsdale, AZ, USA

ACM ISBN 978-1-4503-4523-1/17/03.

DOI: http://dx.doi.org/10.1145/3029806.3044199

2. QUESTIONS TO CONSIDER

This panel will discuss the current state of trustworthy data science, and explore possible technical, legal, and cultural solutions that can increase our trust in the input, systems, and results of data science. Some broad questions that the panel will consider follows.

- How does the community view research that replicates other studies? Should we, as a community, encourage reproducible research and replication studies? If so, how can we actually achieve this goal?
- What prevents researchers from sharing data sets? What could be done to alleviate these problems? Are there technical research solutions that could prove useful?
- With many of our computing systems relying on machine learning algorithms, how much trust can we place in these systems? Can there be a trustworthy machine learning future, and if so how can we create that future?
- How does one know that a dataset is representative and unbiased? Are there techniques in collecting data that can improve the representativeness of the underlying population?
- How can we trust the output of research that is performed on datasets that the researchers are under legal obligations not to share (due to NDAs, privacy issues, etc.)?