

UNIVERSITY OF CHICAGO
DEPARTMENT OF COMPUTER SCIENCE

PRESENTS:

“Reliably Managing Unreliable Data”



Oliver Kennedy
University of Buffalo, SUNY

Abstract:

We've all seen news articles where data errors have severe implications on the real world, like a 12-year old being labeled a terrorist by "the system" and denied boarding on a flight with his parents. Heuristic guesses like this are often treated as facts because databases simply have no better way of representing them. In this talk, I outline our efforts to make uncertain data management easier through a probabilistic database system called Mimir. First, I will outline how rethinking uncertainty as a form of provenance metadata facilitates more natural and more efficient interactions with the database. Then, I will introduce Mimir's declarative approach to uncertainty in data, which allows it to decouple query processing from more mechanical concerns like how to present uncertainty in query results, or which algorithms to use.

(This work is supported by NSF Award ACI-1640864, NPS Award #N00244-16-1-0022, and gifts from Oracle)

Bio:

Oliver Kennedy has been an assistant professor at the University at Buffalo, SUNY since 2012. He earned his PhD from Cornell University in 2011. Oliver's Online Data Interactions (ODIn) lab operates at the intersection of database and programming language techniques. The ODIn lab is currently exploring issues related to the usability of uncertain data management systems, making index structures adapt to changing workloads, and modeling SQL workloads.

Wednesday, October 4, 2017
2:30 pm
Ryerson 251
Host: Mike Franklin