

# Converting data into knowledge

November 17 2014, by Jessica Stoller-Conrad

---



Yisong Yue, assistant professor of computing and mathematical sciences Credit: Lance Hayashida/Caltech Marketing and Communications

When a movie-streaming service recommends a new film you might like, sometimes that recommendation becomes a new favorite; other

times, the computer's suggestion really misses the mark. Yisong Yue, assistant professor of computing and mathematical sciences, is interested in how systems like these can better "learn" from human behavior as they turn raw data into actionable knowledge—a concept called machine learning.

Yue joined the Division of Engineering and Applied Science at Caltech in September after spending a year as a research scientist at Disney Research. Born in Beijing and raised in Chicago, Yue completed a bachelor's degree at the University of Illinois in 2005, a doctorate at Cornell University in 2010, and an appointment as a postdoctoral researcher at Carnegie-Mellon in 2013.

Recently he spoke with us about his research interests, his hobbies, and what he is looking forward to here at Caltech.

## **What is your main area of research?**

My main research interests are in [machine learning](#). Machine learning is the study of how computers can take [raw data](#) or annotated data and convert that into knowledge and actionable items, ideally in a fully automated way—because it's one thing to just have a lot of data, but it's another thing to have knowledge that you can derive from that data.

## **Is machine learning a general concept that can be applied to many different fields?**

That's right. Machine learning is becoming a more and more general tool as we become a more digital society. In the past, some of my research has been applied to applications such as data-driven animation, sports analytics, personalized recommender systems, and adaptive urban transportation systems.

## **What application of this work are you most excited about right now?**

This is tough because I'm excited about all of them, really, but if I had to just pick one, it would be human-in-the-loop machine learning. The idea is that although we would love to have computers that can derive knowledge from data in a fully automated way, oftentimes the problem is too difficult or it would take too long. So machine learning with humans in the loop acknowledges that we can learn from how humans behave in a system.

I think that we are entering a society where we depend on digital systems for basically everything we do. And that means we have an opportunity to learn from humans how to optimize our daily lives. Because human interaction with digital systems is so ubiquitous, I think learning with humans in the loop is a very compelling research agenda moving forward.

## **Can you give an example of humans-in-the-loop machine learning that we experience on a daily basis?**

One example of humans-in-the-loop that we experience fairly regularly is a personalized recommender system. Many websites have a recommendation system built into them, and the system would like to provide personalized recommendations to maximize feedback and engagement of that user with the system. However, when there is a brand-new user, the system doesn't really understand their interests. What the system can do is recommend some stuff and see if the user likes it or not, and their response—thumbs up, thumbs down, or whatever—is an indicator of the topics or content this user is interested in. You see this sort of closed loop between a machine learning system that's trying to learn how best to personalize to a user and a user that's using the system

and providing feedback on the fly.

## **You also mentioned animation. How is your work applied in that field?**

Before I came to Caltech, I spent one year as a research scientist at Disney Research. I worked on both [sports analytics](#) and data-driven animation. With regard to the animation, the basic idea is as follows: you take data about how humans talk in a natural sentence-speaking setting, and then you try to automatically generate natural lip movements or [facial movements](#) that correspond to the types of sentences that people would normally say. This is something that people at Disney Research have been working on for a while, so they have a lot of expertise here.

One of the things that you notice many times with animation is that either the character's lip movements are fairly unrealistic—like their mouths just open and close—or in the big-budget movies, it takes a team of artists to manually animate the character's lips. An interesting in-between technology would be to have fairly realistic automatically generated [lip movements](#) and facial movements to any type of sentence.

## **What are you looking forward to now that you're at Caltech?**

Here I have a combination of research independence, talented colleagues, and support for my research endeavor—and a great culture for intellectual curiosity.

It's such a tight-knit community. It's one of the smallest institutions that I'm familiar with, and what that implies is that basically everyone knows everyone else. The great thing about that is that if you have a question about something that you may not be so knowledgeable about, it's really

not that big of a deal to go down the block to talk to someone who works in that field, and you can get information and insight from that person.

## **Have you already begun collaborating with any of your new colleagues?**

I'm starting a collaboration with Professor Pietro Perona [Allen E. Puckett Professor of Electrical Engineering] from electrical engineering and Professor Frederick Eberhardt [Professor of Philosophy]. In that collaboration, we'll be addressing a problem that biologists and neuroscientists at Caltech face in assessing how genes affect behavior. These researchers modify the genes of animals—such as fruit flies—and then they video the animal's resulting behaviors. The problem is that researchers don't have time to manually inspect hours upon hours of video to find the particular behavior they're interested in. Professor Perona has been working on this challenge in the past few years, and I was recently brought in to become a part of this collaboration because I work on machine learning and big-data analysis.

The goal is to develop a way to take raw video data of animals under various conditions and try to automatically digest, process, and summarize the significant behaviors in that video data, such as an aggressive attack or attempt to mate.

## **Tell us a little bit about your background.**

It is a bit all over the place. I was born in Beijing. I moved to Chicago when I was fairly young, and I spent most of my childhood in Chicago and the surrounding areas. But my parents actually moved out of Chicago after my sister and I left for college, and so I really don't have any relatives or strong ties to Chicago anymore. Where I call home is ... I don't really know where I call home. I guess Pasadena is my home.

## **Do you have any hobbies outside of your research?**

I like hiking and photography, and I'm really excited to try some of the hiking trails in the area and to bring my camera and my tripod with me.

I have a few other hobbies, although I don't really have the time to do them as much now. I was part of an improv group in high school, and I did a fair amount of comedic acting. I wasn't very good at it, so it's not something I can really brag about, but it was fun. I am also an avid eSports fan. For instance, I love watching and playing StarCraft.

Provided by California Institute of Technology

Citation: Converting data into knowledge (2014, November 17) retrieved 18 April 2024 from <https://phys.org/news/2014-11-knowledge.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.