BUILDING INTELLIGENCE

From Numbers to Knowledge

For a person, the ability to recognize and apply knowledge and skills learned in previous tasks to new endeavors is a natural occurrence. After understanding how one card game works, it is easier to pick up another. For a computer, such learning is incredibly hard. This is the specialist domain of Prof Qiang Yang, an expert in data mining, artificial intelligence, machine learning, transfer learning and deep reinforcement learning.

Prof Yang has spent 20 years fathoming algorithms that seek to endow computers with similar capabilities to humans in retaining and reusing previously learned knowledge in order to "think" and "decide" how to extract information and patterns from the rivers of data flooding our digital age. Prof Yang and his team have improved the accuracy of computers' performance through devising versatile frameworks for such "transfer learning". He has developed Instance-based Transfer Learning, which uses individual instances to bridge different domains, and Heterogeneous Transfer Learning, where the computer learns in one knowledge domain (for example, text) then transfers what it has learned to a separate or more difficult domain (for example, images).

Prof Yang has made these frameworks open source, enabling other researchers and the field overall to develop at a faster pace. He was also the first to propose the use of transfer learning in collaborative filtering and recommender systems. Applications have ranged from early online advertising directed at users to improvements in recommendation systems, including a state-of-the-art recommendation system for ICT global giant Huawei's App Store.

Recent research at the WeChat-HKUST Joint Lab on Artificial Intelligence Technology (WHAT LAB), set up with Mainland China internet giant Tencent in 2015, has inspired a novel application to improve machine reading capabilities. Books, news articles, and blogs are used as input to train a machine learning model that can produce an abstract of such readability that it doesn't appear to have been written by a machine. The objective is to assist people with information overload on social networks or boost company productivity by enabling a computer to develop an abstract of a long report or integrate data and highlight the main points the reader needs to know. By reading books by the same author, the demonstration model designed by Prof Yang and his students has even written a high-quality novel of its own in the writer's style, taking just a few seconds to do so.

In improving such machine reading abilities, Prof Yang's team has become the first to integrate a reinforcement learning algorithm that leverages users' feedback related to positive comments on prior abstracts with transfer learning and deep learning (recurrent neural networks) to help the computer make a more intelligent decision on what abstract to generate. The innovation has improved the quality substantially. With information to hand quicker, it can also speed up report-writing as well as learning.

"The work at HKUST seeks to increase the knowledge we can get from data by making the process of moving from data source to understanding faster and more efficient, accurate and useful to people," Prof Yang said.

> "We are inventors, always thinking of how to use data in a new way"

PROF QIANG YANG
New Bright Professor of Engineering,
Head, Department of Computer Science and Engineering,
Director, HKUST Big Data Institute,
Inaugural Editor-in-Chief of
IEEE Transactions on Big Data

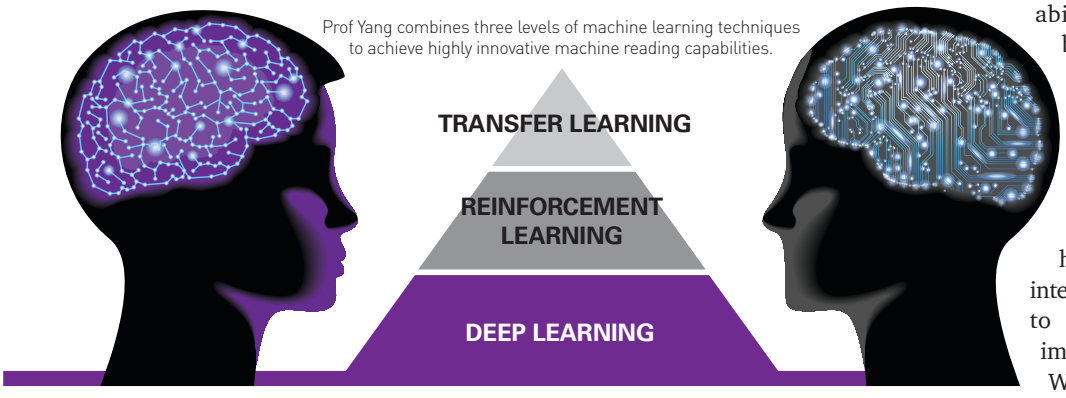

Prof Yang combines three levels of machine learning techniques to achieve highly innovative machine reading capabilities.

Humans have the ability to apply knowledge and skills learned in previous tasks (e.g. cycling) to new tasks (e.g. motorcycling). Prof Yang's transfer learning algorithms help computers to acquire the ability to retain and transfer knowledge from a source domain to a target domain.

> All these charts and visuals are like a movie. The actors are the same, but when you combine them together differently, you can tell a new story

PROF HUAMIN QU
Professor of Computer Science and Engineering

Seeing the Larger Picture

The power of the visual to impart information plays a hugely significant part in our lives, shaping our understanding of the world through "seeing with our own eyes" and through a variety of media, ranging from art over the ages to today's selfies and YouTube videos. Prof Huamin Qu and his team are leveraging such visual impact to mine the digital world of big data, by combining computational power to detect patterns and extract information from vast quantities of data with cutting-edge graphics and virtual reality techniques. In this way, they are uncovering previously unknown relationships, including those related to our own behavior. "We call it amplifying cognition," Prof Qu said.

One recent outcome of such data visualization is VisMOOC, the first visual analysis system for discerning e-learning behavior. The intuitive HKUST web app offers fine-grained analysis of video "clickstream" data, drawn from learners watching lectures for Massive Open Online Courses (MOOCs). VisMOOC pinpoints where learners play a section multiple times (indicating difficulty in comprehension), where they pause (to consider or take notes) and what they skip through (lack of interest or not challenging enough), among other details. Such clickstream data are matched with statistics from chat groups (forums), demographics, and grading for assignments and exams. Results are then provided in a novel visual form, labeled a "seek diagram".

Following VisMOOC's success, Prof Qu's team and collaborators are developing an open source platform with advanced visualization interfaces for individual institutions to do detailed analysis on e-learning behavior and course design.

On a wider scale, Prof Qu is integrating cutting-edge visualization with large-scale telecommunications data to create applications contributing to smart city understanding, for example, route planning, crowd management for transportation, analysis of visitor traffic for shopping centers at different times of the day, and even tracking of disease outbreaks. In 2016, such work saw Prof Qu receive the Distinguished Collaborator Award from Huawei's Noah's Ark Lab, the company's long-term, big-impact research lab. Working with WeChat, Mainland China's dominant messenger app, Prof Qu has also solved the challenging problem of visualizing the propagation of information over a large social media network, involving multiple attributes/dimensions and dynamic evolution. Analyzing users' behavior can assist in finding common communication patterns adopted by the public.

According to Prof Qu, a good visualization design must be effective in serving as a magnifying glass for what the data patterns show, aesthetic and intuitive. In addition, it should not be a pie chart or bar graph but a new visual form that carries interest for the viewer. Such integration of computational power in pattern recognition and mining and human expertise in visual pattern recognition, is a key area for further exploration, he noted. "Many real-world problems cannot be easily formulated as a computer algorithm so we need to keep humans in the loop."

Visual Analysis for Massive Open Online Courses (MOOCs)

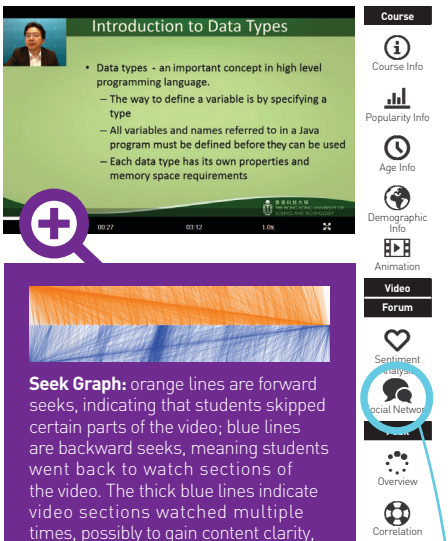

Seek Graph: orange lines are forward seeks, indicating that students skipped certain parts of the video; blue lines are backward seeks, meaning students went back to watch sections of the video. The thick blue lines indicate video sections watched multiple times, possibly to gain content clarity, thus alerting instructors to evaluate the course content and delivery.

Event Graph: showing six different types of clickstream data - play, pause, sought, rate change, stalled, and error - of the same course during the same time period but for students from different countries. By filtering demographic info on the dashboard, instructors can explore and compare online learning behaviors internationally.

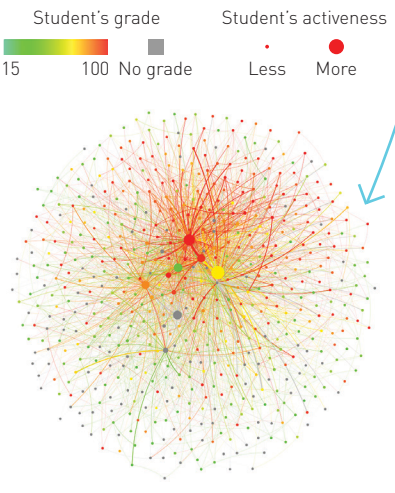

Forum Social Network: each dot represents a student. The size indicates the degree of activeness, and the color indicates the grade. The large gray dot shows that the student is active in the forum but does not achieve any score. The visualization gives valuable insight about students and their learning interactions, and provides MOOC instructors with useful information to improve course content.