STANFORD UNIVERSITY SCHOOL OF ENGINEERING / COMPUTER SCIENCE DEPARTMENT

Stanford

Artificial Intelligence Lab

FEBRUARY 2018

Stanford
ARTIFICIAL
INTELLIGENCE

Maneesh Agrawala
Professor, Computer Science; Director, the Brown Institute for Media Innovation

Computer graphics, human-computer interaction, visualization

To improve the effectiveness of media of all kinds

Serafim Batzoglou
Professor, Computer Science

Computational genomics

To analyze large-scale genomic data

Gill Bejerano
Associate Professor, Computer Science, Developmental Biology, Pediatrics (School of Medicine)

Computational genomics

To discover the relationship between the genotype and phenotype

Jeanette Bohg
Assistant Professor, Computer Science

Robotics, machine learning/deep learning, computer vision

To understand the fundamental principles of complex, sensorimotor behavior and how it can be generated on robots.

Emma Brunskill
Assistant Professor, Computer Science

Machine learning/deep learning

To advance the frontiers of reinforcement learning

Ron Dror
Associate Professor, Computer Science

Computational biology

To determine spatial structure and dynamics at the molecular and cellular levels

John Duchi
Assistant Professor, Electrical Engineering, Statistics

Machine learning, optimization and statistics

To understand the limits and application of statistics and AI

Stefano Ermon
Assistant Professor, Computer Science

Machine learning, statistics, sustainability

To solve impactful problems in sustainability

Ron Fedkiw
Professor, Computer Science

Computer graphics, computer vision, fluid dynamics, solid mechanics, biomechanics

To design computational algorithms

Michael Genesereth
Associate Professor, Computer Science

Computational logic

To solve problems in business, law, and game playing

Noah Goodman
Associate Professor, Psychology, Linguistics (courtesy), Computer Science (courtesy)

Computational psychology, machine learning, linguistics

To understand cognition, language and social behavior

Leo Guibas
Professor, Computer Science

Computer vision, computer graphics, geometry

To accurately describe the real physical 3D world

Thomas Icard
Assistant Professor of Philosophy and (by courtesy) of Computer Science

Machine learning/ deep learning, big data, knowledge bases, logic

To understand how reasoning works, and how we might like it to work

Dan Jurafsky
Professor, Linguistics, Computer Science

Computational linguistics

To solve problems and provide insights in behavioral and social sciences

Oussama Khatib
Professor, Computer Science

Robotics

To enable a new generation of robots that cooperate with humans and other robots in complex and unpredictable environments

Mykel Kochenderfer
Assistant Professor, Aeronautics and Astronautics, Computer Science (courtesy)

Machine learning, decision theory

To develop safe and efficient systems for air traffic, drones, and autonomous vehicles

Daphne Koller
Adjunct Professor, Computer Science; Chief Computing Officer, Calico

Machine learning

To solve real-world problems involving complexity and uncertainty

Anshul Kundaje
Assistant Professor, Genetics (School of Medicine), Computer Science
Computational biology, machine learning
To analyze all kinds of genomic and genetic data to understand gene regulation

James Landay
Anand Rajaraman and Venky Harinarayan Professor, School of Engineering; Professor, Computer Science
Human-computer interaction, NLP, autonomous vehicles
To design user-centered AI systems that augment and support people rather than replace them

Jure Leskovec
Associate Professor, Computer Science
Data mining, machine learning
To study the workings of large social and information networks

Fei-Fei Li
Associate Professor, Computer Science, Psychology (courtesy); Director, Stanford Artificial Intelligence Lab
Computer vision, machine learning, computational cognitive neuroscience
To help computers see better to help and work with humans

Percy Liang
Assistant Professor, Computer Science, Linguistics (courtesy)
Machine learning, natural language processing
To build systems that allow humans and computers to communicate

Tengyu Ma
Assistant Professor, Computer Science
Machine learning/ algorithms
To understand and develop machine learning algorithms

Chris Manning
Professor, Computer Science, Linguistics
Natural language processing, machine learning
To develop computers that can process, understand and generate human language

Andrew Ng
Adjunct Associate Professor, Computer Science
Machine learning
To solve problems in autonomous driving, robots, image analysis and language

Juan Carlos Niebles
Senior Research Scientist, Computer Science
Computer vision, machine learning
To allow computers to understand objects, scenes, activities and events in images and videos

Vijay Pande
Professor, Structural Biology, Computer Science (courtesy)
Statistical mechanics, Bayesian statistics, biophysics, biochemistry
To push the limits of simulation and statistics to answer questions in biophysics and biochemistry

Christopher Ré
Associate Professor, Computer Science
Database, machine learning, theory
To create the future of data systems for unstructured and structured data

Dorsa Sadigh
Assistant Professor, Computer Science, Electrical Engineering
Machine learning/deep learning, robotics, autonomous vehicles
To design algorithms for robots that safely interact with people

Ken Salisbury
Professor (Research), Computer Science
Robotics
To develop useful robots for surgery, imaging, haptics and personal assistance

Silvio Savarese
Associate Professor, Computer Science
Computer vision, robotics, geometry, machine learning
To push the limits of computer vision for object, scene and human behavior understanding and applications in social robots and autonomous vehicles

Sebastian Thrun
Adjunct Professor, Computer Science; CEO, Udacity
Robotics, machine learning
To improve robotics, autonomous vehicles, smart homes, healthcare and drones

Dan Yamins
Assistant Professor, Psychology, Computer Science (courtesy)
Computational neuroscience
To reverse engineer the human brain and build more effective AI systems

ARTIFICIAL INTELLIGENCE

Introduction

Dear Friends,

Welcome to the Stanford Artificial Intelligence Lab (SAIL)!

SAIL was founded by Prof. John McCarthy, one of the founding fathers of the field of AI. While the discipline of AI has transformed in many fundamental ways since its inception in 1950s, SAIL remains a proud leading intellectual hub for generations of scientists and engineers, an education mecca for students, and a center of excellence for cutting edge research work. With this brochure, we hope to share with you some the latest research and activities at SAIL.

Reflecting on the history of AI, the past fifty years are mostly what I call the "AI in vitro" times, during which most of the AI research was conducted in academic laboratories. This is the time that AI researchers laid the foundations for our fields, including the questions we are pursuing, the methodologies, the measurements and metrics, and the potential applications. AI grew from a small set of ideas to important areas such as robotics, natural language processing, computer vision, computational genomics and medicine, among others.

An important development in the field of AI was the emergence of machine learning around early 1980s. The recent convergence of modern computing hardware, big data and powerful machine learning algorithms has led to some breathtaking breakthroughs in many applications of AI, from speech recognition, to image recognition, to self-driving cars. We have now entered the time of "AI in vivo," where AI technologies are changing people's everyday lives as products of chatbots, photo organizers, assistive driving technology, and much more.

My colleagues at SAIL have been at the forefront of this AI revolution, just like they were the pioneers of establishing the AI field half a century ago. They continue to push the boundaries of fundamental research and innovative applications in many different areas of AI, from teaching computers to translate between languages,

Contents

Introduction 4
Machine Learning / Deep Learning 6
Natural Language Processing 10
Computer Vision 14
Robotics 18
Genomics 22
Autonomous Vehicles 26
AI Salon 30
Stanford AI4ALL 32
Faculty Focus 34

to developing a swimming robot that can explore shipwrecks deep under water following haptic instructions from an archeologist at the surface, to enabling computers to read satellite images that can help governments to monitor the state of regional economy, to creating a smart visual robot that can navigate among crowds in socially courteous ways that have only been unique to humans, or to developing algorithms to sort through millions of gene sequences to discover the genetic makeup of fatal diseases. And the list goes on.

Equally important to the mission of technology innovation, we believe in the education of the next generation technologists. As part of the Stanford Computer Science Department, SAIL faculty and researchers have been teaching students across the Stanford campus and from outside. Computer Science has become the largest undergraduate major at Stanford since 2015, and many of them have taken AI classes. In the various research groups of SAIL, undergraduates, including women and under-represented minority students, are often working alongside researchers on the most cutting edge research projects. In this brochure, we highlight two unique programs at SAIL that are beloved by the students. One is AI Salon, the regular SAIL-wide discussion. The other is SAILORS (Stanford AI Lab OutReach Summers), a first-of-its-kind AI summer camp targeted for high school girls focusing on humanistic AI.

This is a historical time for AI researchers, technologists, educators and students. There has never been so much excitement and hope for the potential promises of AI. But equally important, there has never been so much need to create benevolent AI technologies and to educate humanistic AI technologists for our world. We have come a long way, but the journey ahead of us is still long. None of today's intelligence machines come close to the breadth and depth of human intelligence. So all of us at SAIL are striving to build better algorithms and machines that will help humans to live better, safer, more productively and healthier.

Sincerely yours,

Fei-Fei Li
Director, Stanford Artificial Intelligence Lab
Associate Professor, Computer Science Department

It's a wonderful time to be at SAIL.

Machine Learning / Deep Learning

Teaching machines to learn has become a bedrock discipline of AI in recent years. Since 2010, it has moved out of academia, where it has been explored for decades, and into industry where it is starting to remake human-computer relationships.

Deep learning, a key sub area of machine learning, is transforming tasks that long stymied researchers. For instance, speech recognition used to be based on algorithms created by smart human programmers. The algorithms told computers to try to match patterns representing specific phonemes in digitized recordings. But in the last five years, as computers became more powerful, it became practical to build very large learning tools called deep neural networks. A neural network iteratively processes data, through layers of nodes. After listening to thousands of recordings and reading the corresponding transcripts, the computer learns for itself what digitized sound waves correspond to the word "Trump."

The key to machine learning is data—lots of it. If there is lots of annotated data that can be examined by a computer, the computer can learn to identify and understand it. Captioned photos, spoken language with transcripts, text, videos and genomes all

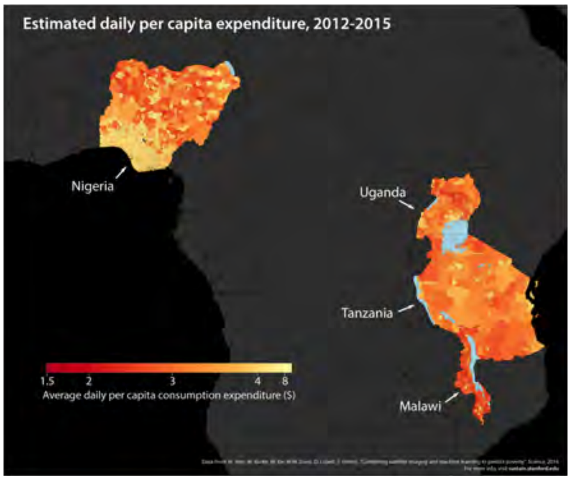

provide rich libraries of information that computers can understand.

The deep learning revolution took off after 2009. One spark occurred when Stanford researcher Fei-Fei Li released ImageNet, a free database of 14 million images that had been labeled by tens of thousands of Amazon Mechanical Turk workers. When AI researchers started using ImageNet to train neural networks to catalog photos

SAIL faculty working in machine learning, deep learning, big data, knowledge base, and logic include:
Jeannette Bohg, Emma Brunskill, John Duchi, Stefano Ermon, Michael Genesereth, Noah Goodman, Leonidas Guibas, Thomas Icard, Mykel Kochenderfer, Daphne Koller, Jure Leskovec, Fei-Fei Li, Percy Liang, Tengyu Ma, Chris Manning, Andrew Ng, Christopher Ré, Dorsa Sadigh, Sebastian Thrun, Dan Yamins.

and annotate their contents, machine learning exploded.

Machine learning replaces the complexity of writing algorithms that cover every eventuality with the complexity of processing lots of data. Processing the data can be a simpler task. With sufficient computing resources, the machines can learn much faster than programmers can teach them.

Machine learning can be applied to tasks that seem far afield. For example, Stanford researcher Stefano Ermon looks for societal problems that can be addressed with machine learning techniques. Identifying poverty levels in underdeveloped countries is a difficult but important task for governments and poverty fighting organizations. Ground level surveys are expensive and difficult. Financial data are scarce.

Ermon obtained high resolution satellite images of villages in Africa where asset levels had been documented by surveys. He used those to train a machine learning system. The system was able to spot

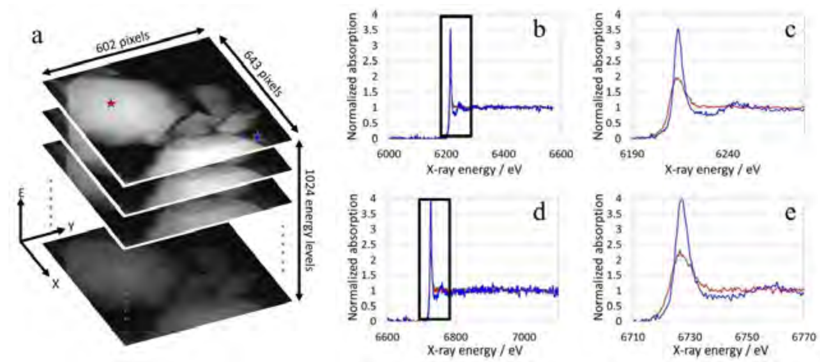

features like roads, plowed fields and metal roofs as signs of relative wealth. After validating the program's conclusions, Ermon published his results and made the program available to governments and NGOs, where it is starting to be used. Because satellite images cover every corner of the world and are frequently updated, agencies can even study areas like Somalia where warfare makes human surveys too dangerous.

He has used similar techniques to look at growth in farm fields. Other applications involve energy and the environment. He hopes to use machine learning to figure out ways to maximize battery life in electric cars. Other researchers are trying to crowd-source sufficient data to get new insights.

Machine Learning / Deep Learning

One target: understanding bird migration and species distribution by getting amateur bird watchers to report their observations. Ermon thinks machine learning could also provide guidelines on how to manage fisheries to benefit fishermen while preserving the population stock.

Still, many humans are likely to be skeptical of relying on machines that have essentially trained themselves. When neural networks learn things on their own, they sometimes err. In 2015, an automated labeling feature of Google Photos mistakenly identified some black people as “gorillas.” Google says it has tweaked the feature, but it’s impossible to predict other unexpected results.

In some cases, people will undoubtedly insist on a full understanding of the machine-learning process. As AI guides robots, cars and planes, if an accident occurs, and engineers can’t reconstruct the decision-making process, liability lawyers will have a field day.

Machine learning will be a core component of solving the next generation of AI problems. Robotics, natural language conversations and understanding videos are all tasks that require machine learning, probably augmented by other techniques that are still under development.

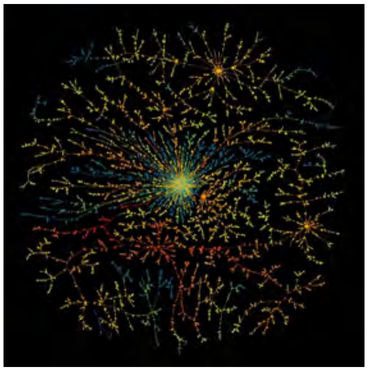

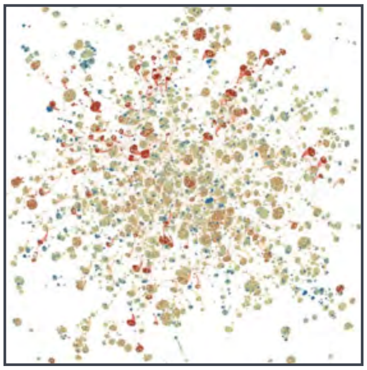

Prof. Kochenderfer's students use partially observable Markov decision processes (POMDPs) to represent problems that involve decision making under uncertainty. Techniques developed by his students are being used to derive optimal control strategies for automatically finding GPS jammers that could interfere with air traffic control.

Machine learning will be a core component of solving the next generation of AI problems.

Natural Language Processing

Artificial intelligence has made dramatic progress in understanding speech. Computers commonly make transcripts and provide phrase-by-phrase translations. But understanding human language is a thornier problem.

Natural language processing is one of the liveliest areas of AI research, in part because it addresses so deeply the issue of what makes humans human. Computers can listen to speech and learn the meanings of individual words and phrases. Processing speech is susceptible to using neural networks that train themselves to recognize the streams of bits that represent spoken words. But really understanding full sentences in context or having a dialogue is still a daunting task. AI methods pioneered at Stanford have made it possible to extract the grammatical structure of a sentence with high accuracy.

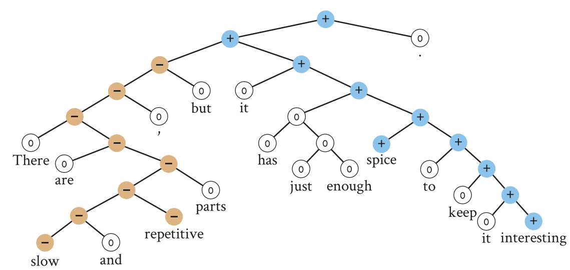

Sentiment tree. Jason Chuang and Richard Socher

Nevertheless, computers are getting better at natural language. AI has made huge progress in improving language translation using deep learning techniques and neural networks. By looking at hundreds of thousands or even millions of documents in English that have been translated to Chinese, a neural network can learn how to translate a full sentence. Stanford NLP Group founder Chris Manning and his students have helped advance this new Neural Machine Translation paradigm.

By using large databases of newspaper articles, neural networks can learn to generate summaries or headlines for articles. That can make it easier for humans to scan search results and find what they need.

Researchers are also applying AI techniques to understand the social meanings behind written and spoken language, people's opinions, attitudes, and emotions. For instance, Stanford researcher Dan Jurafsky works on sentiment analysis, to understand

SAIL faculty working in natural language processing include:
Noah Goodman, Dan Jurafsky, James Landay, Percy Liang, Chris Manning.

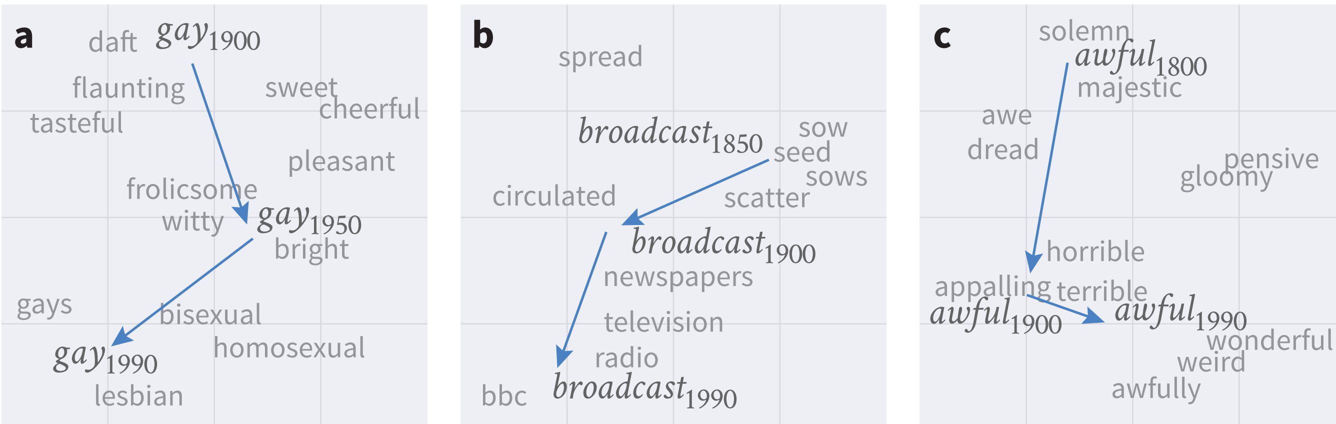

how groups of people feel about a subject discussed on social media like Reddit or Twitter. One challenge is that words have sharply different connotations in different contexts. For example, he notes, “soft” is a good thing when people are discussing puppies; it’s negative when discussing football players.

Jurafsky and other researchers developed a lexicon of accurately labeled sentiment words by identifying a few positive and negative words that are roughly consistent across hundreds of domains. Then they used those seed words to let their algorithm create a much larger lexicon of related words specific to each domain. Creating such lexicons is extremely time-consuming for humans, and it has been a barrier to progress in sentiment analysis. One interesting byproduct of the technology is that it has made it practical to do sentiment analysis on historical documents like newspapers and letters. That is tricky because words’ meanings change over time. “Awful” once was a positive word, connoting awe and majesty, but today it has a negative sentiment. Furthermore, the sentiment of a whole sentence is a complex function of the meanings of the individual words. Chris Manning has shown how to deal with these contextual aspects of sentiment by using tree-structured neural networks to combine the sentiment meanings of individual words.

Natural Language Processing

Analyzing natural language can give new insights in the field of computational social science. The language people use about everyday culture like food can tell social scientists a lot. Examining one million Yelp reviews of restaurants, Jurafsky found that disappointed diners write in the past tense and use words similar to those used by people who have been traumatized. Looking at online menus from restaurants, he found that words used in meal descriptions are a strong indicator of price. For example, “fresh” is used in downscale menus; in expensive restaurants it’s assumed the food is fresh.

One goal of natural language processing is using insights to improve the evaluation and training of humans. Stanford professor Jure Leskovec has applied natural language processing to online chats at a crisis hotline, to learn what successful counselors do to aid troubled callers.

Another goal of natural language researchers is to make it possible for machines to answer questions and carry on conversations with humans. While computers have long been able to fool some people into thinking they are conversing, actually working reliably with people remains an elusive goal.

Stanford AI researcher Percy Liang says “language is very context dependent. It’s about the connection between language and the world.” For example, if someone says “the election” the listener understands what election and probably what opinion the speaker has. A computer probably doesn’t. “So much of their inability to understand is that they don’t have the world context,” Liang says.

Industry is increasingly unveiling chatbots that answer questions in specific domains from obtaining visas to picking a dress. But carrying on conversations with machines on unspecified subjects often causes frustration for humans. Liang notes that the goal is to create a personal assistant that can respond to under-specified requests like “find me a flight,” by incorporating understanding of the user’s past preferences for non-stop flights, and avoiding the “red-eye” flight as well as showing the lowest price.

Liang says having a general conversation is “much more difficult. We’re not there yet.” But, “we’re making good progress” toward solving the conversation problems for specific domains like travel.

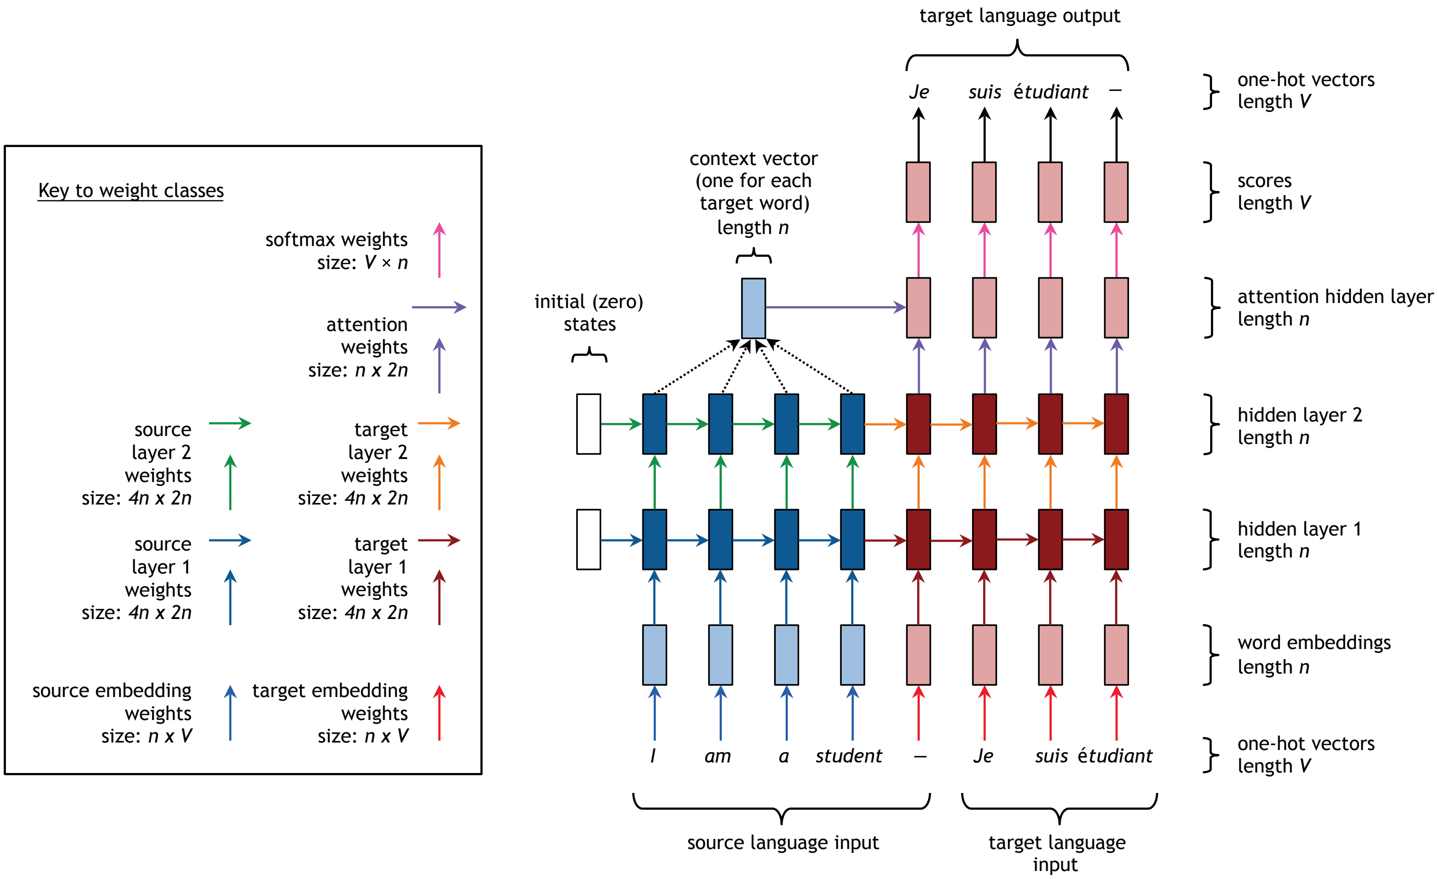

Analyzing natural language can give new insights in the field of computational social science.

COMPUTER VISION

Computer Vision

Computer vision has become one of the most fruitful areas of AI research and development in the past decade. And its growing maturity is opening up new pathways to intelligence in fields like robotics, machine learning and intelligent vehicles.

Humans are producing and storing an exploding amount of visual information from visual sensors like mobile phones or surveillance cameras that take beautifully detailed photos and videos with sound. But computers continue to struggle to make sense of it all. Recording the series of numbers that represent pixels in a digital photograph doesn't enable a computer to understand their meaning.

Humans look with their eyes, but they see with their brains. Computer vision lets computers see by processing the information in individual pixels and putting it together in the computer to create concepts. Researchers need to teach computers to see.

Computer vision took a great leap forward in 2009 when researchers led by Fei-Fei Li, the head of Stanford's Vision Lab, organized ImageNet, a database of 14 million photographs of tens of thousands of types of objects that were labeled by nearly 50,000 online human workers through the Amazon Mechanical Turk platform. By training neural networks using ImageNet, researchers developed the ability to identify objects almost as well as humans can.

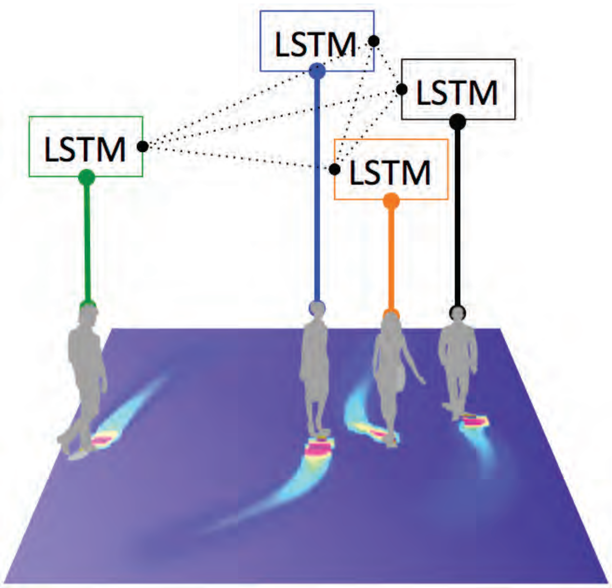

SAIL faculty working in computer vision, haptics, and sensing include:
Maneesh Agrawala, Jeannette Bohg, Ron Fedkiw, Leonidas Guibas, Oussama Khatib, Fei-Fei Li, Juan Carlos Niebles, Ken Salisbury, Silvio Savarese, Sebastian Thrun, Dan Yamins.

Neural networks have remade the science of computer vision. Just as children learn to identify objects by seeing hundreds of millions of images of the real world, neural networks are trained by examining labeled images from the Internet. The neural networks used to develop machine vision have 24 million nodes and almost 15 billion connections.

The next step in evolving computer vision was to identify multiple objects in an image. That often involves interpreting partial images such as the top of a head. Stanford researchers have developed vision systems that are better than humans at some tasks such as instantly identifying thousands of cars by make, model and year. Using car images from Google StreetView, the researchers have correlated car prices with certain city neighborhoods. That turns out to predict not only wealth, but also crime patterns and voting preferences.

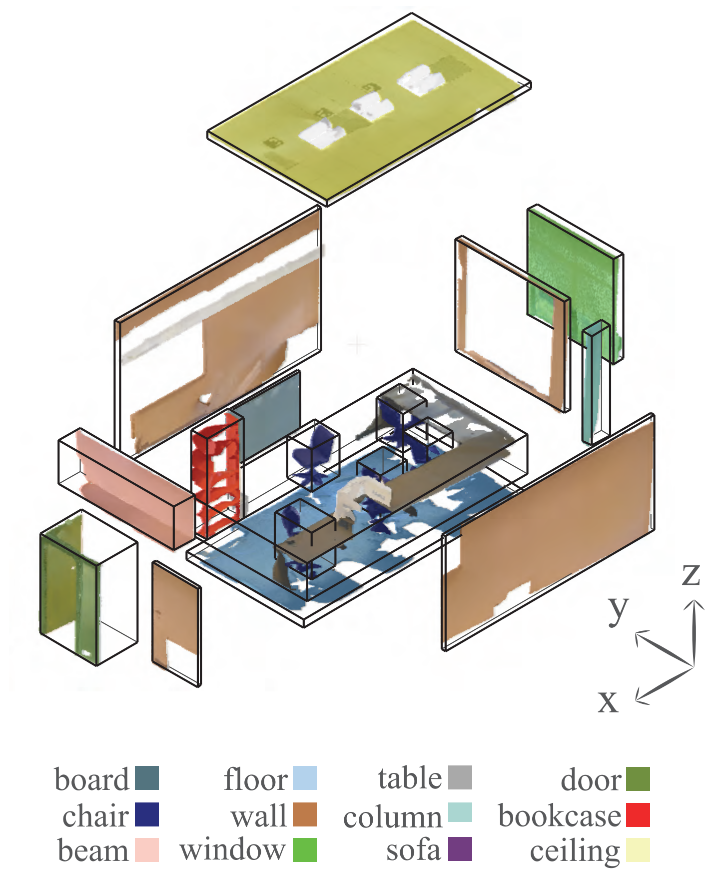

Computer Vision

The next step is teaching computers to describe the relationships in a picture, such as “a cat lying on a bed next to a laptop,” and make sentences and even stories about them. High-level visual recognition and reconstruction problems such as understanding a scene or recognizing human behavior in the complex 3D world requires more sophisticated algorithms. Even a 4-year-old child knows that a boy blowing out candles on a cake is celebrating his birthday. Helping computers reach that level of understanding is something that Stanford researchers are starting to achieve today.

Silvio Savarese, a SAIL researcher and the head of Stanford’s Computational Vision and Geometry Lab, says that one goal is to help robots have real relationships with humans. To interact with people, “You need to infer intention and have a good understanding of how to put together all the components into a coherent interpretation,” he says.

In his lab, researchers are using neural networks to interpret videos of scenes, so that a robot will learn to act appropriately. That requires more understanding than merely avoiding collisions. For instance,looking at a video of a cocktail party, the algorithm learns what distance people maintain while having a conversation. It has to learn that when getting drinks, people form a queue at the bar, and when crossing a room, they avoid passing between two people who are talking.

SAIL researchers are also trying to use machine vision to provide expertise that humans can’t. One project involves construction-project monitoring. A vision system learns a building by scanning a detailed 3-D model of the completed building. Then it goes through the building week-by-week while it is under construction and compares progress to the project-management plan. By identifying delays and incomplete areas early, the technology can help avoid delays and costly overruns.

When machines can see, they will be able to improve security by being tireless watchers, continuously observe medical operations as a separate set of eyes, and improve the ability of cars to interpret and react to changing conditions. Seeing a pedestrian on a street corner and predicting whether he will cross or wait will help make independent vehicles safer. Computer vision is a crucial aspect of AI that is helping machines and people coexist in an intelligent way. ▣

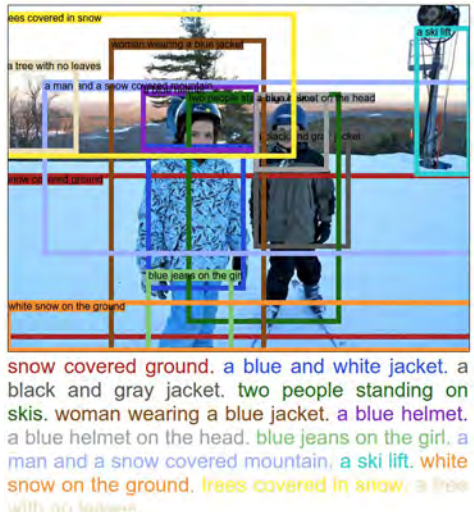

Humans look with their eyes, but they see with their brains. Computer vision lets computers see and understand concepts.

Robotics

Robotics goes beyond the goal of traditional AI—replicating human intelligence—by attempting to build machines that physically act like humans. It extends well beyond computer science, into the fields of mechanical engineering, bioengineering and even psychology of learning.

Robotics developed on factory floors with programmable arms in safety cages doing tightly constrained activities. But around the beginning of the 21st century, sensors and actuators improved to the point that scientists could start to develop robots that perform tasks in a human environment.

Stanford's Robotics Center at SAIL specializes in developing robots that can work in conjunction with humans. Robots can be dangerous. For safety, they must be able to sense humans near them both visually and by touch to avoid hurting them inadvertently. They need to be able to sense their tasks so they can pick up iron bars or porcelain vases with the same grippers.

Roboticists have long instructed robots how to move by creating computer programs that describe the movement of each joint. But that isn't how people move. Oussama Khatib, who heads SAIL's robotics lab, says, "Humans don't move precisely. They move to make contact and then modify their position." Stanford's robotics researchers have worked with biomechanics researchers to model the human musculoskeletal system in order to make robots move the way humans do.

Rather than program the motion of every limb, robotics researchers today try to mimic biological systems. Sensors on

SAIL faculty working in robotics include:
Jeannette Bohg, Oussama Khatib, Mykel Kochenderfer, Andrew Ng, Dorsa Sadigh, Ken Salisbury, Silvio Savarese, Sebastian Thrun.

ROBOTICS

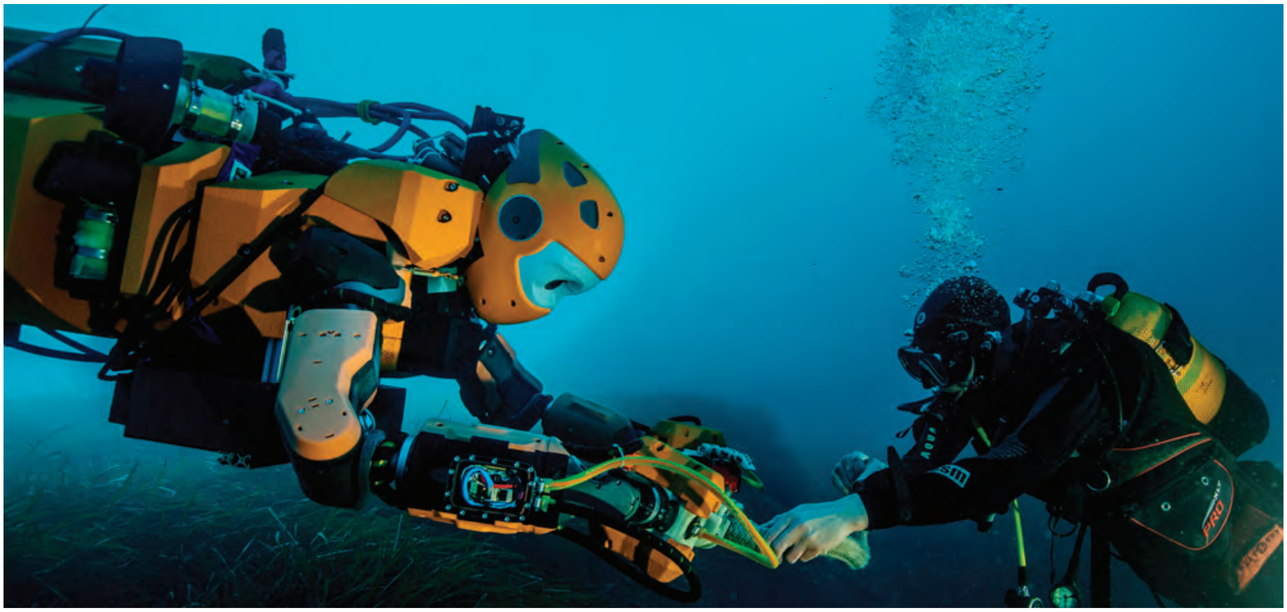

robots respond to stimuli. They provide the information that controls or modifies the next movement.

With those capabilities provided to robots, researchers can use different techniques for instructing robots to move. Based on computer models of human motion that they had developed through careful analysis, “we could understand the human strategy behind the motions,” Khatib says. Then they encoded the strategies into the robot. As a result they are able to encode complex motions into a robot without writing extremely complex code.

One of the most exciting examples of work by the robotics lab is Ocean One—a humanoid diver. Khatib says it is actually more like a mermaid with two arms, but no legs. Several propellers around its torso give Ocean One a range of motion. The robot was designed to assist archaeologists who want to explore shipwrecks hundreds of feet beneath the surface. Human divers can’t go that deep and are limited to shorter times. Submersible vessels can’t maneuver the way divers can. The robot provides a solution.

Robotics

Ocean One has already explored a French ship that sank in 1664 in 320 feet of water in the Mediterranean Sea. Ocean One is a slave robot whose actions are controlled by a human sitting in the ship with haptic controls. Through the haptic system, the controller can feel the same forces in his hands that the robot's sensors transmit. In its maiden voyage, it found and brought to the surface a ceramic pot. Soon it is scheduled to go to the Red Sea off Egypt to investigate deepwater coral reefs.

Robots that can work with humans have other uses. In factories, they can transport goods from one workstation to another, avoiding the need to build expensive assembly lines. Robots that respond to simple instructions could be used in small factories that can't afford expensive programming staffs. In surgery, haptic controls can provide surgeons much more information and ability to guide surgical instruments.

Stanford researchers have used the improved haptic controls to explore the way the brain instructs muscles and limbs to move. Moving a hand, for example, requires signaling 50 different muscles. To study how that occurs, researchers need to analyze brain activity in conjunction with physical activity. They placed a subject wearing haptic controllers on his hands in an MRI machine. The subject can respond to virtual reality while the scanner records the related brain activity.

Increasingly, robotics researchers are focused on teaching robots to respond to their experiences and learn from them. When the robot is told to carry an object across a room, the researchers want it to observe its environment and adapt to changes like people walking around. Algorithmic programming and learning techniques need to combine to produce robots that function with the skills they will need in the future.

Increasingly, robotics researchers are focused on teaching robots to respond to their experiences and learn from them.

GENOMICS

Genomics

Genomics is where medicine meets computer science. Mapping the human genome is something only machines can do. And once a genome has been mapped, AI techniques are being used to decode it.

Traditionally, doctors have tried to understand diseases by looking at phenotypes, and figuring out what caused the phenotypes. Many physical ailments are rooted in the gene sequence that defines every human.

Looking at diseases through the lens of genomics opens up countless new avenues to understand genotypes and phenotypes. Instead of looking at outcomes, genomicists look for the gene sequences that cause them. Stanford researchers have found that variances from normal in gene sequences are correlated with diseases from hypertension to narcolepsy.

The cutting edge of genomics today involves machine learning: teaching machines to sort through millions of gene sequences. Then they try to identify the sequences that correlate with particular phenotypes—the observable characteristics linked to the genes. One fruitful avenue pioneered at Stanford is examining just the regulatory

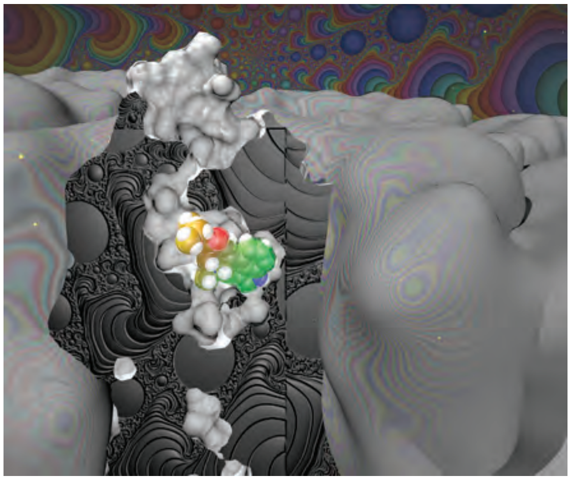

regions of a person's gene sequence. For example, mutations in a region associated with regulating cardiac output appear in the gene of a person with a family history of sudden cardiac failure.

Starting with the genomic code reverses the way people traditionally view the human body and its diseases. Physicians and clinical researchers usually start by trying to understand body parts such as organs and systems. Then they try to understand

SAIL faculty working in genomics, medicine, and healthcare include:
Serafim Batzoglou, Gill Bejerano, Ron Dror, Daphne Koller, Anshul Kundaje, Fei-Fei Li, Vijay Pande, Ken Salisbury, Dan Yamins.

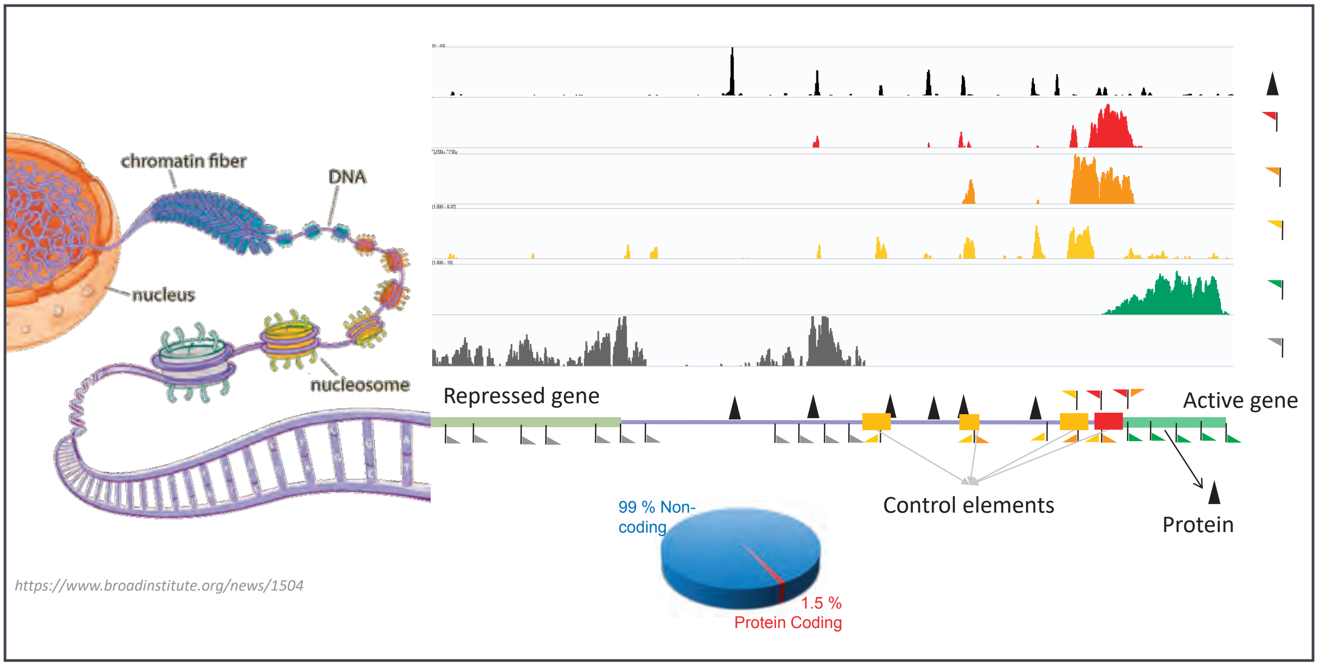

Biochemical markers of cell-type specific functional elements

phenotypes related to disease of those parts. But it turns out that variations in the genome may damage multiple body parts in different ways. Understanding the full impact may help doctors devise individualized medical treatments or predict side effects.

Biologists once thought that decoding the human genome would provide a clear path to understanding many human characteristics. They hoped that identifying anomalies in a gene sequence would explain a specific disease. They hoped that would suggest pathways for cures.

But the genome is more complex than expected. Figuring out which gene sequences matter is a huge challenge. Moreover, they seem to interact in many ways, some of which have unnoticed significance.

Genomics

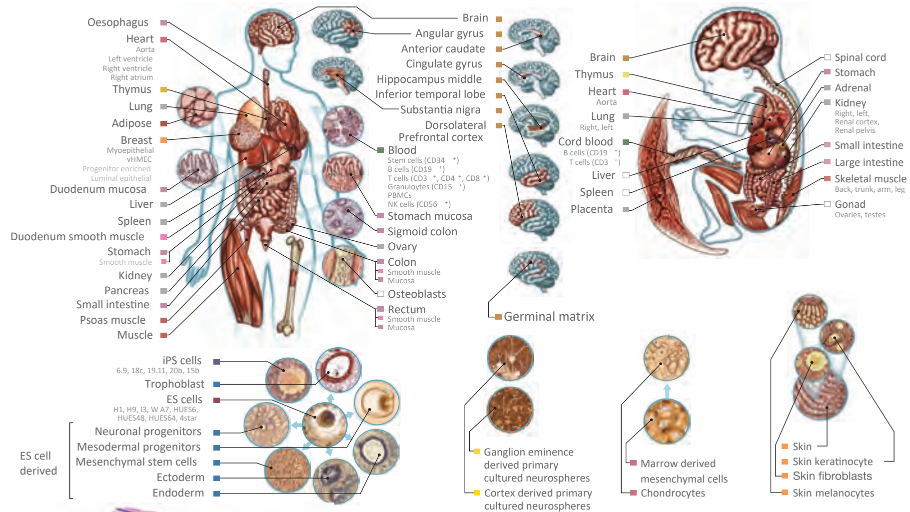

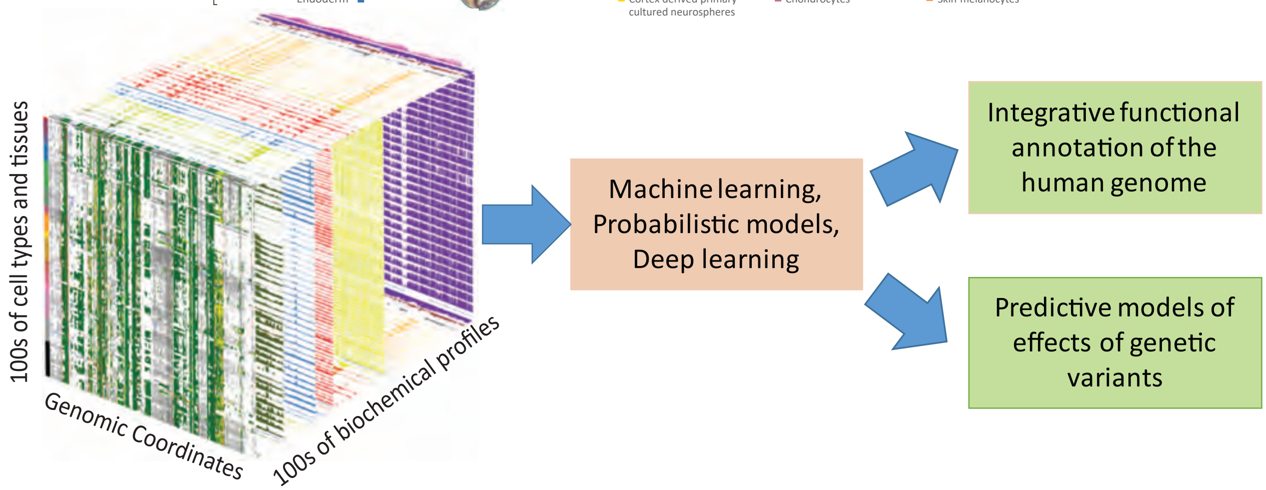

Gill Bejerano, a Stanford researcher, says the biology is complicated. In a typical patient's three billion nucleotide genome there might be four million mutations, and 300 of them might look suspicious—far too many to test experimentally.

Now, he says, researchers in his lab and elsewhere are focusing on using machine learning neural network techniques to find correlations among genotypes and phenotypes—the genes and their outcomes. "We are trying to take patient groups with a single disease and ask if they have anything in their genome in common." Using machine learning techniques, he has been able to demonstrate ways to eliminate all but a few of the 300 suspicious mutations.

This research avenue has become much more fruitful with the development of

CRISPR-Cas9 gene editing technology. CRISPR allows researchers to go to a specific point on a genome and cut it or block its enzymatic activity in order to see how the loss of that part of the code affects the cell. When a gene has been identified as the likely cause of a disease, it can be tested. If the gene sequence exists in a mouse or in a blood cell, the gene can be replaced using CRISPR, and researchers can see whether changing it had the expected effect.

Even without finding a cure, identifying rogue genes in a patient can help provide genetic counseling. When a child has a genetic disease, parents want to know whether future siblings might be affected as well. That requires mapping the parents' genes. If one or both have the same gene sequence, it makes it more likely future children will inherit it. But if neither has the gene, the anomaly is probably just random and is unlikely to appear in other children.

One of the most fascinating areas of genomic research involves comparing genetic variation among species. Many genes are common even among vastly different species like humans and rats. Understanding the differences can give valuable insight into species variation.

Bejerano's lab has studied dolphin genomes, and found that they lack a gene that is present in many mammal species and is associated with virus fighting capabilities. He suspects that mutation may help explain the periodic die-offs of pods of dolphins. He hopes further genomic research can determine what in the gene allows dolphins "to echo-locate and what allows them to live in the ocean." ▣

Every single genome sequenced is a treasure chest of secrets…this is an amazing time to be a genomicist.

Autonomous Vehicles

Driverless cars, planes and boats are going to profoundly change people's lives and our environment over the next decades. Making them work in a world of unpredictability and ambiguous information remains a challenge. Researchers at SAIL are exploring ways to help vehicles operate safely in an environment that is constantly changing due to unpredictable human behaviors.

Autonomous cars are already on the roads in some places, but there are significant restrictions on their performance.
AV researchers say it will be a number of years before cars and trucks can be set loose on city streets where pedestrians, bicyclists, arm-waving traffic cops and double-parked delivery trucks constantly alter the driving environment.

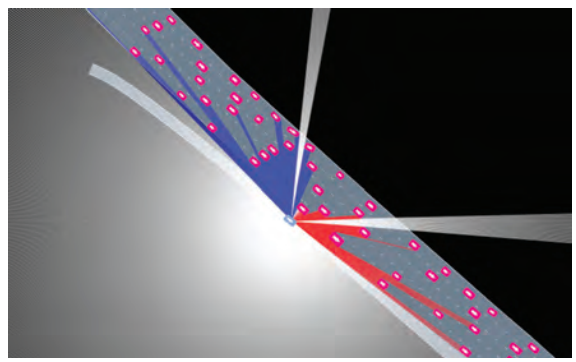

Stanford researchers from fields including robotics, computer vision, human-computer interaction, machine-learning and decision-making are working on ways to make cars interact safely with humans—no easy task for a 3,500 pound robot on wheels that can travel 100 miles per hour.

Teaching cars to observe and predict human actions is vital. There is bound to be a lengthy transition period before all 260 million U.S. cars and trucks are replaced by fully autonomous vehicles. Even then, cars will share spaces with bicyclists and pedestrians. Moreover, for a number of years, new cars will have increased autonomy while still requiring humans to take over in some situations such as navigating unmarked parking lots, following police directions that override traffic signals or merging multiple traffic streams.

Some Stanford researchers are working on issues related to the car's interaction

SAIL faculty working in autonomous vehicles and intelligent machines include:
Maneesh Agrawala, John Duchi, Stefano Ermon, Leonidas Guibas, Oussama Khatib, Mykel Kochenderfer, James Landay, Fei-Fei Li, Juan Carlos Niebles, Christopher Ré, Dorsa Sadigh, Ken Salisbury, Silvio Savarese, Sebastian Thrun.

with its own driver. They study drivers in an immersive AV simulator and see how they react to various situations. They are developing vision systems that observe the driver's attentiveness. One issue with part-time automation is that drivers tend to nod off while the car is handling things. When the car isn't sure what to do, it has to make sure the transfer of control to the driver is fail-safe.

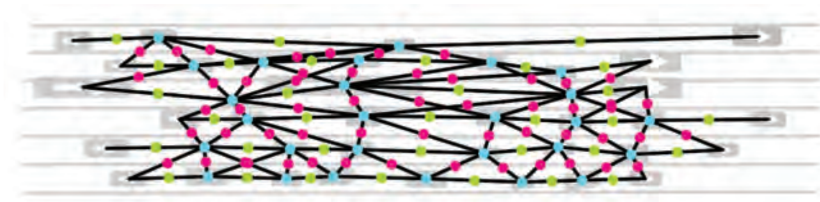

Prof. Kochenderfer's students are developing new methods for randomly generating realistic driving scenes and trajectories.

Human drivers easily understand many things happening in the outside environment, but teaching them to a car requires more than just identifying objects. Senior Research Scientist Juan Carlos Niebles, a computer vision expert who is Associate Director of research at the Stanford-Toyota Center for AI Research, says: "We have very good cameras today. The bottleneck is the software that correctly interprets the world and what the pixels are showing."

For example a child wobbling along in a bike lane requires more caution than a spandex-clad bike messenger. A ball rolling onto the street may be followed by a child. A driver in the next lane who is texting or putting on makeup is a bigger risk than someone with both hands on the wheel. Drivers and pedestrians will need to accept the fact that a vision system in every passing car is watching and analyzing them, (although they may not be storing the video).

Decision-making is one of the key issues for AVs, because there are far more possible situations than programmers can ever model in advance. Good human drivers are told to follow the rules of the road. But good driverless cars may be better off with innovative software that follows the probabilities of the road.

Past work on decision-making by Mykel Kochenderfer, an AI researcher at Stanford, has led to a rethinking of the way airplanes' collision avoidance software is designed. He has proved that in domains where uncertainty is unavoidable, probabilistic models are safer than rules-based models.

Autonomous Vehicles

A similar approach may be needed for the vastly more complicated problem of making cars that can safely navigate roadways.

In some cases, the car must decide to swerve slightly to one side in order to give itself a better view of an intersection. Writing an expert system that gives yes-no answers to every question would be impossible given the variety of situations that develop. But Kochenderfer acknowledges that consumers and regulators are generally more comfortable with rules-based systems that humans can understand than they are with turning over control to a computer, whose decision-making process is obscure.

Recognizing the leadership position of Stanford's AI research community, SAIL and Toyota came together in 2015 to create a multi-year research center focusing on human-centric AI research in autonomous vehicles and robotics. Today, the SAIL-Toyota Center hosts more than a dozen research projects led by twenty principle

As part of the Toyota SAIL project, students of Mark Cutkosky and Chris Gerdes test applications of a skin stretch haptic steering wheel on experimental testbed X1 for navigation cues, collision avoidance warnings, and autonomous previews.

The Prius V is used as part of the Toyota SAIL project to create driver-centric datasets to support machine learning around driver emotion, physiology, and elicited responses during naturalistic driving.

investigators and their students and postdoctoral scholars. Research topics span from theoretical estimation of uncertainties, to large-scale database querying of anomalies, to visual-motor deep reinforcement learning, to in-car human behavior studies. The world will have to take incremental steps toward fully autonomous vehicles. But the benefits of getting there will be immense. Fully autonomous vehicles will liberate elderly and disabled people who can't drive. And AVs can be programmed to use the streets much more efficiently, reducing congestion. Most important, roads will be much safer. Driver error is the main cause of the six million accidents that occur annually in the U.S. and permanently injure two million people. ▣

This is a Mercedes-Benz E-Class outfitted with lidar and radar sensors. This vehicle is being used to collect data for the development of deep stochastic sensor models, which will allow the efficient simulation of autonomous vehicles in realistic environments. This work is being conducted by Tim Wheeler in Prof. Kochenderfer's lab in collaboration with the Technical University Darmstadt in Germany. The vehicle is owned by TU Darmstadt.

Creating the intelligence for fully autonomous vehicles is one of the great opportunities of AI.

AI Salon

A special piece of SAIL's culture is the AI Salon, which is a biweekly event modeled after the 18th century French enlightenment salons. At each salon, we invite two Stanford graduate students, faculty, or guests to share contrasting thoughts on a topic of relevance to artificial intelligence. The goal is to foster discussion that takes a wider view than typical day-to-day research; for instance, past topics have included "AI and the Legal System," "Software Engineering for Machine Learning," "Trust in AI Techniques/Algorithms," and "Diversity in AI." While the Salon culture is driven by graduate students and faculty, we also regularly have invited guests ranging from visiting professors to CEOs, journalists, and judges.

Salon attendees are treated to wine and cheese. In exchange, we enforce a strict rule of absolutely no electronics, both to remind us of the enlightenment era, and so that everyone is fully present during the discussion. We take this seriously—time is even kept with an hourglass rather than a clock. Everyone is invited to participate in the discussion; while two hosts introduce a topic at the beginning, the majority of discussion comes from the audience at large.

Salon topics typically lead popular awareness. For instance, in April and October 2015 we held Salons discussing how filtered newsfeeds shape society; this later became a major topic of discussion during the 2016 U.S. election season. In April of 2015, Elon Musk participated in a discussion of the future of AI, prompting lively debates with students on both the promises and perils that advanced AI technology will bring to the society. In January 2016, Tino Cuellar, Stanford affiliate faculty and one of the Supreme Court justices of the State of California, came to talk to SAIL members about AI and its challenges to our legal

system. Several SAIL emeritus faculty have participated in our Salons, including Ed Feigenbaum and Nils Nilsson. They had dialogues with current students at SAIL on both the history and the prospects of AI.

By bringing together the many bright minds at Stanford, as well as expert guests, we hope to give Stanford graduate students and faculty the resources to become leaders in the public dialogue around artificial intelligence. Given AI's growing impact on society, engaging with this dialogue is more important than ever.

The AI salon was started in 2014 by graduate student Jacob Steinhardt and SAIL Director Prof. Fei-Fei Li. ▣

Given AI's growing impact on society, engaging with this public dialogue is more important than ever.

Stanford AI4ALL, formerly known as SAILORS

THE STANFORD ARTIFICIAL INTELLIGENCE LAB OUTREACH SUMMER

As the field of AI continues to make a bigger impact in the world, researchers and educators at SAIL believe that to develop the most inclusive, humanistic, and benevolent technologies, it is imperative that the field of AI includes students, researchers, and technologists from all walks of life. With this mission in mind, SAILORS (short for the "Stanford Artificial Intelligence Laboratory's Outreach Summer Program") was created in 2015 to expose high school students in underrepresented populations to the field of artificial intelligence. This program has since been renamed to Stanford AI4ALL, signalling a partnership between SAIL and education non-profit AI4ALL. This annual summer program is built upon the hypothesis that to increase diversity representation in the field of AI, CS, and STEM at large, it is critical to introduce the technology along with its humanistic mission statements. And in turn, the long-term vision for AI and STEM is for the fields to train a more diverse generation of technologists who have humanistic goals in mind when designing next generation technologies.

In 2015 and 2016, SAILORS was aimed at rising 10th-grade young women. Originally two-weeks, this three-week full-time program continues to provide both broad exposure to AI topics through faculty lectures and industry field trips, as well as an in-depth experience with a research area through hands-on projects. Every part of the Stanford AI4ALL curriculum is designed to combine rigorous technical exposure with important humanistic applications. For example, in prior years, the robotics team worked on self-driving cars to help aging seniors. The natural language processing team worked on document analysis using Twitter data for disaster relief. The computer vision team worked on clinician hand hygiene behavior analysis using depth-image videos from hospitals. And the computational genomics team worked on leukemia classification. The students also had a chance to visit local companies, as well as the Computer History Museum, personally curated by SAIL emeritus professor, Turing Award winner Professor Edward Feigenbaum.

Stanford AI4ALL/SAILORS was co-founded and co-directed by Prof. Fei-Fei Li; her former PhD student Dr. Olga Russakovsky (assistant professor at Princeton University) herself a past participant of another Stanford Pre-Collegiate Studies program, the Stanford University Mathematics Camp; and Dr. Rick Sommer (Executive Director of Stanford Pre-Collegiate Studies). Now renamed, this program will continue to be based in the Stanford Artificial Intelligence Laboratory, but with a new, updated curriculum developed by AI4ALL, an education nonprofit organization dedicated to training the next generation of AI researchers. Both SAILORS and AI4ALL were founded by Dr. Fei-Fei Li and Dr. Russakovsky. Juan Carlos Niebles, Senior Research Scientist at SAIL, was named as the Director for Stanford AI4ALL 2018.

More than 40 members of the Stanford CS department helped make the program possible each year, including Professors Gill Bejerano, Stefano Ermon, Noah Goodman, Oussama Khatib, Mykel Kochenderfer, Anshul Kundaje, Percy Liang, Chris Manning, Ken Salisbury, and Silvio Savarese, as well as undergraduates, graduate students and postdoctoral fellows.

Renowned computer scientists Dr. Ruzena Bajcsy (Robotics professor at UC–Berkeley) and Dr. Maria Klawe (President of Harvey Mudd College) delivered the keynote speeches in respective years. In 2015, two undergraduate students of Computer Science, Marie Eve Vachovsky and Grace Wu, performed a rigorous study of Stanford AI4ALL as a summer undergraduate research project through the CURIS program. Their finding was published in a research paper at SIGCSE2016, a premier conference in computer science education. Wired Magazine had a feature story on Stanford AI4ALL in their August issue in 2015. In 2015 and 2016, Stanford AI4ALL was free for commuter participants thanks to the generous support of Dropbox, Google, Bloomberg, Oculus, Intel, Airbnb, Baidu, Pinterest, and many other companies and individuals. Since then, the program has evolved into a residential program that requires tuition, however, Stanford AI4ALL does offer financial aid in order to best serve a global, diverse audience.

Efforts for Stanford AI4ALL 2018 are well underway now. This will be the second year that the program is residential, having expanded to 32 students. AI4ALL is now partnering with several other universities, creating similar programs aimed at increasing diversity and access to computer science research.

AI will change the world.
Who will change AI?

Faculty Focus

The Stanford Artificial Intelligence Laboratory (SAIL) has been a center of excellence for Artificial Intelligence research, teaching, theory, and practice since its founding in 1962.

MANEESH AGRAWALA

MANEESH AGRAWALA
Maneesh Agrawala works on computer graphics, human computer interaction and visualization. His focus is on investigating how cognitive design principles can be used to improve the effectiveness of audio/visual media. The goals of this work are to discover the design principles and then instantiate them in both interactive and automated design tools.

SERAFIM BATZOGLOU

SERAFIM BATZOGLOU
Serafim's research focuses on computational genomics: developing algorithms, machine learning methods, and systems for the analysis of large scale genomic data. He received a BS in Computer Science, BS in Mathematics, and MEng in EE&CS from MIT in June 1996, and a PhD in Computer Science from MIT in June 2000. He joined Stanford in 2001. He is a recipient of the Sloan Fellowship, the NSF Career Award, and was named Technology Review's "Top 100 Young Technology Innovators" in 2003. He is also co-founder of DNAnexus.

JEANETTE BOHG

JEANNETTE BOHG

Jeannette Bohg is an assistant professor in the Department of Computer Science at Stanford University. She leads the Interactive Perception and Robot Learning Lab that seeks to understand the underlying principles of robust sensorimotor coordination by implementing them on robots. The lab is specifically interested in developing methods that are goal-directed, real-time and multi-modal such that they can provide meaningful feedback during action execution and learning. The research is at the intersection of robotics, computer vision and machine learning applied to the problem of autonomous manipulation and grasping. Topics include real-time visual tracking, learning to grasp, vision-based control and robotic system building as well as interactive perception.

GILL BEJERANO

GILL BEJERANO

The Bejerano Lab studies genome function in human and related species. We are deeply interested in the following broad questions: Mapping genome sequence (variation) to phenotype (differences) and extracting specific genetic insights from deep sequencing measurements. We take a particular interest in gene cis regulation. We use our joint affiliation to apply a combination of computational and experimental approaches. We collect large scale experimental data; write computational analysis tools; run them massively to discover the most exciting testable hypotheses; which we proceed to experimentally validate. We work in small teams, in house or with close collaborators of experimentalists and computational tool users who interact directly with our computational tool builders.

EMMA BRUNSKILL

EMMA BRUNSKILL

Emma Brunskill is an assistant professor in the Department of Computer Science at Stanford University in the Statistical Machine Learning Group and Stanford Artificial Intelligence Laboratory. Her work focuses on reinforcement learning in high stakes scenarios—how can an agent learn from experience to make good decisions when experience is costly or risky, such as in educational software, healthcare decision making, robotics or people-facing applications. Key areas of focus in her group include the theoretical foundations of efficient reinforcement learning, "What if" reasoning for sequential decision making, and human-in-the-loop systems. She is the recipient of multiple early faculty career awards (NSF, ONR, Microsoft) and her group has received several best paper nominations and awards.

Faculty Focus
continued

STEFANO ERMON

STEFANO ERMON
We research innovative computational approaches to help address the societal and environmental challenges of the 21st century. We combine research on the foundations of artificial intelligence and machine learning with applications in science and engineering. Our work enables computers to act intelligently and adaptively in increasingly complex and uncertain real world environments.

RON DROR

RON DROR
We are a highly multidisciplinary lab whose work spans fields ranging from high-performance computing, computer vision, and machine learning to biochemistry, cell biology, and drug discovery. We study the spatial organization and dynamics of biomolecules and cells in order to develop better medicines and quantitatively explain the workings of living systems.

RON FEDKIW

RON FEDKIW
Fedkiw's group generally works on designing numerical algorithms for the simulation of physics based systems spanning a range from classical applied mathematics and engineering all the way to Hollywood special effects for feature films. He has two Academy Awards: one for his work on water simulation for Pirates of the Caribbean, Harry Potter, etc. as well as a second award for destruction simulation in Transformers, The Hulk, Avengers, etc. More recently, he has been focused on a variety of special effects simulation problems that require real world computer vision and related data in order to augment the simulations themselves to a new level of fidelity. Current efforts include the flying of drones to capture, reconstruct, and subsequently simulate very large scale detailed trees, novel cloth simulation technologies that allow for the capture, interpretation, and further use of detailed wrinkling processes, and a hybrid approach to 3D face and body simulation and data capture methods incorporating biomechanical systems as priors. His most recent work on face simulations that target and improve real world data capture has just been used on King Kong's face in the new Skull Island movie that releases Spring 2017.

JOHN DUCHI

JOHN DUCHI
I am an assistant professor of Statistics and Electrical Engineering at Stanford University. I completed my PhD in computer science at Berkeley in 2014. My research interests are a bit eclectic, and they span computation, statistics, optimization, and machine learning; if you like any of these, we can probably find something interesting to chat about. At Berkeley, I worked in the Statistical Artificial Intelligence Lab (SAIL) under the joint supervision of Michael Jordan and Martin Wainwright. I obtained my master's degree (MA) in statistics in Fall 2012. I was also an undergrad and a masters student at Stanford University, where I worked with Daphne Koller in her research group, DAGS. I also spend some time at Google Research, where I had (and continue to have) the great fortune to work with Yoram Singer.

MICHAEL GENESERETH

MICHAEL GENESERETH

The focus of work in the Logic Group is Computational Logic. (Computational Logic is that branch of Computer Science concerned with the representation and processing of information in the form of logical statements.) Current research topics include logical languages, automated deduction, automated reformulation, paraconsistent reasoning, and dynamic logic. While the group's work is primarily theoretical, a considerable amount of effort is devoted to applications of research results. Application areas of special interest at this time include Logical Spreadsheets, Data Integration, Logic Programming, Enterprise Management, Computational Law, and General Game Playing.

NOAH GOODMAN

NOAH GOODMAN

The Computation and Cognition Lab at Stanford studies the computational basis of natural and artificial intelligence, merging behavioral experiments with formal methods from statistics and programming languages. Our research topics include language understanding, social reasoning, and concept learning. We also create technologies to ease the use of complex techniques for cognitive modeling and artificial intelligence applications, such as probabilistic programming languages.

LEONIDAS GUIBAS

LEONIDAS GUIBAS

The Geometric Computation Group, headed by Professor Leonidas Guibas, addresses a variety of algorithmic problems in modeling physical objects and phenomena, and studies computation, communication, and sensing as applied to the physical world. Current foci of interest include the analysis of shape or image collections, geometric modeling with point cloud data, the development and annotation of the ShapeNet 3D model repository, machine learning for geometric objects, the analysis of mobility data, and modeling the shape and motion biological structures. More theoretical work is aimed at investigating fundamental computational issues and limits in geometric computing and modeling, including the handling of uncertainty, the analysis of networks of maps and correspondences, and the geometric properties of deep networks. The group is also interested in automated interactive teaching tools.

Faculty Focus

continued

DAN JURAFSKY

DAN JURAFSKY

Dan Jurafsky is Professor and Chair of Linguistics and Professor of Computer Science at Stanford University. He is the recipient of a 2002 MacArthur Fellowship, is the co-author with Jim Martin of the widely-used textbook “Speech and Language Processing,” and co-created with Chris Manning one of the first massively open online courses, Stanford’s course in Natural Language Processing. His new trade book “The Language of Food: A Linguist Reads the Menu” came out on September 15, 2014, and was a finalist for the 2015 James Beard Award. Dan received a BA in Linguistics in 1983 and a PhD in Computer Science in 1992 from the University of California at Berkeley, was a postdoc 1992-1995 at the International Computer Science Institute, and was on the faculty of the University of Colorado, Boulder until moving to Stanford in 2003. His research ranges widely across computational linguistics; special interests include natural language understanding, machine translation, spoken language and conversation and the relationship between human and machine processing.

THOMAS ICARD

THOMAS ICARD

Thomas works in formal philosophy on topics at the intersection of philosophy, artificial intelligence, and cognitive science, including causal reasoning, natural language semantics, and action theory and decision making. The main technical focus is on formalisms combining rich logical structure with probabilistic and statistical reasoning. An ongoing philosophical interest, connected with central issues in philosophy of AI, is how descriptive questions (“What is in fact the case?”) relate to normative questions (e.g., “What ought we do?”).

OUSSAMA KHATIB

OUSSAMA KHATIB

Oussama has made fundamental contributions to robotics, control, and human motion analysis. His work includes the development of potential fields for control, the operational space control framework, whole body multi-contact control with prioritized null spaces, elastic planning, articulated body dynamic simulation, haptic rendering, and biomechanics based analysis of human motion.

MYKEL KOCHENDERFER

MYKEL KOCHENDERFER

The Stanford Intelligent Systems Laboratory (SISL) researches advanced algorithms and analytical methods for the design of robust decision making systems. Of particular interest are systems for air traffic control, automated vehicles, and other applications where decisions must be made in uncertain, dynamic environments while maintaining safety and efficiency. Research at SISL focuses on efficient computational methods for deriving optimal decision strategies from high-dimensional, probabilistic problem representations.

DAPHNE KOLLER

DAPHNE KOLLER

My main research focus is on the use of machine learning to analyze and extract valuable insight from large biomedical data sets, and on the development of novel machine learning methods to address the unique challenges of these data. The techniques used include traditional and Bayesian classification and regression, deep learning, and probabilistic graphical models. My work involves data at the cellular scale (including multiple types of 'omics data and microscopy), physiological scale (including diverse phenotypes as well as imaging), and population cohorts over time. A particular focus is on understanding the mechanisms underlying aging and age-related diseases, with the goal of identifying interventions that can help people live longer, healthier lives.

ANSHUL KUNDAJE

ANSHUL KUNDAJE

The Kundaje lab is an interdisciplinary computational biology lab in the departments of Genetics and Computer Science at Stanford University. Our research focuses on deciphering the molecular and genetic basis of disease by integrative analysis of diverse types of large-scale genomic data. The Kundaje lab develops statistical and machine learning methods to map functional elements in the human genome across diverse cell types and tissues, model the three-dimensional genome architecture, reverse engineer mechanisms of gene regulation and interpret the molecular impact of natural and disease-associated genetic variation. We have led the integrative analysis efforts for two of the largest functional genomics consortia—The Encyclopedia of DNA Elements (ENCODE) and The Roadmap Epigenomics Project resulting in the most comprehensive functional annotation of the human genome to date. We specialize in developing interpretable machine learning methods based on deep neural networks to integrate diverse types of large-scale genomic data and infer the genomic and molecular basis of health and disease.

JAMES LANDAY

JAMES LANDAY

James Landay was named the Anand Rajaraman and Venky Harinarayan Professor in the School of Engineering in 2016. He specializes in human-computer interaction – the study, design, and implementation of the interaction between people and computers—and how it can be used to improve and inform people's health, sustainability, education and more.

Faculty Focus
continued

FEI-FEI LI

Dr. Fei-Fei Li is an Associate Professor in the Computer Science Department at Stanford, and the Director of the Stanford Artificial Intelligence Lab and the Stanford Vision Lab. She is also the Director of the recently established Stanford Toyota Center for Human-Centric AI Research. Dr. Fei-Fei Li's main research areas are in machine learning, deep learning, computer vision and cognitive and computational neuroscience. She has published more than 150 scientific articles in top-tier journals and conferences, including Nature, PNAS, Journal of Neuroscience, CVPR, ICCV, NIPS, ECCV, IJCV, IEEE-PAMI, etc. Dr. Fei-Fei Li obtained her B.A. degree in physics from Princeton in 1999 with High Honors, and her PhD degree in electrical engineering from California Institute of Technology (Caltech) in 2005. She joined Stanford in 2009 as an assistant professor, and was promoted to associate professor with tenure in 2012. Prior to that, she was on faculty at Princeton University (2007-2009) and University of Illinois Urbana-Champaign (2005-2006). Dr. Li is the inventor of ImageNet and the ImageNet Challenge, a critical large-scale dataset and benchmarking effort that has contributed to the latest developments in deep learning and AI. In addition to her technical contributions, she is a national leading voice for advocating diversity in STEM and AI. She is co-founder of Stanford's renowned SAILORS outreach program for high school girls. For her work in AI, Dr. Li is a speaker at the TED2015 main conference, a recipient of the IAPR J.K. Aggarwal Prize, the 2016 nVidia Pioneer in AI Award, 2014 IBM Faculty Fellow Award, 2011 Alfred Sloan Faculty Award, 2012 Yahoo Labs FREP award, 2009 NSF CAREER award, the 2006 Microsoft Research New Faculty Fellowship and a number of Google Research awards. Work from Dr. Li's lab have been featured in a variety of popular press magazines and newspapers including New York Times, Wall Street Journal, Fortune Magazine, Science, Wired Magazine, MIT Technology Review, Financial Times, and more.

JURE LESKOVEC

Jure Leskovec studies massive complex networks, which allows him to model large complex interconnected systems at all scales, from interactions of proteins in a cell to interactions between humans in a society. He and his research group work on machine learning and data mining approaches to study complex networks, their evolution, and the diffusion of information and influence over them. Computation over massive data is at the heart of his research and has applications in computer science, social sciences, economics, marketing, and healthcare.

PERCY LIANG

Percy Liang's research focus spans understanding theoretical foundations of machine learning to building practical natural language understanding systems. An overarching goal is to build virtual assistants that can perform complex tasks based on natural language instructions. A few technical themes include: program induction: inferring how to perform a task given what to do; interactive learning: how to learn with humans in the loop; and reliable learning: how to design machine learning systems that are robust and understandable when deployed in the real world.

TENGUY MA

TENGYU MA

Tengyu Ma's research focuses on algorithm design and machine learning. The overarching goal is to achieve the fundamental understanding of machine learning and other areas of artificial intelligence, which can lead to practical implications. His work results in novel efficient algorithms with theoretical guarantees as well as wide applications. In particular, he has worked on non-convex optimization, representation learning, deep learning, distributed optimization, and high-dimensional statistics.

CHRIS MANNING

CHRIS MANNING

Christopher Manning is a professor of computer science and linguistics at Stanford University. His PhD is from Stanford in 1995, and he held faculty positions at Carnegie Mellon University and the University of Sydney before returning to Stanford. His research goal is computers that can intelligently process, understand, and generate human language material. Manning concentrates on machine learning approaches to computational linguistic problems, including syntactic parsing, computational semantics and pragmatics, textual inference, machine translation, and deep learning for NLP. He is an ACM Fellow, a AAAI Fellow, and an ACL Fellow, and has co-authored leading textbooks on statistical natural language processing and information retrieval.

ANDREW NG

ANDREW NG

Dr. Andrew Ng is a globally recognized leader in AI (Artificial Intelligence). He was until recently Chief Scientist at Baidu, where he led the company's ~1300 person AI Group and was responsible for driving the company's global AI strategy and infrastructure. He led Baidu Research in Beijing and Silicon Valley as well as technical teams in the areas of speech, big data and image search. Dr. Ng is an Adjunct Professor at Stanford University's Computer Science Department. In 2011 he led the development of Stanford's Massive Open Online Course (MOOC) platform and taught an online machine learning class that was offered to over 100,000 students. This led to the co-founding of Coursera, the world's leading MOOC (Massive Open Online Courses) platform, where he continues to serve as chairman. Previously, Dr. Ng was the founding lead of the Google Brain deep learning project. Dr. Ng has authored or co-authored over 100 research papers in machine learning, robotics and related fields. He holds degrees from Carnegie Mellon University, MIT and the University of California, Berkeley.

JUAN CARLOS NIEBLES

JUAN CARLOS NIEBLES

Juan Carlos Niebles received a B.S. degree in Electronics Engineering from Universidad del Norte (Colombia) in 2002, an M.Sc. degree in Electrical and Computer Engineering from University of Illinois at Urbana-Champaign in 2007, and a Ph.D. degree in Electrical Engineering from Princeton University in 2011. He is a Senior Research Scientist at the Stanford AI Lab and Associate Director of Research at the Stanford-Toyota Center for AI Research since 2015. He is also an Assistant Professor of Electrical and Electronic Engineering in Universidad del Norte (Colombia) since 2011. His research interests are in computer vision and machine learning, with a focus on visual recognition and understanding of human actions and activities, objects, scenes and events. His computer vision research has been sponsored by a Google Faculty Research award (2015), the Microsoft Research Faculty Fellowship (2012), a Google Research award (2011) and the Colombian science agency, COLCIENCIAS.

Faculty Focus
continued

VIJAY PANDE

VIJAY PANDE

Vijay Pande's research centers on novel cloud computing simulation techniques to address problems in chemical biology. In particular, he has pioneered distributed computing methodology to break fundamental barriers in the simulation of protein and nucleic acid kinetics and thermodynamics. As director of the Folding@home project, Prof. Pande has, for the first time, directly simulated protein folding dynamics, making quantitative comparisons with experimental results, often considered a "holy grail" of computational biology. His current research also includes novel computational methods for drug design, especially in the area of protein misfolding and associated diseases such as Alzheimer's and Huntington's Disease.

CHRISTOPHER RÉ

CHRISTOPHER RÉ

Christopher Ré is an associate professor in the Department of Computer Science at Stanford University in the InfoLab who is affiliated with the Statistical Machine Learning Group, Pervasive Parallelism Lab, and Stanford AI Lab. His work's goal is to enable users and developers to build applications that more deeply understand and exploit data. His contributions span database theory, database systems, and machine learning, and his work has won best paper at a premier venue in each area, respectively, at PODS 2012, SIGMOD 2014, and ICML 2016. In addition, work from his group has been incorporated into major scientific and humanitarian efforts, including the IceCube neutrino detector, PaleoDeepDive and MEMEX in the fight against human trafficking, and into commercial products from major web and enterprise companies.

DORSA SADIGH

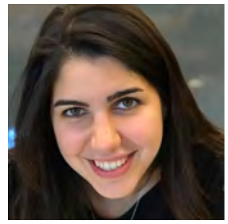

DORSA SADIGH

Dorsa Sadigh is an assistant professor in Computer Science and Electrical Engineering at Stanford University. Her research interests lie in the intersection of robotics, control theory, formal methods, and human-robot interaction. Specifically, she works on developing efficient algorithms for autonomous systems that safely and reliably interact with people. Dorsa has received her doctoral degree in Electrical Engineering and Computer Sciences (EECS) at UC Berkeley in 2017, and has received her bachelor's degree in EECS at UC Berkeley in 2012. She is awarded the NSF and NDSEG graduate research fellowships as well as the Leon O. Chua departmental award, the Arthur M. Hopkin departmental award, and the Google Anita Borg Scholarship.

KEN SALISBURY

KEN SALISBURY

Professor Salisbury received his PhD from Stanford in Mechanical Engineering in 1982. At MIT from 1982-1999, he served as Principal Research Scientist in Mechanical Engineering and as a member of the Artificial Intelligence Laboratory. Some of the projects with which he has been associated include the Salisbury (Stanford-JPL) Robot Hand, the JPL Force Reflecting Hand Controller, the MIT-WAM arm, and the Black Falcon Surgical Robot. His work with haptic interface technology led to the founding of SensAble Technologies Inc., producers of the PHANTOM haptic interface and 3D FreeForm software. In 1997 he joined the staff of Intuitive Surgical, in Mountain View CA, where as a Fellow and Scientific Advisor his efforts focused on the development of telerobotic systems for the operating room. In the fall of 1999 he joined the faculty at Stanford in the Departments of Computer Science and Surgery where his research focuses on medical robotics and surgical simulation, and the design for robots for interaction with and near humans. The most recent spinout from his lab resulted in the founding of the personal robotics program at www.willowgarage.com.

SEBASTIAN THRUN

SEBASTIAN THRUN
We are a team of expert computer scientists with the singular aim of significantly helping society through artificial intelligence technologies, and are constantly on the lookout for high-impact projects. We have worked on robotics, self-driving cars, automated homes, healthcare, drones, and a number of other applications. We currently focus on three areas: AI for healthcare, AI for people-prediction, and smart-homes.

SILVIO SAVARESE

SILVIO SAVARESE
Savarese's research addresses the theoretical foundations and practical applications of computational vision. His group's interest lies in discovering and proposing the fundamental principles, algorithms and implementations for solving high level visual recognition and reconstruction problems such as object and scene understanding as well as human behavior recognition in the complex 3D world. Savarese earned his Ph.D. in Electrical Engineering from the California Institute of Technology in 2005 and was a Beckman Institute Fellow at the University of Illinois at Urbana-Champaign from 2005–2008. He joined Stanford in 2013 after being Assistant and then Associate Professor (with tenure) of Electrical and Computer Engineering at the University of Michigan, Ann Arbor, from 2008 to 2013. He is recipient of several awards including the Best Student Paper Award at CVPR 2016, the James R. Croes Medal in 2013, a TRW Automotive Endowed Research Award in 2012, an NSF Career Award in 2011 and Google Research Award in 2010. In 2002 he was awarded the Walker von Brimer Award for outstanding research initiative.

DAN YAMINS

DAN YAMINS
Daniel Yamins is a computational neuroscientist at Stanford University, where he's an assistant professor of Psychology and Computer Science (by courtesy), a faculty scholar at the Stanford Neurosciences Institute, and an affiliate of the Stanford Artificial Intelligence Laboratory. His research group focuses on reverse engineering the algorithms of the human brain, both to learn both about how our minds work and build more effective artificial intelligence systems. He is especially interested in how brain circuits for sensory information processing arise via the optimization of high-performing cortical algorithms for key behavioral tasks. Most recently, he has used performance-optimized deep neural networks to build neurophysiologically accurate models of higher visual and auditory cortex. He received his AB and PhD degrees from Harvard University, was a postdoctoral research at MIT, and has been a visiting researcher at Princeton University and Los Alamos National Laboratory. He is a recipient of the James S. McDonnell Foundation award in Understanding Human Cognition.

About the Stanford Artificial Intelligence Lab

Artificial Intelligence comprises the complete loop from sensing to perception, learning, communications, and action. Stanford's Artificial Intelligence Lab is devoted to the design of intelligent machines that serve, extend, expand, and improve human endeavor, making life more productive, safer, and healthier. These intelligent machines will learn everything about anything using multi-sensory information and the entire cyber world of information and knowledge.

The faculty members of the Stanford AI Lab are changing the world. Their research includes deep learning and machine learning; robotics; natural language processing; vision, haptics, and sensing; big data and knowledge base; and genomics, medicine, and healthcare. The approach is personalized, adaptive, anticipatory, communicative, and context aware.

2018 PRINTED ON RECYCLED PAPER

Stanford Artificial Intelligence Lab
Gates Computer Science Building
353 Serra Mall
Stanford, CA 94305

ai.stanford.edu

This brochure is the result of contributions by many people. We are particularly grateful to Kate Matney for graphic design and layout, Bill Bulkeley for interviews and writing, Tin Tin Wisniewski and Crystle Palafox Haarman for project management, leadership, and coordination, and Luanne Hakanson for administrative support and content management.