



**Stanford**  
ARTIFICIAL  
INTELLIGENCE  
LABORATORY



# CONTENTS

LETTER FROM OUR DIRECTOR .....	3
RESEARCHERS .....	5
MACHINE LEARNING / DEEP LEARNING.....	7
ROBOTICS .....	9
AI + SUSTAINABILITY .....	11
NATURAL LANGUAGE PROCESSING .....	13
GENOMICS AND HEALTHCARE.....	15
COMPUTER VISION .....	17
AUTONOMOUS VEHICLES.....	19
STANFORD AI4ALL .....	21
AI SALON .....	23
RESEARCH CENTERS.....	25

## Letter from our Director



**Chris Manning**

Director of Stanford Artificial Intelligence Lab

Dear Friends,

Welcome to the Stanford Artificial Intelligence Lab.

The Stanford Artificial Intelligence Lab (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 the 1950s, SAIL remains a proud leading intellectual hub for 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 of 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 AI research was conducted in academic laboratories. This is the time when 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 in the 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 in 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, to developing a swimming robot that can explore shipwrecks deep under water following haptic instructions from an archaeologist 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 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.



# The SAIL research lab and community are blossoming.

Equally important to the mission of technology innovation, we believe in the education of the next generation of technologists. As part of the Stanford Computer Science Department, SAIL faculty and researchers have been teaching students across the Stanford campus and from outside. It is well known that Computer Science has become the largest undergraduate major at Stanford since 2015. Less known is the fact that except for the introductory classes, all the most popular CS classes are now in AI: AI Principles and Techniques, Machine Learning, Natural Language Processing and Computer Vision. 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 Stanford AI4ALL, which was the first in a now growing number of summer camps for high schoolers. Stanford AI4ALL targets high school girls focusing on humanistic AI.

This is a historic 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. To address this need, SAIL is strongly involved in creating a new university-wide Human-centered Artificial Intelligence Initiative and we will have much more to tell you about that in the coming year.

We have come a long way, but the journey ahead of us is still long. None of today's intelligent 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 productive and healthier lives.

Sincerely yours,



Christopher Manning  
Director, Stanford Artificial Intelligence Lab  
Thomas M. Siebel Professor in Machine Learning



**Maneesh  
Agrawala**

Professor, Computer Science;  
Director, The Brown Institute For  
Media Innovation  
**Computer Graphics, Human-  
Computer Interaction,  
Visualization**



**John  
Duchi**

Assistant Professor, Electrical  
Engineering, Statistics  
**Machine Learning, Optimization  
And Statistics**



**Noah  
Goodman**

Associate Professor, Psychology,  
Linguistics (Courtesy),  
Computer Science (Courtesy)  
**Computational Psychology,  
Machine Learning, Linguistics**



**Mykel  
Kochenderfer**

Associate Professor,  
Aeronautics And Astronautics,  
Computer Science (Courtesy)  
**Machine Learning, Decision  
Theory**



**Gill  
Bejerano**

Associate Professor,  
Computer Science, Developmental  
Biology, Pediatrics (School Of  
Medicine)  
**Computational Genomics**



**Stefano  
Ermon**

Assistant Professor,  
Computer Science  
**Machine Learning, Statistics,  
Sustainability**



**Leonidas  
Guibas**

Professor, Computer Science  
**Geometric Computing And  
Learning, 3D Computer Vision,  
Computer Graphics**



**Daphne  
Koller**

Adjunct Professor,  
Computer Science  
**Machine Learning**



**Jeannette  
Bohg**

Assistant Professor,  
Computer Science  
**Robotics, Machine Learning,  
Computer Vision**



**Ron  
Fedkiw**

Professor, Computer Science  
**Computer Graphics, Computer  
Vision, Fluid Dynamics,  
Biomechanics**



**Tatsunori  
Hashimoto**

Assistant Professor,  
Computer Science  
**Machine Learning, Natural  
Language Processing**



**Anshul  
Kundaje**

Assistant Professor, Genetics  
(School Of Medicine),  
Computer Science  
**Computational Biology, Machine  
Learning**



**Emma  
Brunskill**

Assistant Professor,  
Computer Science  
**Reinforcement Learning,  
Interactive Learning**



**Chelsea  
Finn**

Assistant Professor,  
Computer Science,  
Electrical Engineering  
**Machine Learning, Deep Learning,  
Robotics**



**Thomas  
Icard**

Assistant Professor, Philosophy,  
Computer Science (Courtesy)  
**Machine Learning, Knowledge  
Bases, Logic**



**Monica  
Lam**

Professor, Computer Science,  
Electrical Engineering (Courtesy)  
**Natural Language Processing,  
Virtual Assistants**



**Ron  
Dror**

Associate Professor,  
Computer Science  
**Computational Biology**



**Michael  
Genesereth**

Associate Professor,  
Computer Science  
**Computational Logic**



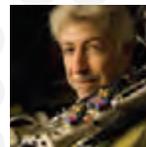
**Dan  
Jurafsky**

Professor, Linguistics,  
Computer Science  
**Computational Linguistics**



**James  
Landay**

Professor, Computer Science  
**Human-Computer Interaction,  
Natural Language Processing,  
Autonomous Vehicles**



**Oussama  
Khatib**

Professor, Computer Science  
**Robotics**



# RESEARCHERS, CONTINUED



**Jure  
Leskovec**

Associate Professor,  
Computer Science  
Data Mining, Machine Learning



**Chris  
Manning**

Professor, Computer Science,  
Linguistics; Director of Stanford  
Artificial Intelligence Lab  
Natural Language Processing,  
Machine Learning, Deep Learning



**Chris  
Piech**

Assistant Professor,  
Computer Science  
Machine Learning, Education



**Sebastian  
Thrun**

Adjunct Professor,  
Computer Science  
Robotics, Machine Learning



**Fei Fei  
Li**

Professor, Computer Science,  
Psychology (Courtesy);  
Co-Director, Human-Centered  
Artificial Intelligence (Hai)  
Computer Vision, Machine  
Learning, Computational  
Cognitive Neuroscience



**Andrew  
Ng**

Adjunct Professor,  
Computer Science  
Machine Learning



**Chris  
Potts**

Professor and Chair, Department  
of Linguistics, and Professor  
(Courtesy),  
Computational Linguistics



**Jiajun  
Wu**

Assistant Professor,  
Computer Science  
Computer Vision, Machine  
Learning, Cognitive Science



**Percy  
Liang**

Associate Professor,  
Computer Science, Linguistics  
(Courtesy)  
Machine Learning, Natural  
Language Processing



**Juan Carlos  
Nieves**

Senior Research Scientist,  
Computer Science  
Computer Vision, Machine  
Learning



**Christopher  
Ré**

Associate Professor,  
Computer Science  
Machine Learning, Theory,  
Databases



**Dan  
Yamins**

Assistant Professor, Psychology,  
Computer Science (Courtesy)  
Computational Neuroscience



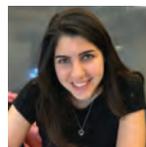
**Scott  
Linderman**

Assistant Professor, Statistics  
Machine Learning, Computational  
Neuroscience



**Vijay  
Pande**

Professor, Structural Biology,  
Computer Science (Courtesy)  
Statistical Mechanics, Bayesian  
Statistics, Biophysics,  
Biochemistry



**Dorsa  
Sadigh**

Assistant Professor,  
Computer Science, Electrical  
Engineering  
Machine Learning, Deep Learning,  
Robotics, Autonomous Vehicles



**Serena  
Yeung**

Assistant Professor,  
Computer Science (Courtesy),  
Electrical Engineering  
Healthcare, Computer Vision,  
Deep Learning



**Tengyu  
Ma**

Assistant Professor,  
Computer Science, Statistics  
Machine Learning, Deep Learning,  
Reinforcement Learning



**Marco  
Pavone**

Assistant Professor, Electrical  
Engineering (Courtesy)  
Robotics; Decision Theory;  
Autonomous Vehicles



**Ken  
Salisbury**

Professor, Computer Science  
Robotics



**James  
Zou**

Assistant Professor,  
Computer Science (Courtesy),  
Electrical Engineering  
Machine Learning, Computational  
Biology, AI for Health



**Silvio  
Savarese**

Associate Professor,  
Computer Science  
Computer Vision, Robotics,  
Geometry, Machine Learning

A key element of AI is enabling machines to learn to improve their perception, cognition, and ability to take actions from experience. Deep learning, a key sub-area of machine learning, is helping computers understand language, perceive the visual world, and be able to take actions better than ever before. A key spark for the deep learning revolution took place at Stanford in 2009: Prof. Fei-Fei Li released ImageNet, a database of 14 million labeled images that set a new benchmark for a machine's ability to understand the visual world. The next few years saw a torrent of progress as researchers from around the world developed new techniques to achieve superhuman performance on ImageNet.

Despite progress on benchmarks, fundamental challenges remain as machine learning moves from lab prototypes to embedded in our daily lives. Researchers at Stanford AI lab in machine learning are working to develop cutting-edge methods to fundamentally improve the robustness and efficiency of machine learning models.

#### ROBUSTNESS/RELIABILITY OF ML MODELS

While ML models obtain super human performance in labs, ML models deployed in the real world may fail in unexpected ways. For example, image classification models have higher error rates on images from outside Europe and North America, because most of images models are trained on are collected from these two continents, and the machine learning models focus on subtle associations in these images that do not generalize to other locations.

Prof. Emma Brunskill, Prof. Tatsunori Hashimoto, Prof. Percy Liang, and Prof. Tengyu Ma work on developing machine learning systems that perform well beyond the data and/or domains for which it was trained. For supervised learning, their works utilize approaches such as distributionally robust optimization or self-training to ensure that models do not learn brittle correlations that hold only in the training data. They do this by learning models that work well across all potential test distributions, including those which do not contain spurious and brittle associations. Liang and Ma have developed algorithms that allow the model to automatically adapt to a sequence of gradual shifting domains over time without using any additional labels for the new domains. (See more in Figure 1.) Ma's group works on understanding and characterizing various types of possible domain shifts and designing customized algorithms for them. In decision making settings, like reinforcement learning and bandits, Brunskill's group works on providing high probability guarantees with

FIGURE 1



In gradual domain adaptation we are given labeled data from a source domain and unlabeled data from a sequence of domains that shift gradually in distribution. For example, here the task is gender classification, and the source labeled domains are images from 1905. Given these labeled images, we are able to adapt the models to classify all the images in later years well using self-training, even though the style of portrait changes over time. (Here, blue = female, red = male, and gray = unlabeled data.)

respect to future performance given prior data in a way that can also satisfy desirable constraints, like fairness or safety.

#### SEQUENTIAL DECISION MAKING UNDER UNCERTAINTY

A fully intelligent ML system should not only predict very accurately, but also be able to make intelligent decisions based on the predictions. It not only requires the understanding of the sequential nature of the decisions—the ramifications of a prior decision to the future and the future decisions—but also requires precise characterization of the risk of each prediction and decision. Researchers at Stanford develop reinforcement learning and control-theoretic algorithms for sequential decision-making problems for applications such as robotics, education, healthcare, and finance.



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

Prof. Brunskill's lab seeks to create continually improving systems towards improving education and healthcare, and the core technical questions that arise in this quest. This includes formally understanding how quickly any agent can learn to make good decisions, how to perform counterfactual reasoning to leverage the enormous opportunity from vast records of past decisions made and their outcomes, and how to work together with humans-in-the-loop to best leverage the strengths of artificial and human intelligence.

Prof. Kochenderfer's group develops decision-making algorithms for situations where there is uncertainty in the current state of the environment and the outcomes of decisions, which is often modeled by is the Partially Observable Markov Decision Process (POMDP). A comprehensive Julia interface, POMDPs.jl, was developed for representing, defining, and solving discrete and continuous POMDPs.

Prof. Finn's lab studies how machine learning and reinforcement learning methods can be used to develop broadly intelligent behavior. This includes algorithms that produce behavior that can be quickly adapted to new environments and objectives, even when those settings are never experienced during training. It further involves algorithms that enable robots to learn autonomously without human supervision and requires agents to understand the intentions of humans.

Prof. Tengyu Ma has been working on improving the sample efficiency of RL algorithms—a big challenge of applying RL to real-world problems—by the model-based reinforcement learning and imitation learning. The model-based RL algorithms simultaneously learns a neural network model as a simulator of the world and the policy that outputs the decisions based on the simulations.

Prof. Stefano Ermon, Prof. Percy Liang, and Prof. Tengyu Ma's labs also work on quantifying the uncertainty of the predictions of machine learning models. This allows us to make more informed and prudent decisions based on the predictions.



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.

## IMPROVING DATA-EFFICIENCY

Machine learning methods are often data hungry—they require a large amount of data with labels, which can be expensive to collect. Improving the efficiency of utilizing limited data is the key towards broader applications of machine learning.

Prof Re's lab has examined how one can program machine learning systems with weaker forms of supervision than traditional machine learning systems, in a system called Snorkel. Engineers have used weaker forms of supervision, like heuristic rules or existing knowledge bases, to generate dramatically cheaper—but lower quality—training data for years. However, machine learning engineers could not reliably use weak sources of supervision as often the signals were too weak and it is difficult to combine many sources of weak supervision that have different quality or are highly correlated with one another. A key idea in Snorkel is a new way to learn the quality and correlations of sources of input, without traditional labeled data. Although still nascent, this new way to build program has caught on with industry and scientific collaborators. Thanks to Stanford collaborations with industry, you've probably used a Snorkel-inspired system in the last few hours. ▣

**R**obotics 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, electrical engineering, bioengineering and even cognitive science and 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 perceive their environment and adapt to their tasks so they can pick up iron bars or porcelain vases with the same grippers.

Roboticians 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 center, 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 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 center 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.

Humans have the innate ability to "read" one another. When people walk in a crowded public space such as a sidewalk, an airport terminal, or a shopping mall, they obey a large number of (unwritten) common sense rules and comply with social conventions. For instance, as they consider where to move next, they respect



# Increasingly, robotics researchers are focused on teaching robots to learn from their experiences.

personal space and yield right-of-way. The ability to model these “rules” and use them to understand and predict human motion in complex real world environments is extremely valuable for the next generation of social robots. Researchers in the Stanford Vision and Learning (SVL) Lab have developed a self-navigating mobile manipulator robot, “Jackrabbot”, that is capable of learning and interacting with people in pedestrian spaces.

In addition to social navigation, researchers in the Intelligent and Interactive Autonomous Systems (ILIAD) group at Stanford study the design of robotics algorithms that safely and reliably interact with people in a diverse set of collaborative settings.

Researchers at ILIAD look into building computational and predictive models of human behavior, and leveraging such learned models in better interaction, collaboration, and coordination with robots. They formalize the interactions between humans and robots in domestic environments as well as in human-robot teaming. Going beyond interaction, Stanford roboticists are interested in social implications of automation on humans. Specifically, how autonomous systems can influence societal objectives, e.g., the role and influences of autonomous vehicles in mixed-autonomy roads, or emergent behaviors in human-robot teams.

Another fundamental challenge in robotics is generalizability. Researchers at the Interactive Perception and Robot Learning (IPRL) lab at Stanford are studying the underlying principles of robust sensorimotor coordination for problems such as autonomous robotic manipulation. Robotics researchers at Stanford develop algorithms and systems that unify reinforcement learning, control theoretic modeling, and visual



scene understanding to teach robots to perceive and to interact with the physical world.

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. ▣

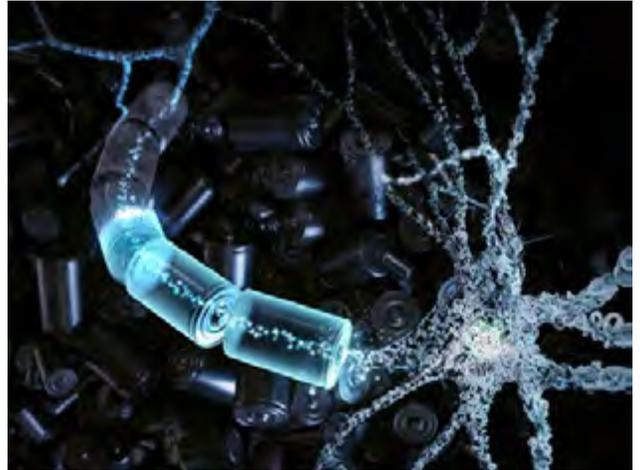


Sustainability challenges such as climate change, food insecurity, and global poverty eradication are among the most pressing issues our society is facing today. Researchers at SAIL are collaborating with domain experts to investigate how techniques from Artificial Intelligence can provide more scalable, computational, and data-driven solutions to these problems.

In the world of sustainable development, decisions are often made without data. As Kofi Annan wrote in 2018, “data gaps undermine our ability to target resources, develop policies and track accountability [...] If you can’t see it, you can’t solve it.” Imagine if governments and citizens could see exactly where social programs were performing, and where they needed improvement. How many more lives could be transformed? New developments in remote sensing and communication technologies are creating many cheap, unconventional data streams that contain information relevant to poverty, health, governance, and broader economic outcomes throughout the world.

Recent efforts at Stanford have begun to explore the potential of big data for development with promising results. For example, Marshall Burke, David Lobell and Stefano Ermon demonstrated that information extracted from high resolution satellite images can be leveraged using machine learning to estimate poverty. Such granular poverty maps, if available at scale and over time, could help aid organizations and policymakers distribute funds more efficiently and enact and evaluate policies more effectively. They would also provide real-time measures of changing economic conditions that would advance social scientific studies of poverty dynamics. This project was recognized as one of the 10 World Changing Ideas of 2016 by Scientific American. Additionally, similar techniques can be used to forecast agricultural productivity from space both in the U.S., where the method is as accurate as USDA forecasts, and in developing countries, where such data is key to address food security challenges. The team is working on extending these capabilities to predict famines, mapping refugee camps, and track health-related outcomes at scale. Given how poorly these outcomes have traditionally been measured, these new data could lead to radical improvements in understanding how to accelerate sustainable development.

Led by Prof. Andrew Ng in the Stanford Machine Learning Group, the AI for Climate Change Bootcamp convenes students across campus to work on pressing problems in climate change using artificial intelligence, including



work on methane prediction from noisy sensors and wind turbine localization from satellite imagery at a global scale. The organization partners with environmental experts to identify important problems at this intersection. These collaborators include the Environmental & Civil Engineering and Earth System Sciences departments at Stanford, government organizations such as NCAR, non-profits such as the Environmental Defense Fund, and commercial entities such as Descartes Labs and Open Climate Fix.

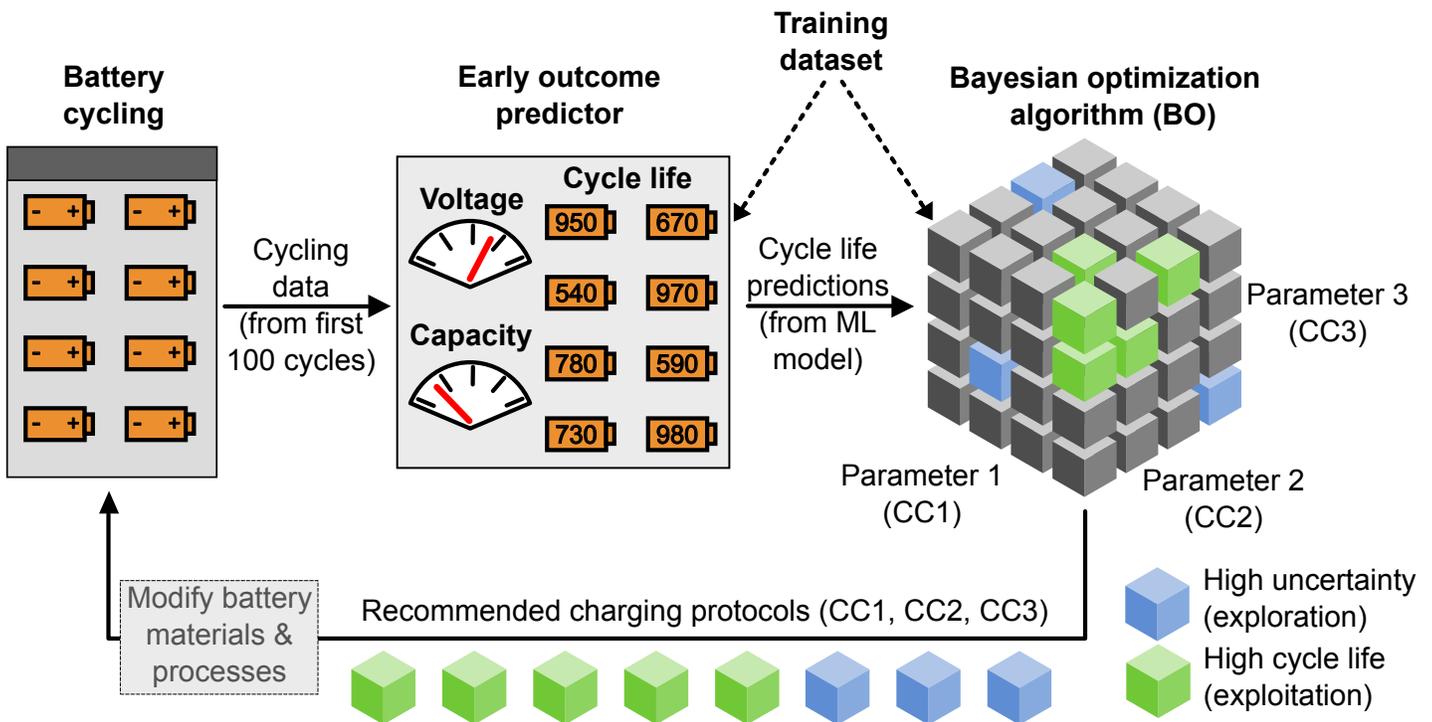
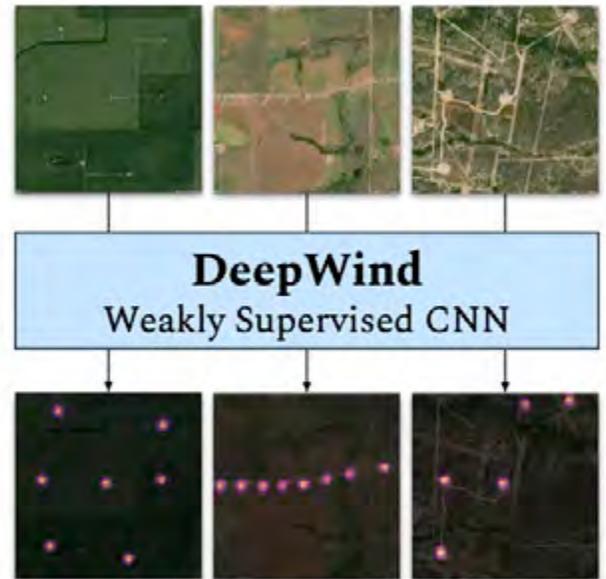
Wind energy is being adopted at an unprecedented rate. Hundreds of thousands of locations of global wind energy sources, however, remain undocumented, which significantly impedes their integration into power systems. These locations are particularly critical to a variety of stakeholders, such as: (i) wind developers to identify the best new areas for deployments; (ii) electricity grid operators to integrate renewable energy, to perform real-time system operation, and to plan capacity expansion; (iii) utilities and city planners to plan the management of local demand; and (iv) policymakers to design and estimate the impact of incentives and other contracting policies. Towards the goal of mapping global wind energy infrastructure, a team led by Andrew Ng developed deep learning models to automatically localize wind



# Recent advances in lithium-ion batteries are driving the adoption of alternative energy.

turbines in satellite imagery. These models enable the efficient construction of a complete, up-to-date, global database of wind energy infrastructure, ultimately assisting with the management and adoption of wind energy worldwide.

Recent advances in lithium-ion batteries are driving the adoption of alternative energy technologies. However, a major bottleneck is maximizing battery lifetime, which can take months to years to assess. Hence, a key challenge in battery development is to reduce both the number and the duration of the experiments required during materials selection, manufacturing, and deployment. A team led by Stefano Ermon and Will Chueh (Materials Science) designed an AI-driven closed-loop system for efficient, near-autonomous optimization of fast charging policies for lithium-ion batteries. The system can autonomously design and perform hundreds of experiments in parallel. Data collected is combined with machine learning to refine the hypothesis space and guide future experiments in a closed-loop way. The system can identify charging policies that outperform existing ones with a 15-fold reduction in optimization time. ▣





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

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.

The advancement of NLP systems in areas such as conversational agents and question answering has begun to close the gap between humans and machines on a wide variety of benchmarks. Despite this, it’s been suspected that such gains become gaming benchmarks, rather than understanding of the underlying task: conversational agents are engineered to return generic responses as a way to mask ignorance, and natural language inference systems for identifying logical relations cheat by learning to identify flaws in the data annotation process.

Stanford Professor Tatsu Hashimoto uses approaches from robust machine learning to identify and fix these types of failures. His recent work uses the idea of adversarial evaluation—ensuring that a model’s outputs are indistinguishable from a human’s—as a way to evaluate natural language generation systems. To prevent a model from exploiting evaluation flaws, his work uses an approach called distributionally robust optimization to ensure that an NLP system performs well not just on a single test set, but on an entire family of potential test sets. This approach captures the intuition that systems which game a benchmark may do well on a test set similar to the training data, but are likely to fail when extrapolating to a very different test set.

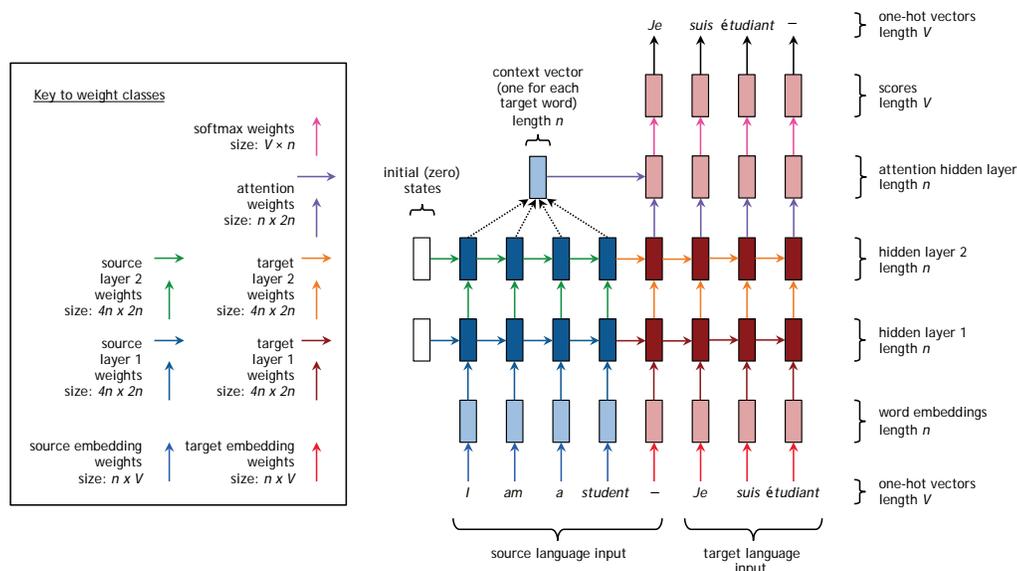
Jure Leskovec has applied natural language processing to online chats at a crisis hotline, to learn what successful counselors do to aid troubled callers.

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.

Having a general conversation is much more difficult and still an unsolved problem. Nevertheless, students working with Chris Manning are attempting just that: They are exploring neural dialogue strategies for a conversational socialbot as part of the Alexa Prize competition. ■



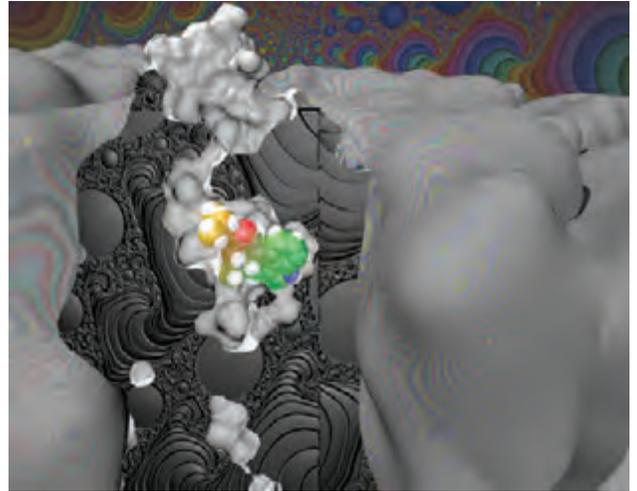
**G**enomics is where medicine meets computer science. Mapping the human genome is something only machines can do. And once a genome has been mapped, AI techniques are being used to decode it.

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.

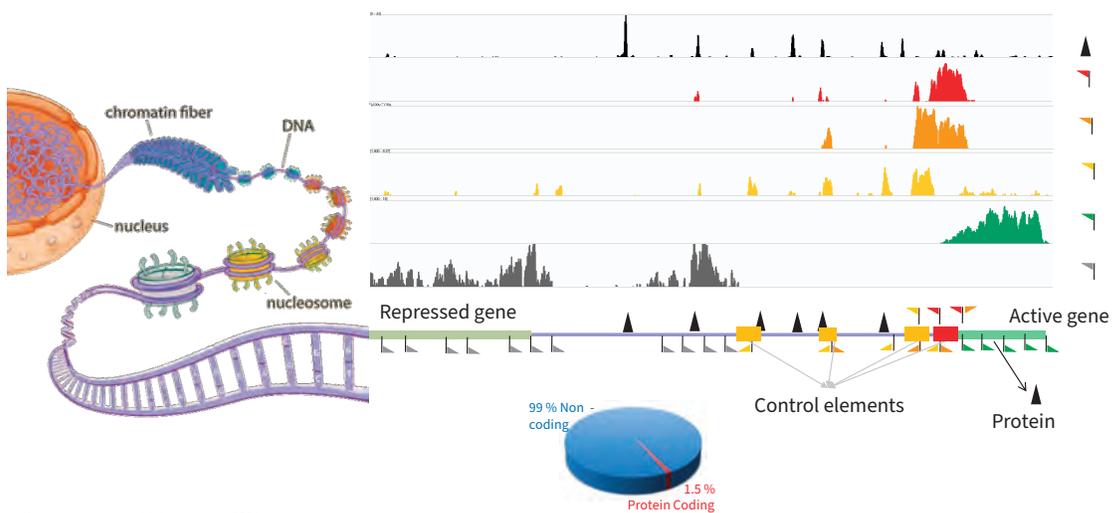
Starting with the genomic code reverses the way people traditionally view the human body and its diseases.



Physicians and clinical researchers usually start by trying to understand body parts such as organs and systems. Then they try to understand 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.



## Biochemical markers of cell-type specific functional elements



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

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.

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.

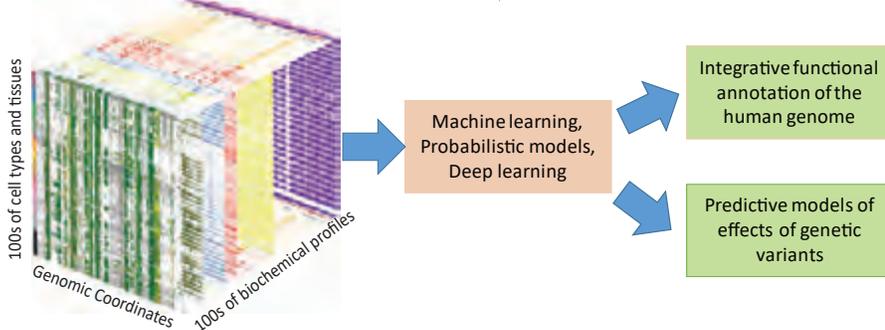
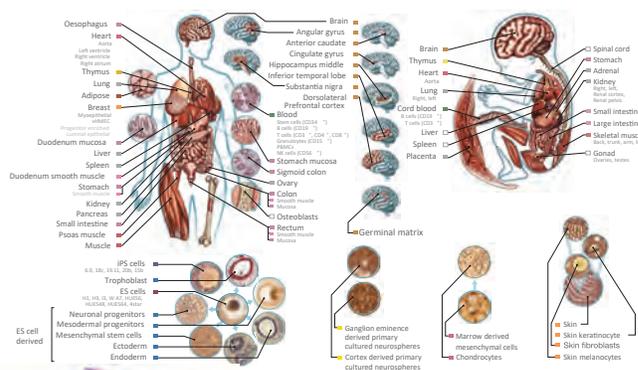
This research avenue has become much more fruitful with the development of CRISPR-Cas9 gene editing technology. CRISPR allows researchers to go to a specific point on a genome and cut it or block its enzymatic activity in order to see how the loss

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.” □



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.

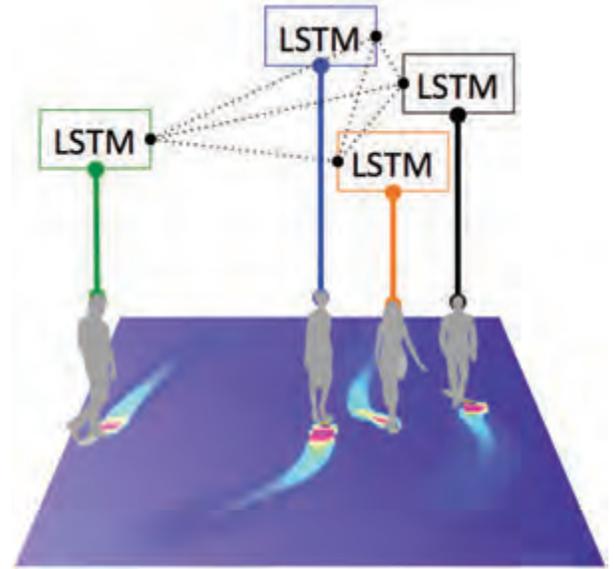
Humans are producing and storing an exploding amount of visual information from visual sensors like mobile phones or surveillance cameras that take beautifully detailed photos and videos with sound. But computers continue to struggle to make sense of it all. Recording the series of numbers that represent 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.

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.



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 professor 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



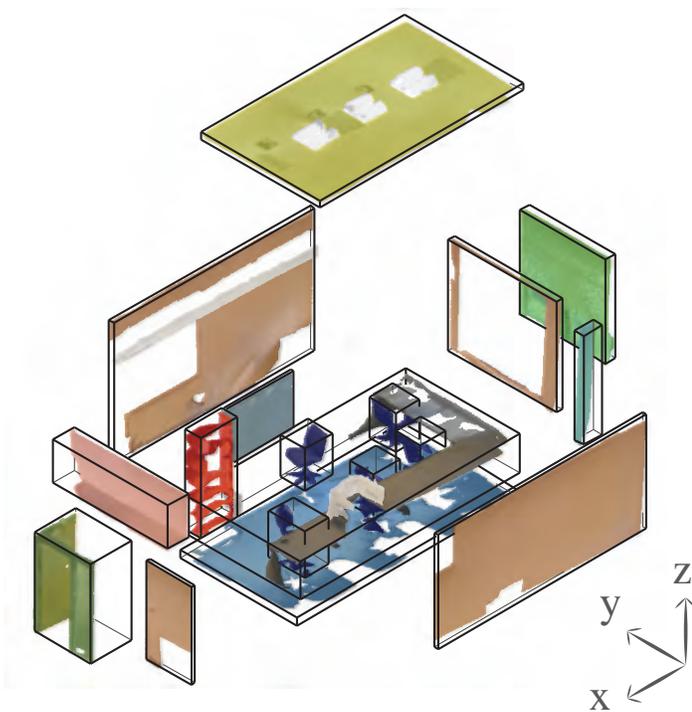
# Humans look with their eyes, but they see with their brains.

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.

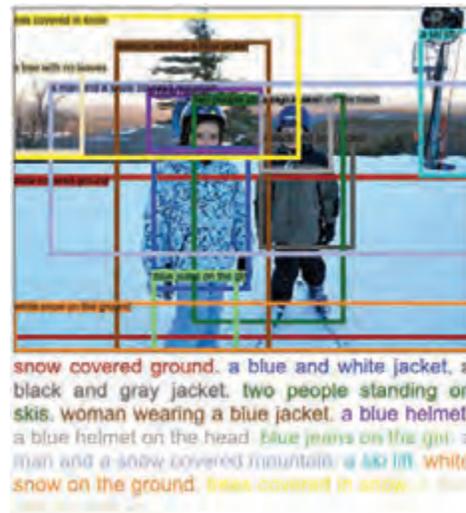
SAIL researchers are also trying to use machine vision to provide expertise that humans can't. One project involves construction-project monitoring. A vision system learns a building by scanning a detailed 3-D model of the completed building. Then it goes through the building week-by-week while it is 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. ▣



- |       |   |        |   |        |   |          |   |
|-------|---|--------|---|--------|---|----------|---|
| board | ■ | floor  | ■ | table  | ■ | door     | ■ |
| chair | ■ | wall   | ■ | column | ■ | bookcase | ■ |
| beam  | ■ | window | ■ | sofa   | ■ | ceiling  | ■ |



**D**riverless 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 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.

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



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.

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 professor 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. A similar approach may be needed for the



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

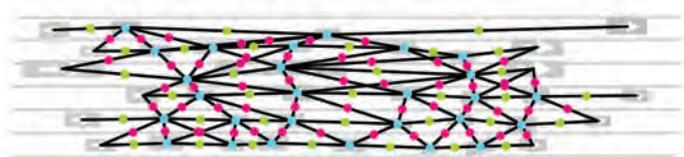


As part of the Toyota SAIL project, students of Mark Cutkosky and Chris Gerdes test applications of a skin stretch haptic steering wheel on experimental testbed X1 for navigation cues, collision avoidance warnings, and autonomous previews.

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-or-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 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.

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.



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

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.



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.

Past work on decision-making by Mykel Kochenderfer, an AI professor 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. 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-or-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. ▣

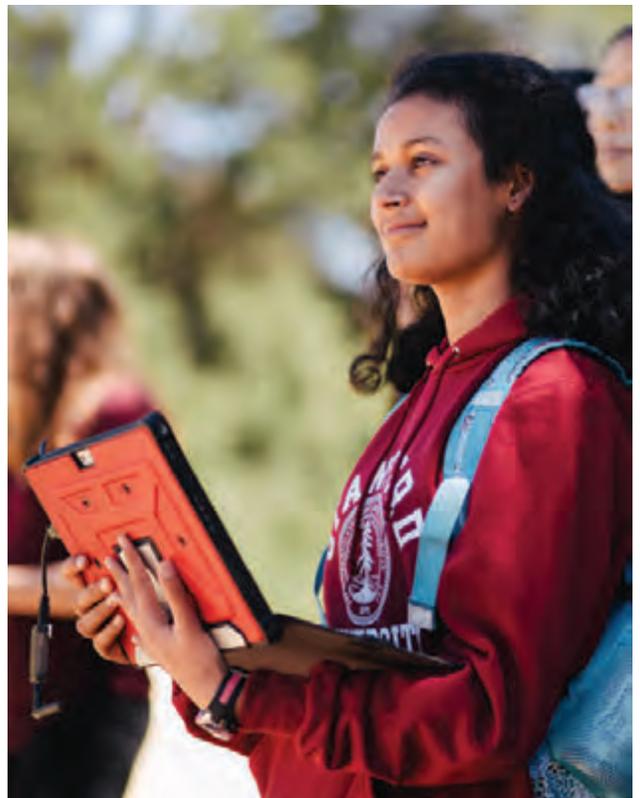
As the field of AI continues to make a bigger impact in the world, researchers and educators at SAIL believe that to develop the most inclusive, humanistic, and benevolent technologies, it is imperative that the field of AI includes students, researchers, and technologists from all walks of life. With 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, signaling 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.

The original SAILORS and the present Stanford AI4ALL is 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 Professor Olga Russakovsky (Assistant Professor at Princeton University), and Dr. Rick Sommer (Executive Director of Stanford Pre-Collegiate Studies). Now renamed, this program continues to be based in the Stanford Artificial Intelligence Laboratory, but with an updated curriculum developed by AI4ALL, an education nonprofit organization dedicated to training the next generation of AI researchers. AI4ALL was also founded by Dr. Fei-Fei Li and Dr. Russakovsky. Juan Carlos Niebles, Senior Research Scientist at SAIL, was named as the Director for Stanford AI4AL 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.

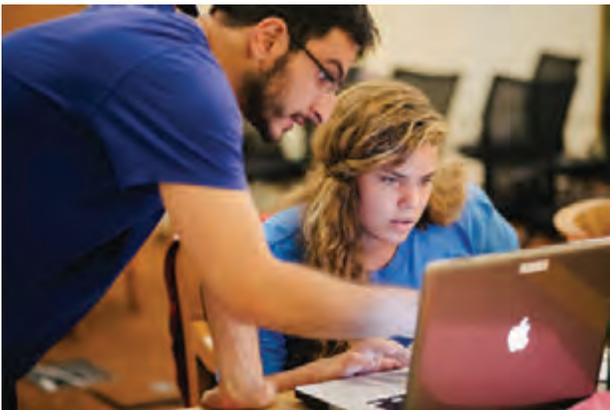
Renowned computer scientists Dr. Ruzena Bajcsy (Robotics Professor at UC–Berkeley) and Dr. Maria Klawe (President of Harvey Mudd College) have delivered the keynote speeches in respective years. In 2015, two undergraduate students of Computer Science,



# AI will change the world. Who will change AI?

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.

In recent years, the Stanford AI4ALL program has evolved into a residential program with 32 students that requires tuition. However, Stanford AI4ALL does offer financial aid in order to best serve a global, diverse audience. AI4ALL has evolved into a nationwide organization, with programs now at over a dozen other universities, providing similar programs aimed at increasing diversity and access to computer science research. AI4ALL is now partnering with several other universities, creating similar programs aimed at increasing diversity and access to computer science research. ▣



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 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.

In December 2018 we had a wide-ranging and widely attended panel on AI and the future of work and human endeavors, with Kai-fu Lee, MIT Professor Erik Brynjolfsson and Stanford professor Susan Athey. We hosted the head



of Google's AI division, Jeff Dean, in March 2019 for a discussion on the importance of flexibility and scaling for modern real-world machine learning systems. In January 2020, renowned cognitive psychologist Barbara Tversky and our very own Fei-Fei Li participated in a lively Salon on the relationship between artificial and human cognition and intelligence, and how the two fields of study can benefit from each other.

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. ▣

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



# Corporate-Sponsored Research Centers and Initiatives

The Stanford AI Lab values open collaboration with industry. AI is an integral part of many exciting business and consumer tools such as speech recognition, semantic search, recommendation systems, machine translation, 3D sensing, safety and more. Industry engagement through the SAIL-Toyota Center for AI Research and the SAIL-JD AI Research Initiative helps faculty and researchers demonstrate and develop translational technologies with real-world applications.

The SAIL-JD AI Research Initiative was established in 2017 to develop and apply artificial intelligence to supply chain management, intelligent customer service, automation in warehouses and delivery as well as other areas of innovation and automation in e-commerce. Research at the SAIL-JD Initiative involves collaborations between faculty in computer vision, natural language processing, information networks and machine learning. Projects have focused on knowledge representation and learning, retail, and robotics for logistics.

## KNOWLEDGE REPRESENTATION AND LEARNING

One key area of collaboration has been around knowledge representation for understanding customers and interactions. The goal of this work is to create a multiscale behavioral model predict future behavior using temporal knowledge graphs and deep embedding methods. Once constructed, this knowledge graph will provide intelligent customer service, recommendations and predictions from structured query and natural language.

Additional projects have focused on advancing the field of Deep Reinforcement Learning that can improve sample complexity in order to apply in cases where collecting environmental interactions is costly.

## RETAIL

With the fast-growing e-commerce industry, online cloth shopping has become increasingly popular across the world, and in particular a large portion of the transactions happen on mobile devices. One of the biggest challenges that e-commerce companies face for selling clothes is to help customers establish realistic expectations of the garments' fit and look on their bodies via virtual fit. The goal of this work is to combine



deep learning with physics-based models to build more realistic virtual clothing.

## AUTONOMOUS VEHICLES AND DRIVING

Future intelligent vehicles require a deep and accurate understanding of the environment under diverse and challenging conditions. This includes recognition of animate and inanimate objects, their properties, and their estimated behaviors. It is particularly important to predict the behavior of other drivers, bicyclists, and pedestrians. Future intelligent vehicles must also look inward and understand the state of their own driver. Such detailed perception of the environment can then be leveraged to perform robust decision-making under the large uncertainty of the real world.

Research in autonomous vehicles includes models for safe navigation through dynamic environments, trajectory forecasting, developing software for reliable machine learning, cloud-aided computer-vision



systems, and provably and practically safe algorithms for autonomous driving.

### ROBOTICS FOR LOGISTICS

Research in logistics focuses on two specific applications; warehouse robotics and robotic delivery. Research in autonomous logistics combines the skills of a diverse team of robotics, perception, and machine learning. Robotics for Logistics is at the frontier of robotics and artificial intelligence, requiring autonomous manipulation of unconstrained objects, often in proximity or in collaboration with humans. Research tasks include pick and place, packing/unpacking, sorting, delivery/retrieval, human activity recognition and understanding, and human and agent intent prediction. Major technical challenges include:

- Mobile robotic hardware including grippers, arms, and on-board visual and haptic sensing
- Manipulation and perception capabilities adaptive to an ever-increasing set of objects through lifelong learning
- Controllers for safe interaction in collaborative and distributed tasks between robots
- Algorithms for understanding high-level human behaviors, intentions and goals

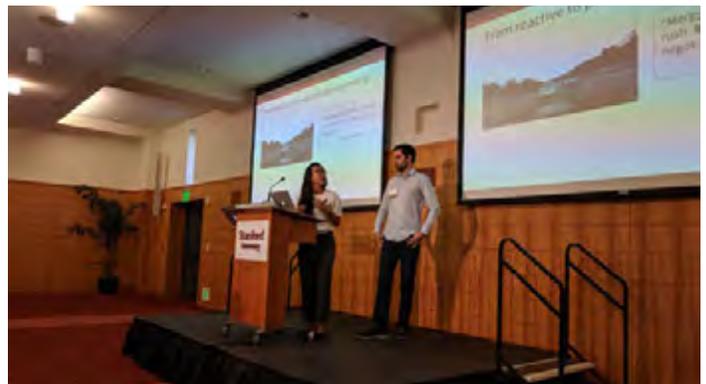
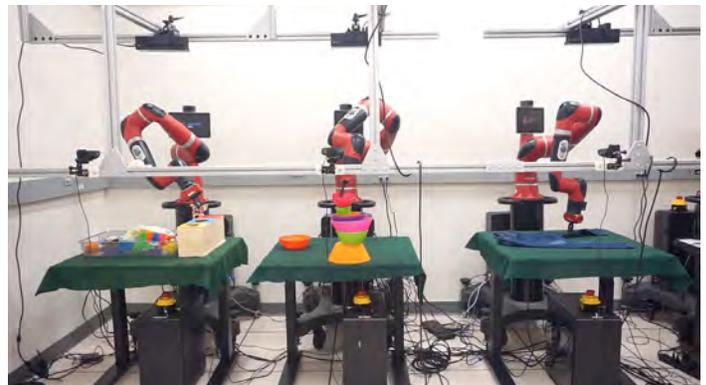
The SAIL-Toyota Center for AI Research was established in 2015 to develop innovative and impactful approaches and algorithms for future intelligent vehicles and agents. The SAIL-Toyota Center brings together researchers from visual computing, machine learning, robotics, human-computer interactions, intelligent systems, decision-making, natural language processing, and dynamic modeling and design. Research at the center focuses on the complete loop from sensing to perception, learning, communication, and action. Projects develop technologies to advance the fields of autonomous vehicles, robotics and human-centric research.

### ROBOTICS

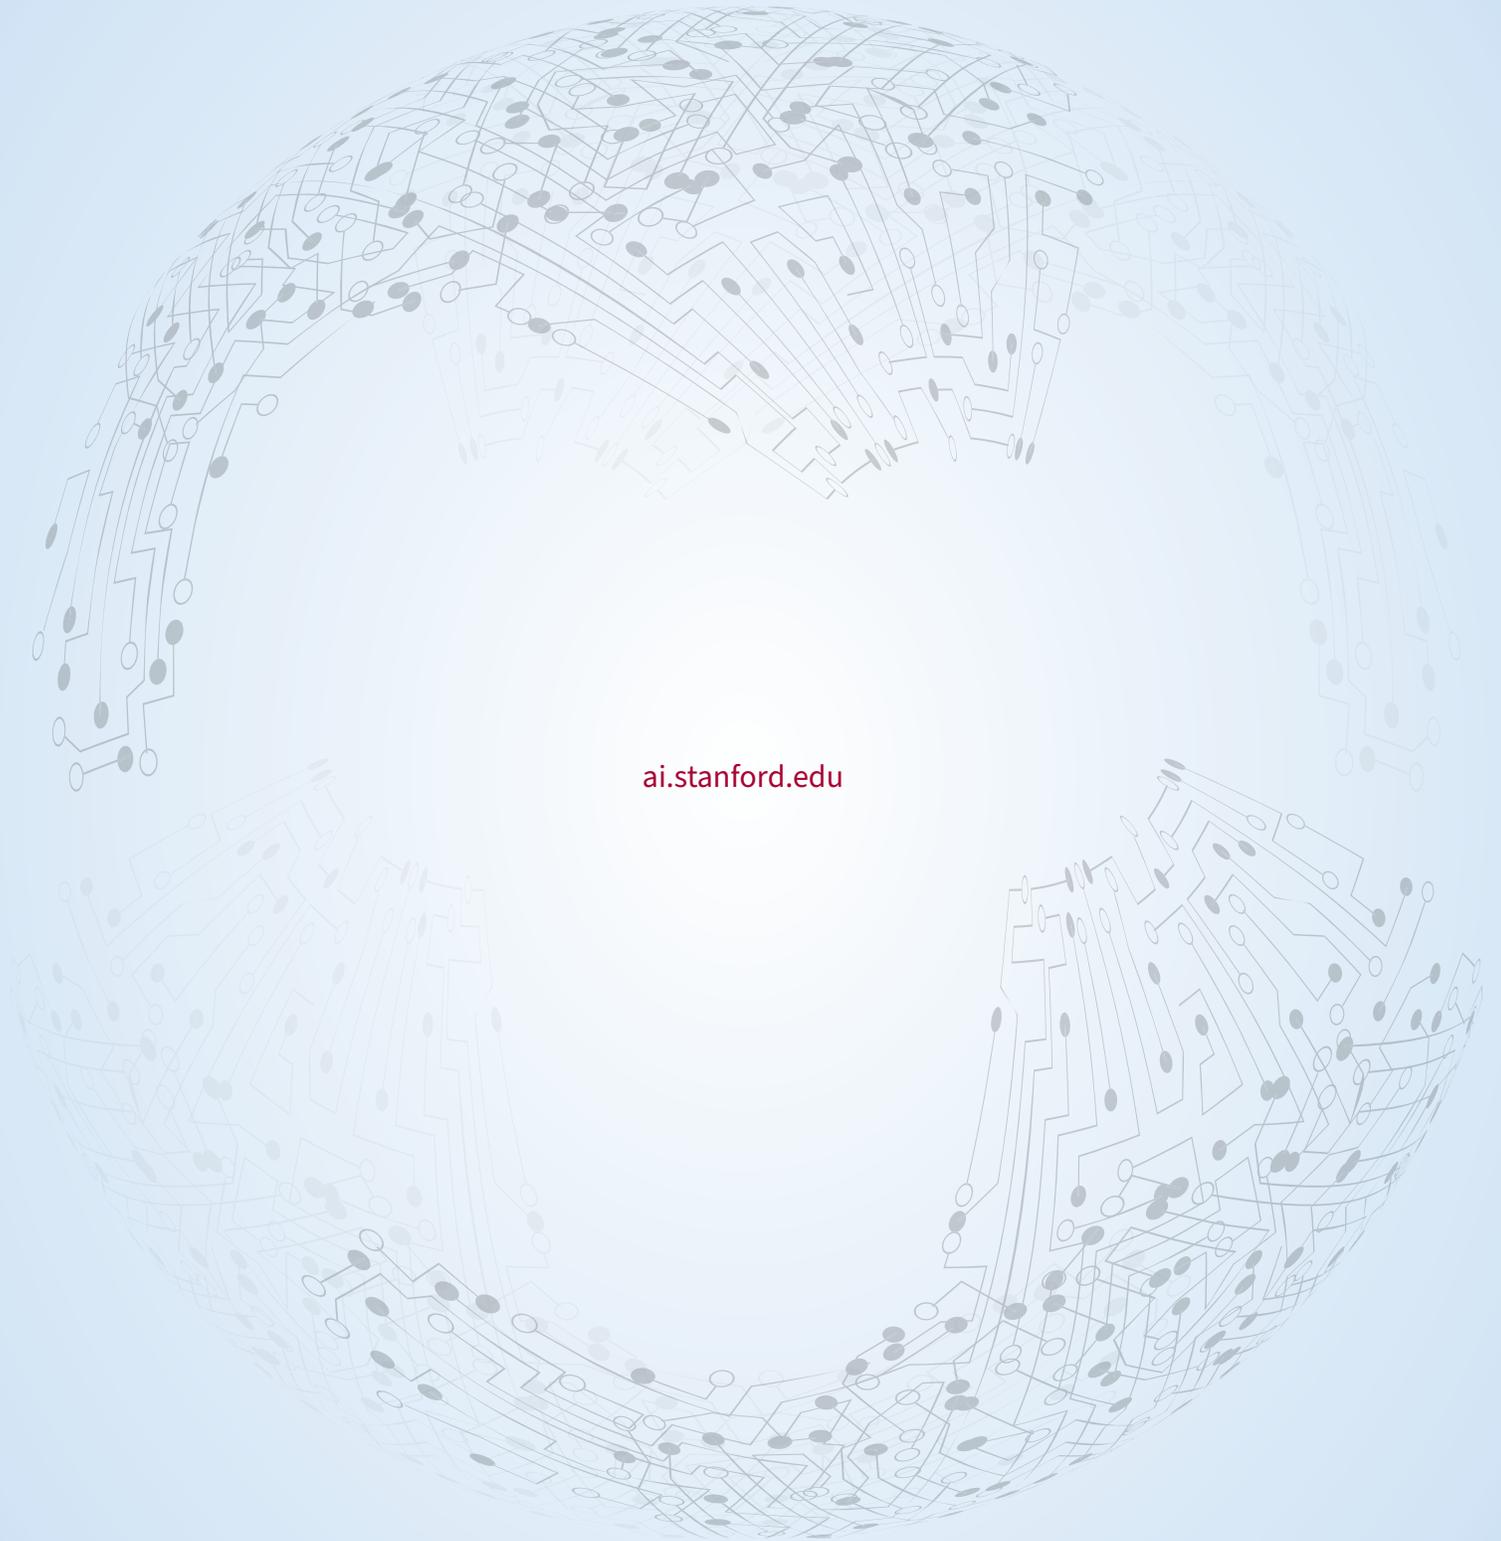
Research in the SAIL-Toyota Center also has applications outside the automotive environment, allowing people to live more independently, particularly in the latter years of life. Research in future robotics includes advanced manipulation in real world home environments, transfer learning mechanisms to efficiently handle new tasks, hybrid imitation learning for planning and scene understanding, learning physical attributes of unknown objects, shared autonomy and teleoperation scenarios, soft robots, haptics, and improving grasping and contact sensing capabilities.

### HUMAN-CENTRIC

Future robotic and mobility systems will not exist in controlled environments; these autonomous agents will



interact frequently with humans. As such, a large focus on the center is on human-centered artificial intelligence. Projects in this area include developing engagement learning methods for data collection via in-vehicle conversational agents, identifying and classifying driver behavior types, modelling the latent structure behind human decision making in the context of multi-agent mixed-autonomy interactions, using natural language to better scale machine learning models, developing data-efficient ML algorithms, and reducing the cost of collecting human labels to improve prediction accuracy. ▣



[ai.stanford.edu](https://ai.stanford.edu)

## About the Stanford Artificial Intelligence Laboratory

Artificial Intelligence comprises the complete loop from sensing to perception, learning, communications, and action. Stanford's Artificial Intelligence Laboratory 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 Laboratory 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.

[ai.stanford.edu](http://ai.stanford.edu)

