HAI

Stanford University
Human-Centered
Artificial Intelligence

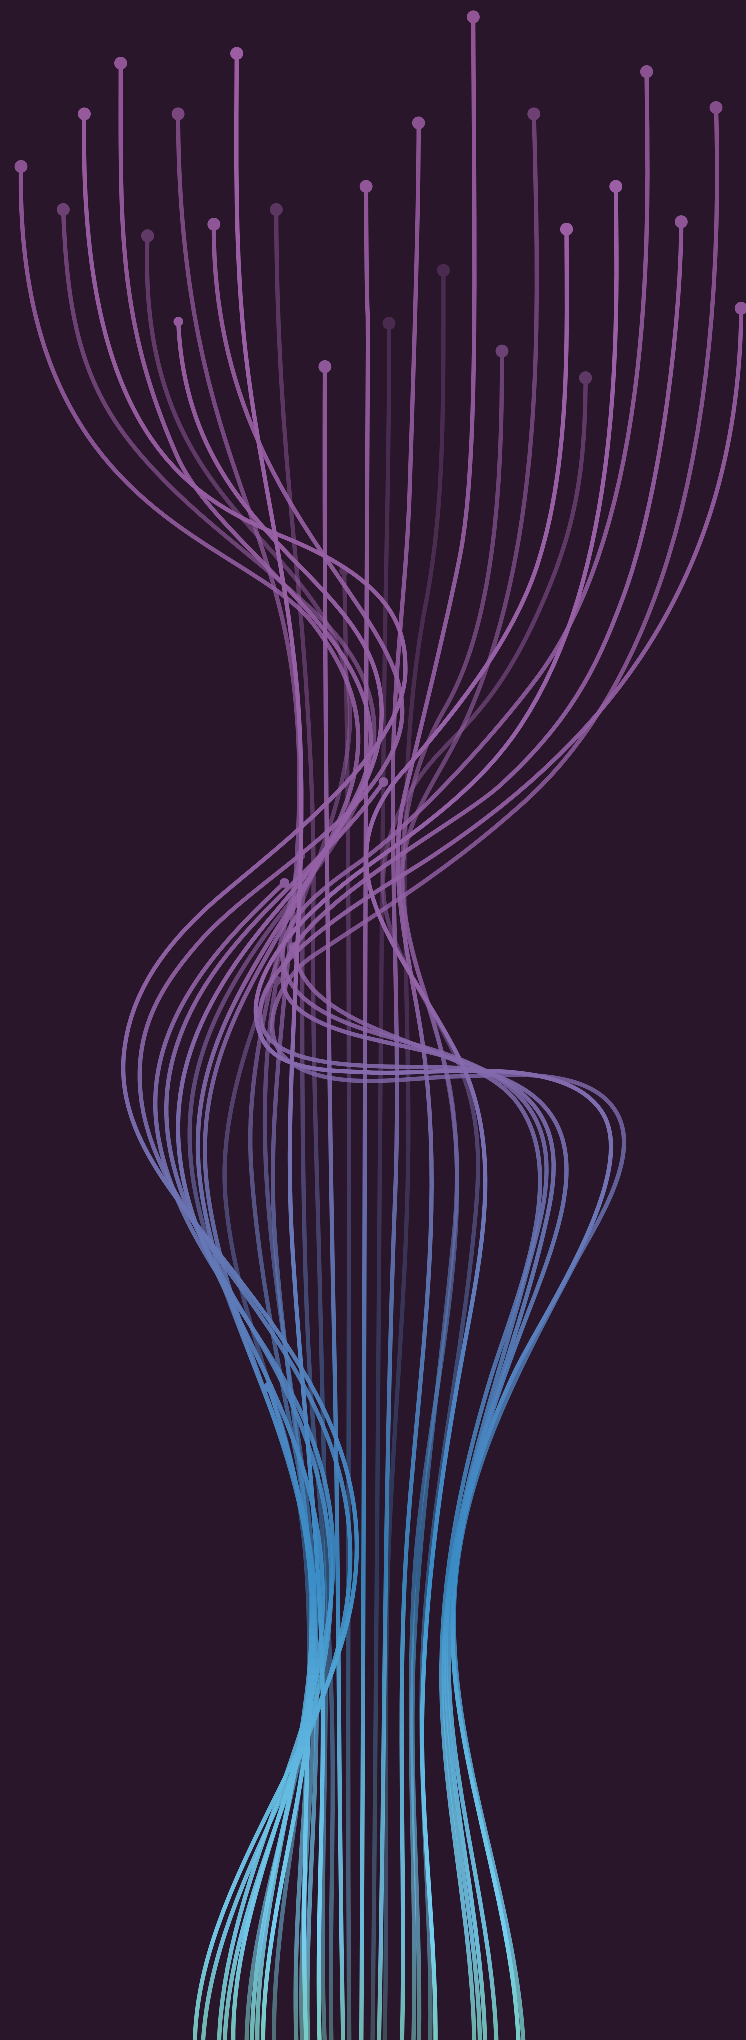

JUNE 2022
INDUSTRY BRIEF

An initiative of HAI Industry Programs & Partnerships

Robotics and AI

Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Introduction

At the Stanford Institute for Human-Centered Artificial Intelligence (HAI), we are focused on leveraging artificial intelligence to improve human wellbeing. Robotics has a great potential to augment human ability, and AI plays a key role in making these robots adaptable, generalizable, and autonomous so that they can truly be useful. With AI, robots are given the ability to learn and adapt so that they can work collaboratively alongside humans and other robots in real-world environments.

In the following pages you can read about some of the research happening at Stanford that leverages AI for robots, allowing them to learn from, collaborate with, and improve the standard of living for all humans. This research spans different industries — including health care, transportation, and home services — and touches upon both consumer and enterprise applications.

Many of us already interact with robots on a day-to-day basis, and in the future the frequency of these interactions will only increase. It is important for industry leaders to be aware of the newest developments in robotics so that they can adopt the technologies that benefit their organizations and their customers. We hope that the following exploration of human-centered AI in robotics provides a clear look into the many applications of these cutting-edge technologies as well as their societal importance and implications.

HAI's mission is to advance AI research, education, policy and practice to improve the human condition. To learn more about HAI, visit hai.stanford.edu

Dorsa Sadigh, Assistant Professor in the Computer Science Department at Stanford University

Learning, Adapting, and Influencing in Human-Robot Interaction

As robots become increasingly prevalent and capable, their interaction with humans and other robots is inevitable. Therefore, a key question we investigate is: how can we effectively build a model of the human partner, anticipate their behavior, and develop robot algorithms to seamlessly coordinate and collaborate with humans. One of the key challenges is that humans adapt their behavior as they interact with robots or other intelligent agents. We thus need to discover the 'changing' strategy of the human in response to autonomy, and identify what norms and conventions emerge through repeated interactions between humans and robots.

Erik Brynjolfsson, Jerry Yang and Akiko Yamazaki Professor and Senior Fellow at the Stanford Institute for Human-Centered AI (HAI), and Director of the Stanford Digital Economy Lab

AI and humans: augmentation versus substitution

AI-powered robots represent a huge opportunity for productivity growth and boosting living standards. But many of the current robotics developers are focused solely on automation; building a robotic substitute for tasks done by humans. In some cases, substitution is the best option - leading to automation of dangerous or tedious work, or in situations where human labor is scarce. But excessive automation can lead to negative impacts for human workers, as well as missed opportunities for larger long-term financial and productivity gains. As machines become better substitutes for human labor, workers lose economic and political bargaining power and become increasingly dependent on those who control the technology.

In contrast, when machines are focused on augmenting humans rather than mimicking them, then humans retain the power to insist on a share of the value created. What's more, augmentation creates new capabilities and new products and services, ultimately generating far more value than purely substitution-focused technology.

While both automation and augmentation technologies can be enormously beneficial, there are currently excess incentives for automation among technologists, business executives, and policymakers. The solution is not to slow down technology, but rather to eliminate or reverse the excess incentives for automation over augmentation.

Robotics Activities on the Stanford Campus

Stanford Vision and Learning Lab

Stanford Robotics Lab

Stanford Intelligent Systems Laboratory

Intelligence through Robotic Interaction at Scale Lab

Stanford Intelligent and Interactive Autonomous Systems Group

Interactive Perception and Robot Learning Lab

Collaborative Haptics and Robotics in Medicine Lab

Assistive Robotics and Manipulation Lab

... and more!

*Listed in no particular order.

Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Robotics and AI

1 Adaptive Learning & Optimization

2 Learning from Human Demonstration

3 Assistive Robotics and Human Robot Interaction

4 Medical Robotics and Human Augmentation

5 Autonomous Vehicles and Robot Navigation

1 ROBOTICS AND AI
ADAPTIVE LEARNING & OPTIMIZATION

WHAT'S NEW?

Intelligent robots have high potential for impact in a wide range of industries including manufacturing, healthcare, transportation, and many more. For these applications, robots need to be capable of many skills and be adaptable to dynamic environments but it's unscalable to train them for each individual task and situation. Researchers have enabled robots to learn on the job and become adaptable to sensor input, allowing them to learn from their own mistakes and act even more autonomously.

WHY DOES THIS MATTER NOW?

The robotics industry has been growing steadily since its start in the 1960s, with advancements in AI posed to improve robotic capabilities even further. The COVID-19 pandemic further inflated industry interest in automation so that human contact and exposure to the virus could be minimized. In 2020, the industrial automation equipment market was worth $16.5 billion globally. For certain industries like groceries, augmenting their workforce with automation allowed them to keep up with higher demand and new sanitation requirements. It's clear that robotic automation will greatly change the landscape of many industries, so organizations must consider the ways that intelligent robots can improve their current processes, especially with robots that can adapt to highly dynamic situations.

Read more: https://www.nytimes.com/2020/04/10/business/coronavirus-workplace-automation.html
https://www.ft.com/content/fa81806d-8ddf-4f2a-9a5a-de6c96844373
https://www.mckinsey.com/~/media/mckinsey/industries/advanced%20electronics/our%20insights/growth%20dynamics%20in%20industrial%20robotics/industrial-robotics-insights-into-the-sectors-future-growth-dynamics.ashx

EYE ON CAMPUS

Researchers on campus are...

- Creating scalable algorithms for decentralized and collision-free coordination among robots working together on sequential assembly tasks.
- Developing a new algorithm so that robots can do Task and Motion Planning in real world environments so that they can adjust to environmental changes and inaccurate perception without being slowed down and having to re-plan.
- Enabling robots to learn from failed object-fitting attempts and distinguish between failure types thus allowing robots to more effectively accomplish assembly tasks.
- Giving robots the power of touch so that robots can use both haptic and visual feedback to more successfully accomplish insertion tasks like those needed for assembly.
- Creating an algorithm to help robots make better trade offs so that they are able to balance optimizing for both safety constraints and task success.

HAI RESEARCH

- Improving upon the previous challenges of video prediction models so that robots can perform a wider variety of manipulation tasks more successfully.

"Recent advancements in machine learning have tremendously expanded the skill set of robots and made them more adaptable to new situations. We can now compress and fuse high-dimensional, raw, and messy sensor data which enables robotic learning from multiple modalities. The next frontier is using AI to develop lifelong learning and continuous adaptation methods so we can create even more intelligent robots."

–Jeannette Bohg, Professor for Robotics and Director of the Interactive Perception and Robot Learning Lab at Stanford University

2 ROBOTICS AND AI
LEARNING FROM HUMAN DEMONSTRATION

WHAT'S NEW?

The goal of autonomous robotics is often to replicate intelligent human behavior. Recent advancements in machine learning have enabled researchers to create better models using huge amounts of messy data. This mirrors real-world conditions, where good data is hard to collect. With these techniques, robots can learn complex tasks from human demonstrations. Additional innovation with human-in-the-loop learning methods means that robots can even learn from human intervention and adjust their behavior for the future, creating robots that can more accurately imitate intelligent human behavior.

WHY DOES THIS MATTER NOW?

With pandemic fueled automation interest, 77% of machinery and automation companies experienced strong increases in demand. Because of social distancing needs, people cannot work as closely together as before and robots help fill the gaps. These robots need to learn how to perform skills previously done by humans and can only do so by learning from human behavior. As we see more robots and humans working together, we need to ensure these robots can learn from real-world data while also remaining flexible enough to adapt to human feedback so they can continue to be successful in the future.

Read more: https://www.mckinsey.com/industries/advanced-electronics/our-insights/beyond-covid-19-rapid-steps-that-can-help-machinery-and-industrial-automation-companies-recover

EYE ON CAMPUS

Researchers on campus are...

HAI RESEARCH

- Training robots to do many different tasks (70+) from messy human demonstration videos expressed in a single multitask learning policy.

- Drawing from "in-the-wild" videos of human behavior to learn generalizable reward functions that enable reinforcement learning for robots across a wide range of tasks.
- Enabling robots to learn complex real world manipulation tasks from a small number of human demonstrations and then further synthesize novel behaviors so they can perform tasks that were not explicitly demonstrated.
- Adding human intervention to human-in-the-loop imitation learning so that robots can improve after failing during complex manipulation tasks.
- Training robots for complex manipulation tasks using a new platform that helps to collect human demonstration data at scale.

"Reinforcement learning (RL) provides a general framework for robots to acquire new skills through learning and interaction. However, learning new tasks in uncertain environments requires extensive exploration, while safety, on the other hand, imposes many limitations on exploration, preventing widespread use of reinforcement learning in real world applications (e.g., obstacle avoidance). New algorithms aim to navigate this tradeoff and introduce safe RL solutions that effectively balance task performance and constraint satisfaction."

–Chelsea Finn, Assistant Professor in Computer Science and Electrical Engineering at Stanford University

3 ROBOTICS AND AI ASSISTIVE ROBOTICS AND HUMAN ROBOT INTERACTION

WHAT'S NEW?

Working to advance AI for object manipulation, task completion, and planning, researchers are developing ways for robots to be helpful in everyday life. Stanford researchers pioneer the necessary skills for robots to be useful in a variety of daily tasks such as dressing, bathing, cooking and cleaning.

Developments in AI learning methods have also unlocked the ability for robots to interact with people and adapt to their preferences, bringing us closer to a world with highly functional assistive robots.

WHY DOES THIS MATTER NOW?

With the Covid-19 pandemic causing many people to spend much more time at home, consumers are needing to rediscover what they need out of their home. Usage of smart home technology is on the rise and consumers across all age groups are interested. Millennials are increasingly investing in this technology while older generations are frequently choosing to stay at home, rather than moving to another facility, increasing the need for assistive home devices. Continued work is necessary to develop assistive robots that can work generally across many different tasks, especially for those involving heavy human-robot interaction.

Read more: https://www-cbinsights-com.stanford.idm.oclc.org/reports/CB-Insights_Future-Smart-Home.pdf

EYE ON CAMPUS

Researchers on campus are...

- Developing robots that can sense human positions to safely help humans in high contact physical interactions such as dressing and bathing.
- Enabling the development of future assistive robots by creating a framework to simulate situations where a robot can assist someone in the activities of daily living (e.g. drinking, dressing, bathing)
- Building a way to help simulate humans more accurately with a physics simulator capable of exploiting data and machine learning in modeling to better simulate rigid bodies.
- Teaching robots how to adapt to human preferences using many diverse modes of feedback, such as physical corrections and imperfect demonstrations.
- Enabling robots to collaborate or compete with both robot and human partners by creating ways to model partner behavior and study how these models change over time. This is particularly suited for discrete decision making and continuous control tasks such as collaborative assembly or moving objects with a partner.

HAI RESEARCH Establishing benchmarks of data and creating a simulation environment to teach robots how to do everyday household tasks like cleaning and cooking.

- Developing methods for robots to make better long-horizon predictions so they can model human behavior, enabling shared decision making
- Defining principles for combining semi-autonomous robotic control and human decision making to make more capable prostheses

"AI-enabled robots have the potential to provide physical assistance that involve applying physical forces to human bodies. This capability will transform healthcare in our aging society and allow older adults to be independent, stay in their homes longer, and have better quality of life. To realize this vision, we must build intelligent, safe, and ethical machines that enhance our sensing and actuating capabilities, but never take away our autonomy to make decisions."

–Karen Liu, Associate Professor in the Computer Science Department at Stanford University

4 ROBOTICS AND AI MEDICAL ROBOTICS AND HUMAN AUGMENTATION

WHAT'S NEW?

Advancements in AI within robotics can greatly improve the quality of medical care available and enhance our human abilities. From creating more intuitive prosthetics to providing more effective tele-operative surgical tools, researchers use robotics to boost the quality of life for patients and strengthen medical practitioners' ability to provide care.

On top of this, robots can play a broader role in human augmentation - amplifying human abilities beyond normal. Using AI-empowered exoskeletons, researchers have discovered ways to improve while also making energy expenditure more efficient. By combining human intelligence and robotic skill, researchers are exploring a new realm of human capability.

WHY DOES THIS MATTER NOW?

By 2019, the medical device industry had already reached a $400B+ market. Even on its own, the global surgical robots market is projected to reach a value of $14.4B by 2026, with a CAGR of 17.6%. AI and robotics represent a great opportunity to improve upon the current state of tools and devices. Furthermore, when robotics are used to augment human ability, we can see even larger effects across many industries, from manufacturing to agriculture.

Read more: https://www.globenewswire.com/news-release/2021/06/30/2255807/0/en/Medical-Devices-Market-Size-to-Hit-Around-US-671-49-Bn-by-2027.html
https://app-cbinsights-com.stanford.idm.oclc.org/research/cmr-surgical-series-d-funding/
https://app-cbinsights-com.stanford.idm.oclc.org/research/robotics-medicine-disruption/

EYE ON CAMPUS

Researchers on campus are...

- Creating an assistive walking device to prevent falls and help recover balance by modeling human gait to predict falling.
- Enabling surgeons to more accurately estimate the force of their movements in tele-operated and robot-assisted surgery
- Developing intuitive and personalizable interfaces for teleoperating assistive robot arms to help users with disabilities with everyday tasks
- Developing state of the art prosthetics such as ankle prosthesis with wrist exoskeleton control
- Upleveling human mobility by using human-in-the-loop methods to create exoskeletons that can increase the speed and reduce the energy expended of people when walking.

HAI RESEARCH

Creating customizable training methods for exoskeletons so that new users can adapt to and benefit from the devices in less time and with less energy expended.

"Advancements in robotics have enabled novel solutions for medical application such as surgery and rehabilitation. In particular, soft robots have emerged as a promising mechanism to access and support the human body due to their continuum geometry and mechanical flexibility. Soft robots are often created with the goal of offering compliant physical interactions between a robot and its environment, which can be safer and more readily accommodate uncertainty compared to traditional rigid robots. The ability of end-users to design and create low-cost soft robots to support human health and wellbeing has the potential to democratize robotic technology."

–Allison Okamura, Professor in the Mechanical Engineering Department at Stanford University

5 ROBOTICS AND AI
AUTONOMOUS VEHICLES AND ROBOT NAVIGATION

WHAT'S NEW?

For autonomous vehicles to work well, they need to be able to adapt to high amounts of change and uncertainty in their environment. With a focus on ensuring safety, researchers are developing methods to improve autonomous vehicle perception, positioning, and navigation to deal with uncertainty even in undesirable situations such as sensor malfunctions and busy environments.

Additionally, many researchers at Stanford are pioneers in the social navigation space — developing methods for robots to better interact with and predict the behavior of other robots or humans so they can safely navigate around one another. Such improvements can help to build trust in the technologies powering AVs and accelerate their creation.

WHY DOES THIS MATTER NOW?

Large companies have long been interested in developing autonomous vehicles, with big players coming from both the tech and automotive industry. On the consumer side, the global advanced driver assistance system market is expected to reach a value of $77.8B by 2027, according to Global Industry Analysts. Within enterprise, autonomous vehicles are posed to have great influence within delivery, transportation, trucking, and the future of logistics. However, current iterations of autonomous vehicles are far from fully autonomous and innovation in autonomous vehicle safety, specifically safe human interaction, will be essential to unlocking these markets.

Read more: https://app-cbinsights-com.stanford.idm.oclc.org/research/momenta-corporate-minority-funding/

EYE ON CAMPUS

Researchers on campus are...

- Developing a better way for AVs to track multiple objects in their environment by augmenting the traditional detection-based approaches with learning based methods.
- Developing methods for robots to navigate with and around humans in a trustworthy way by communicating intent, avoiding collision, and planning efficient routes.
- Creating safer autonomous vehicle navigation and more trustworthy methods for localization by training deep learning networks to learn sensing uncertainties and integrating deep learning with model-driven approaches.
- Guaranteeing safe autonomous route-planning and navigation even with sensor errors and environmental uncertainty.
- Identifying the most likely failure events for autonomous vehicles using probabilistic models and adaptive stress testing so that we can improve AV safety.
- Creating methods to better analyze and certify neural network collision avoidance systems for aircraft.
- Combining imitation and reinforcement learning techniques so that robots can navigate with and around humans for social navigation.

"Autonomous vehicles and robots have the potential to greatly enhance productivity, but there are major technical challenges in ensuring that the resulting systems are trustworthy. We have seen some major breakthroughs over the past few years in developing systems that are robust to various sources of uncertainty and validating their safety."

– Mykel Kochenderfer, Associate Professor of Aeronautics and Astronautics and Associate Professor, by courtesy, of Computer Science at Stanford University

Corporate Engagement

Stanford HAI seeks corporate members who will enable it to lead with the unparalleled interdisciplinary breadth and excellence of Stanford University, and a laser focus on human-centered development and deployment of AI technology. We invite engagement from companies that share our mission to advance AI research, education, policy, and practice to improve the human condition.

CORPORATE ENGAGEMENT

Are you prepared for the next wave of change?

THINGS YOU SHOULD KNOW.

88%

Respondents who planned to increase their investment in industrial robotics in 2018, across automotive, electronics, and pharmaceutical industries.[1, 2]

30%

productivity increase from robotics in some countries.[2]

$510B

possible global robotics market expansion by 2026.[2]

References:
1 https://www.mckinsey.com/~/media/mckinsey/industries/advanced%20electronics/our%20insights/growth%20dynamics%20in%20industrial%20robotics/industrial-robotics-insights-into-the-sectors-future-growth-dynamics.ashx
2 https://partners.wsj.com/global-x-etfs/charting-disruption/?utm_source=LinkedIn&utm_medium=social&utm_content=charting-disruption-health-care%C2%A7ion%3Dautomation-meets-innovation&section=automation-meets-innovation

Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

CORPORATE ENGAGEMENT

Become a corporate member today.

50% of companies are likely to miss the window of opportunity.[1] Let's talk.

Learn more about the Corporate Members Program and the Stanford advantage.

Panos Madamopoulos,
Managing Director for Industry Programs and Partnerships

1 Jacques Bughin, "Wait-and-See Could Be a Costly AI Strategy," *MIT Sloan Management Review*, June 15, 2018.

Industry Take

Industry plays a critical role in creating and scaling novel applications of AI in robotics. It is, therefore, a goal and privilege of HAI to convene stakeholders from industry in addition to those in academia, government, and civil society to address the technical and societal challenges posed by AI in various robotics applications.

Leading venture investors, positioned at the frontlines of startup innovation in robotics, can provide a unique perspective on the impact and role of AI technologies in robotics design and application problems.

Industry Take

"Humans are wonderfully flexible creatures — in any given job, even the most 'repetitive,' we're doing many things. A line cook in a restaurant doesn't just make food, they prep ingredients, sharpen tools, clean their stations. Automation is great for doing one thing a million times, but a robot trying to do a human job has to be able to do a thousand things a thousand times each. Thanks to advances in software and AI, robots are now able to approach that flexibility, and that means they can address trillion-dollar labor markets instead of just billion-dollar automation hardware markets."
–Haomiao Huang, Investor, Kleiner Perkins

"Robotics is the most complex application area within AI, for the simple reason that it involves navigating the real world. Dealing with atoms, not just bits, makes the problem space vastly more challenging. Today's AI can defeat the world's best human player at Go and accurately predict the three-dimensional structure of a protein, but good luck getting it to manipulate physical objects with the dexterity of a human toddler. But cutting-edge AI research, in particular in reinforcement learning, is rapidly pushing forward the state of the art in robotics. As robots become capable of increasingly flexible and general-purpose behavior, they will become increasingly interwoven into the fabric of our daily lives, from our roads to our hospitals to our schools to our homes. The economic transformations and opportunities will be tremendous."
–Rob Toews, Partner at Radical Ventures

"Deep learning and neural networks have been around for over 30 years, but we have only begun to scratch the surface of what is achievable when it comes to AI applications in robotics. Breakthroughs over the last decade, brought on by the explosion of data sets and advancements in cloud computing, have shifted what is possible. We continue to see massive opportunities as this research graduates from academic to commercial use, paving the way for generational businesses to emerge in the field."
–Mike Volpi, Partner at Index Ventures

"Artificial intelligence systems leveraging breakthroughs in machine learning and robotics are already shaping our present and future! Our investing strategy is focused on companies that innovate with ML because we believe that automation technology will transform the way we work, whether it be in the office, in factories, or on farms. With applications such as workflow automation and robotics, artificial intelligence is critical to improving productivity in all parts of the economy."
–Lan Xuezhao, Founder & Managing Partner at Basis Set

Stanford University
Human-Centered
Artificial Intelligence

INDUSTRY BRIEF

Appendix

Artificial Intelligence Definitions

Intelligence might be defined as the ability to learn and perform a range of techniques to solve problems and achieve goals—techniques that are appropriate to the context in an uncertain, ever-varying world. A fully pre-programmed factory robot is flexible, accurate, and consistent, but not intelligent.

Artificial Intelligence (AI), is a term coined in 1955 by John McCarthy, Stanford's first faculty member in AI, who defined it as "the science and engineering of making intelligent machines." Much research has human program software agents with the knowledge to behave in a particular way, like playing chess, but today, we emphasize agents that can learn, just as human beings navigating our changing world.

Autonomous systems can independently plan and decide sequences of steps to achieve a specified goal without being micromanaged. A hospital delivery robot must autonomously navigate busy corridors to succeed in its task. In AI, autonomy doesn't have the sense of being self-governing common in politics or biology.

Machine Learning (ML) is the part of AI that studies how computer systems can improve their perception, knowledge, decisions, or actions based on experience or data. For this, ML draws from computer science, statistics, psychology, neuroscience, economics, and control theory.

In **supervised learning**, a computer learns to predict human-given labels, such as particular dog breeds based on labeled dog pictures. **Unsupervised learning** does not require labels, but sometimes adopts **selfsupervised learning**, constructing its own prediction tasks such as trying to predict each successive word in a sentence. **Reinforcement learning** enables autonomy by allowing an agent to learn action sequences that optimize its total rewards, such as winning games, without explicit examples of good techniques.

Deep Learning is the use of large multi-layer **(artificial) neural networks** that compute with continuous (real number) representations, similar to the hierarchically organized neurons in human brains. It is successfully employed for all types of ML, with better generalization from small data and better scaling to big data and compute budgets. A recent breakthrough is the **transformer**, a neural net architecture which flexibly incorporates context via an attention mechanism, allowing powerful and computationally efficient analysis and generation of sequences, such as words in a paragraph.

Foundation models are an emerging class of models, often transformers trained by self-supervision on large-scale broad data, that can be easily adapted to perform a wide range of downstream tasks. The best-known examples are large pretrained language models like GPT-3, but the term extends to models for all modalities of data and knowledge.

An **algorithm** is a precise list of steps to take, such as a computer program. AI systems contain algorithms, but typically just for a few parts like a learning or reward calculation method. Much of their behavior emerges via learning from data or experience, a fundamental shift in system design that Stanford alumnus Andrej Karpathy dubbed **Software 2.0**.

Narrow AI is intelligent systems for particular tasks, e.g., **speech** or **facial recognition**. **Human-level AI**, or **artificial general intelligence (AGI)**, seeks broadly intelligent, context-aware machines. It is needed for effective, adaptable **social chatbots** or **human-robot interaction**.

Human-Centered Artificial Intelligence is AI that seeks to augment the abilities of, address the societal needs of, and draw inspiration from human beings. It researches and builds effective partners and tools for people, such as a robot helper and companion for the elderly.

Text by Professor Christopher Manning, v 1.2, April 2022

References

ROBOTICS AND AI

Adaptive Learning & Optimization

Corkery, M., & Gelles, D. (2020, April 10). Robots Welcome to Take Over, as Pandemic Accelerates Automation. New York Times. https://www.nytimes.com/2020/04/10/business/coronavirus-workplace-automation.html

Inagaki, K. (2020, May 12). Coronavirus Will Usher in the Industrial Robots. Financial Times. https://www.ft.com/content/fa81806d-8ddf-4f2a-9a5a-de6c96844373

Bolz, L., Ludwig-Dehm, P. M., Teulieres, M., Tilley, J., & Wägner, S. (2019, September 13). Growth Dynamics in Industrial Robotics. McKinsey & Company https://www.mckinsey.com/industries/advanced-electronics/our-insights/growth-dynamics-in-industrial-robotics

Brown, K., Peltzer, O., Sehr, M. A., Schwager, M., & Kochenderfer, M. J. (2020, May). *Optimal sequential task assignment and path finding for multi-agent robotic assembly planning.* In *2020 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 441-447). IEEE. https://arxiv.org/abs/2006.08845

Migimatsu, T., & Bohg, J. (2020). *Object-Centric Task and Motion Planning in Dynamic Environments. IEEE Robotics and Automation Letters*, 5, 844-851. https://arxiv.org/abs/1911.04679

Zachares, P. A., Lee, M. A., Lian, W., & Bohg, J. (2021, May). Interpreting contact interactions to overcome failure in robot assembly tasks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3410-3417). IEEE. https://arxiv.org/abs/2101.02725

Lee, M. A., Zhu, Y., Srinivasan, K., Shah, P., Savarese, S., Fei-Fei, L., ... & Bohg, J. (2019, May). Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. In *2019 International Conference on Robotics and Automation (ICRA)* (pp. 8943-8950). IEEE. https://arxiv.org/abs/1810.10191

Thananjeyan, B., Balakrishna, A., Nair, S., Luo, M., Srinivasan, K., Hwang, M., ... & Goldberg, K. (2021). Recovery rl: Safe reinforcement learning with learned recovery zones. IEEE Robotics and Automation Letters, 6(3), 4915-4922. https://arxiv.org/abs/2010.15920

HAI-Funded: Wu, B., Nair, S., Martin-Martin, R., Fei-Fei, L., & Finn, C. (2021). Greedy hierarchical variational autoencoders for large-scale video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 2318-2328). https://arxiv.org/abs/2103.04174

Learning from Human Demonstration

Becker, M., Distel, D., Freund, H., Herring, D. (2020, May 26). Beyond Covid-19: Rapid Steps That Can Help Machinery and Industrial Automation Companies Recover. McKinsey & Company. https://www.mckinsey.com/industries/advanced-electronics/our-insights/beyond-covid-19-rapid-steps-that-can-help-machinery-and-industrial-automation-companies-recover

HAI-Funded: Shao, L., Migimatsu, T., Zhang, Q., Yang, K., & Bohg, J. (2021). Concept2robot: Learning manipulation concepts from instructions and human demonstrations. *The International Journal of Robotics Research*, 40(12-14), 1419-1434. http://www.roboticsproceedings.org/rss16/p082.pdf

Chen, A. S., Nair, S., & Finn, C. (2021). Learning generalizable robotic reward functions from" in-the-wild" human videos. arXiv preprint arXiv:2103.16817. https://arxiv.org/abs/2103.16817

Mandlekar, A., Xu, D., Martín-Martín, R., Savarese, S., & Fei-Fei, L. (2020). Learning to generalize across long-horizon tasks from human demonstrations. arXiv preprint arXiv:2003.06085. https://arxiv.org/abs/2003.06085

Mandlekar, A., Xu, D., Martín-Martín, R., Zhu, Y., Fei-Fei, L., & Savarese, S. (2020). Human-in-the-loop imitation learning using remote teleoperation. arXiv preprint arXiv:2012.06733. https://arxiv.org/abs/2012.06733

Mandlekar, A., Zhu, Y., Garg, A., Booher, J., Spero, M., Tung, A., ... & Fei-Fei, L."Robotic Skill Learning with Large Human Datasets." *RoboTurk*, https://roboturk.stanford.edu/

Assistive Robotics and Human Robot Interaction

CB Insights. (2020, December 17). The Future Of The Smart Home. Retrieved from CB Insights database.

Erickson, Z., Clever, H. M., Gangaram, V., Turk, G., Liu, C. K., & Kemp, C. C. (2019, June). Multidimensional capacitive sensing for robot-assisted dressing and bathing. In *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)* (pp. 224-231). IEEE. https://ckllab.stanford.edu/multidimensional-capacitive-sensing-robot-assisted-dressing-and-bathing

Erickson, Z., Gangaram, V., Kapusta, A., Liu, C. K., & Kemp, C. C. (2020, May). Assistive gym: A physics simulation framework for assistive robotics. In *2020 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 10169-10176). IEEE. https://ieeexplore.ieee.org/abstract/document/9197411

Werling, K., Omens, D., Lee, J., Exarchos, I., & Liu, C. K. (2021). Fast and feature-complete differentiable physics for articulated rigid bodies with contact. arXiv preprint arXiv:2103.16021. https://arxiv.org/abs/2103.16021

References

ROBOTICS AND AI

Bıyık, E., Losey, D. P., Palan, M., Landolfi, N. C., Shevchuk, G., & Sadigh, D. (2021). Learning reward functions from diverse sources of human feedback: Optimally integrating demonstrations and preferences. *The International Journal of Robotics Research*, 02783649211041652. https://arxiv.org/abs/2006.14091

Li, M., Canberk, A., Losey, D. P., & Sadigh, D. (2021, May). Learning human objectives from sequences of physical corrections. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2877-2883). IEEE. https://arxiv.org/abs/2104.00078

Cao, Z., & Sadigh, D. (2021). Learning from imperfect demonstrations from agents with varying dynamics. *IEEE Robotics and Automation Letters*, 6(3), 5231-5238. https://ieeexplore.ieee.org/abstract/document/9387082

Xie, A., Losey, D. P., Tolsma, R., Finn, C., & Sadigh, D. (2020). Learning latent representations to influence multi-agent interaction. arXiv preprint arXiv:2011.06619. https://iliad.stanford.edu/pdfs/publications/xie2020learning.pdf

HAI-Funded: Srivastava, S., Li, C., Lingelbach, M., Martín-Martín, R., Xia, F., Vainio, K. E., ... & Fei-Fei, L. (2022, January). BEHAVIOR: Benchmark for Everyday Household Activities in Virtual, Interactive, and Ecological Environments. In *Conference on Robot Learning* (pp. 477-490). PMLR. https://arxiv.org/abs/2108.03332

Li, C., Xia, F., Martín-Martín, R., Lingelbach, M., Srivastava, S., Shen, B., ... & Savarese, S. (2021). Igibson 2.0: Object-centric simulation for robot learning of everyday household tasks. arXiv preprint arXiv:2108.03272. https://arxiv.org/abs/2108.03272

Li, A. H., Wu, P., & Kennedy, M. (2021, May). Replay Overshooting: Learning stochastic latent dynamics with the extended kalman filter. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 852-858). IEEE. https://arm.stanford.edu/publications/replay-overshooting-learning-stochastic-latent-dynamics-extended-kalman-filter

Guptasarma, S., & Kennedy III, M. Considerations for the Control Design of Augmentative Robots. https://arm.stanford.edu/publications/considerations-control-design-augmentative-robots

Medical Robotics and Human Augmentation

Precedence Research. (2021, June 30). Medical Devices Market Size to Hit around US$ 671.49 Bn by 2027. GlobeNewswire News Room, Precedence Research. https://www.globenewswire.com/news-release/2021/06/30/2255807/0/en/Medical-Devices-Market-Size-to-Hit-Around-US-671-49-Bn-by-2027.html

CB Insights. (2021, June 30). This Softbank-Backed Company Raised $600M To Scale Its Surgery Robot. Retrieved from CB Insights database.

CB Insights. (2019, July 24). Paging Dr. Robot: How Robotics Is Changing The Face Of Medicine. Retrieved from CB Insights database.

Kumar, V. C., Ha, S., Sawicki, G., & Liu, C. K. (2020, May). Learning a control policy for fall prevention on an assistive walking device. In 2020 *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4833-4840). IEEE. https://arxiv.org/abs/1909.10488

Chua, Z., Jarc, A. M., & Okamura, A. M. (2021, May). Toward force estimation in robot-assisted surgery using deep learning with vision and robot state. In 2021 *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 12335-12341). IEEE. https://arxiv.org/abs/2011.02112

Losey, D. P., Jeon, H. J., Li, M., Srinivasan, K., Mandlekar, A., Garg, A., ... & Sadigh, D. (2022). Learning latent actions to control assistive robots. *Autonomous robots*, 46(1), 115-147. https://iliad.stanford.edu/pdfs/publications/losey2021learning.pdf

Welker, C. G., Chiu, V. L., Voloshina, A. S., Collins, S. H., & Okamura, A. M. (2020). Teleoperation of an ankle-foot prosthesis with a wrist exoskeleton. *IEEE Transactions on Biomedical Engineering*, 68(5), 1714-1725. https://ieeexplore.ieee.org/abstract/document/9301430

Song, S., & Collins, S. H. (2021). Optimizing exoskeleton assistance for faster self-selected walking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29, 786-795. https://ieeexplore.ieee.org/abstract/document/9408629

Franks, P. W., Bryan, G. M., Martin, R. M., Reyes, R., Lakmazaheri, A. C., & Collins, S. H. (2021). Comparing optimized exoskeleton assistance of the hip, knee, and ankle in single and multi-joint configurations. *Wearable Technologies*, 2. https://www.biorxiv.org/content/10.1101/2021.02.19.431882v1.abstract

HAI-Funded: Poggensee, K. L., & Collins, S. H. (2021). How adaptation, training, and customization contribute to benefits from exoskeleton assistance. *Science Robotics*, 6(58), eabf1078. https://www.biorxiv.org/content/10.1101/2021.04.25.440289v1.abstract

References

ROBOTICS AND AI

Autonomous Vehicles and Robot Navigation

CB Insights. (2021, September 27). GM Invested $300M In Momenta, A Developer Of Autonomous Vehicle Software Used By Automakers, To Enhance Vehicles In China. Retrieved from CB Insights database.

Chiu, H. K., Li, J., Ambruş, R., & Bohg, J. (2021, May). Probabilistic 3D multi-modal, multi-object tracking for autonomous driving. In 2021 *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 14227-14233). IEEE. https://arxiv.org/abs/2012.13755

Che, Y., Okamura, A. M., & Sadigh, D. (2020). Efficient and trustworthy social navigation via explicit and implicit robot–human communication. *IEEE Transactions on Robotics*, 36(3), 692-707. https://arxiv.org/abs/1810.11556

Dai, A., Gupta, S., Kanhere, A., Mohanty, A., & Shetty, A. (2021) Trustworthy Vehicle Localization with Deep Learning. Stanford Engineering NAV Lab. https://navlab.stanford.edu/research/trustworthy-deep-learning

Bhamidipati, S., Dai, A., Mina, T., Shetty, A., Tzikas, A. (2021) Verifiable Perception and Navigation. Stanford Engineering NAV Lab. https://navlab.stanford.edu/research/verifiable-perception-and-navigation

Lee, R., Mengshoel, O. J., Saksena, A., Gardner, R. W., Genin, D., Silbermann, J., ... & Kochenderfer, M. J. (2020). Adaptive stress testing: Finding likely failure events with reinforcement learning. *Journal of Artificial Intelligence Research*, 69, 1165-1201. https://www.jair.org/index.php/jair/article/view/12190

Julian, K. D., & Kochenderfer, M. J. (2021). Reachability Analysis for Neural Network Aircraft Collision Avoidance Systems. *Journal of Guidance, Control, and Dynamics*, 44(6), 1132-1142. https://arc.aiaa.org/doi/abs/10.2514/1.G005233

Pokle, A., Martín-Martín, R., Goebel, P., Chow, V., Ewald, H. M., Yang, J., ... & Vázquez, M. (2019, May). Deep local trajectory replanning and control for robot navigation. In *2019 International Conference on Robotics and Automation (ICRA)* (pp. 5815-5822). IEEE. https://arxiv.org/abs/1905.05279

Pérez-D'Arpino, C., Liu, C., Goebel, P., Martín-Martín, R., & Savarese, S. (2021, May). Robot navigation in constrained pedestrian environments using reinforcement learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1140-1146). IEEE. https://arxiv.org/abs/2010.08600

Corporate Engagement

Bolz, L., Ludwig-Dehm, P. M., Teulieres, M., Tilley, J., & Wägner, S. (2019, September 13). Growth Dynamics in Industrial Robotics. McKinsey & Company https://www.mckinsey.com/industries/advanced-electronics/our-insights/growth-dynamics-in-industrial-robotics

(2021, December 13). Charting Disruption. The Wall Street Journal, Dow Jones & Company. https://partners.wsj.com/global-x-etfs/charting-disruption/?section=intro

Acknowledgments

This brief was produced by the HAI Industry Programs & Partnerships team as one of its strategic initiatives. Research and writing were led by Jessica Hong in coordination with Ahmad Rushdi, Marc Gough, Alice Hau, Victoria Moore, and Panos Madamopoulos. We are grateful for the generosity of the following people for providing their time, helpful suggestions, and constructive feedback in the creation of this brief (names listed in no particular order): Erik Brynjolfsson, Jeannette Bohg, Chelsea Finn, Dorsa Sadigh, Karen Liu, Allison Okamura, and Mykel Kochenderfer. We would also like to thank the HAI Staff team including Vanessa Parli, Jeanina Casusi, Joe Hinman, Shana Lynch, Stacy Peña, and Michi Turner for their help preparing the publication.

Stanford University
Human-Centered
Artificial Intelligence