RE•WORK

AI FOR SOCIAL GOOD

RE•WORK

Yazmin How, Digital Content Creator

ABOUT THE PAPER

With the rapid advancements and applications of AI, conversations and concerns have increased around the intentions of this technology. There are concerns that AI could be used with malicious intent rather than for the benefit of human-kind. This paper explores areas where artificial intelligence can benefit society and tackle global challenges such as the environment, education, healthcare and sustainability.

This June, RE•WORK will be hosting the AI for Good Summit in San Francisco. Visit the website for more information here.

ABOUT RE•WORK

As well as creating digital content, RE•WORK is first and foremost the global leader in AI and Deep Learning events. RE•WORK create and organize globally renowned summits, workshops and dinners, bringing together the brightest minds in AI from both industry and academia. At each RE•WORK event, we combine the latest technological innovation with real-world applications and practical case studies. You can learn from global pioneers and industry experts, and network with CEOs, CTOs, data scientists, engineers and researchers disrupting their industries with AI. We also provide an analysis of current trends and innovations, through podcasts, white papers and video interviews. Additionally, we have an extensive on-demand video library of presentations from world-leading experts in AI. We cover topics such as Deep Learning, Machine Learning, Robotics, AI in Healthcare, Women in Tech, AI in Finance, Reinforcement Learning, Computer Vision, Autonomous Vehicles, AI Assistants, AI for Good, Industrial Automation and more.

CONTENTS

1. CONTRIBUTORS

BABAR BHATTI
CEO & Co-Founder, MutualMind

Babar works on leveraging advanced analytics and proprietary data to create AI-powered solutions for our customers. He is also an advisor on AI strategy, technology, policy and fair and ethical use of AI, AI education for C-level executives and board members.

DIMITRI KANEVSKY
Research Scientist, Google

Dimitri started his career at Google working on speech recognition algorithms for a YouTube speech recognition system. MIT Technology Review, recognized Dimitri's conversational biometrics based security patent as one of five most influential patents for 2003. In 2012 Dimitri was honored at the White House as a Champion of Change for his efforts to advance access to science, technology, engineering, and math.

MARK CROWLEY
Assistant Professor, University of Waterloo

Mark is an Assistant Professor in the Pattern Recognition and Machine Intelligence group in the Department of Electrical and Computer Engineering. Mark's research focuses on algorithms, tools and theory at the intersection of Machine Learning, Optimization and Probabilistic Modeling.

DEVIN KROTMAN
Manager for Global Learning, XPRIZE

Devin serves as the Director of both the IBM Watson AI XPRIZE and Global Learning XPRIZE. He is passionate about tackling the world's grandest challenges from education to disaster prediction and firmly believes technology will help humanity get there.

DANIEL GOLDEN
Director of ML, Arterys

Daniel is focusing on deep learning and convolutional neural network-based segmentation, classification, and decision support for MRI, CT, and other imaging modalities. Arterys was the first ever FDA cleared, cloud-based clinical product with deep learning functionality.

KARRY LU
Sr. Data Scientist, WeWork

Graduate of the University of Michigan, Karry specializes in problem solving and building data products. He is currently working on machine learning powered recommendation systems to improve the user experience and drive engagement by delivering personalized content, such as music, retail products, and social content at WeWork.

DAVID J. GUNKEL
Presidential Teaching Professor, Northern Illinois University

David is an award-winning educator and scholar, specializing in the ethics of new and emerging tech. He has lectured and delivered award-winning papers throughout North and South America and Europe and is the founding co-editor of the International Journal of Žižek Studies and the Indiana University Press book Digital Game Studies.

MICHAEL MARTIN
Head of Communications, SignalFire

Michael works across diverse disciplines, specializing in community building, AI, and public policy. A city planner by trade, he currently is the Head of Communities at SignalFire, an AI-driven, early-stage venture capital fund. In his previous work, he managed the IBM Watson AI XPRIZE and led Global Community Relations at the XPRIZE Foundation.

MAREK HAVRDA
Strategy Advisor, GoodAI

Marek is a social entrepreneur, social scientist and consultant with expertise in applications of current technologies for data analysis. Interested in any data-driven (behavioral) research which delivers practical results. His specific interests lie in artificial intelligence, education, health, retail, non-profits, information technology, telecommunications and non-ferrous metals.

FIONA MCEVOY
Founder & CEO, YouTheData.com

Fiona's thesis paper, "Decisions, decisions: Big Data and the Future of Autonomy" has been published by the Annals of the University of Bucharest as part of their series on emerging technology, and presented at Stanford University to the International Association of Computing and Philosophy. Fiona founded YouTheData.com as a platform for the discussion of technology, AI and the use of data in our society.

HASSAN MURAD
CEO & Co-Founder, Intuitive AI

Hassan is working on developing AI to empower a zero waste world. He was driven by the world's massive waste problem to begin creating solution sand divert waste away from landfills, oceans or incineration. His Robotics background, that includes working on autonomous submarines, farm- mapping drones, and self-driving cars at Tesla, has inspired him to develop Oscar, AI for Zero Waste.

TESS POSNER
CEO, AI4ALL

Tess is working to make AI more diverse and inclusive and to ensure that AI is developed responsibly. Earlier in her career, Tess built and ran Samaschool, a nonprofit that equips low-income people to find work in the digital economy through an online platform training 50,000+ students worldwide and a dozen diverse locations from New York City to rural Arkansas and East Africa.

DEVAL PANDYA
Data Scientist, Shell

Deval is one of the 100 global Future Energy Leaders at World Energy Council and Data Scientist at Data Science CoE in Shell. His work focuses on various predictive analytics problems in the energy industry including seismic processing, predictive maintenance, GHG accounting, soft sensors, energy platform and nature based carbon mitigation mechanism.

MAITHRA RAGHU
Research Scientist, Google Brain

Maithra's research interests are in developing principled tools to empirically study the representational properties of deep neural networks, and apply these insights to deep learning applications in healthcare. She was named one of the Forbes 30 Under 30 in Science and a Rising Star in EECS.

CATHY PEARL
Head of Conversation Design Outreach, Google

Cathy has been designing and creating Voice User Interfaces (VUIs) for nearly 20 years and is passionate about helping everyone make the best conversational experiences possible. Previously, Cathy was VP of User Experience at Sensely, whose virtual nurse avatar helps people stay healthy.

MARK WEBER
Research Staff Member, MIT-IBM Watson AI Lab

Marks' expertise is connecting dots across disciplines to bridge academic research with compelling real-world applications. Working for IBM Research, Mark is focused on developing capabilities in emergent technologies such as Artificial Intelligence and blockchain technology, focused on improving our financial and economic systems for integral human flourishing.

2. MAKING THE MOST OF AI

Artificial Intelligence (AI) is transforming our lives, and will redesign life as we know it for future generations. Although AI is not a silver bullet for overcoming global challenges, there is huge potential to conquer some of the biggest obstacles and offer meaningful solutions by applying AI. The scope for social impact is vast, and we are already seeing countless industries implement machine learning technologies, with some of the most notable being healthcare, education, environment and sustainability, energy, and transport amongst others. Applying AI is one thing, but ensuring it's engineered with a positive social outcome for all must be at the core of every design and engineering step, else the potential for adverse ramifications is of concern. Issues that have previously challenged experts, such as curing cancer or predicting natural disasters, are being addressed by teams of pioneers across the globe who are harnessing AI in an attempt to make some of the most severe obstacles concerns of the past.

The United Nations Sustainable Development Goals (SDGs) outlines 17 global goals to achieve 3 major milestones to transform the world by 2030. The overarching goals are to End Extreme Poverty, Fight Inequality and Injustice, and to Tackle Climate Change. The global goals are: No Poverty, Zero Hunger, Good Health and Well-being, Quality Education, Gender Equality, Clean Water and Sanitation, Affordable and Clean Energy, Decent Work and Economic Growth, Industry, Innovation and Infrastructure, Reduced Inequality, Sustainable Cities and Communities, Responsible Consumption and Production, Climate Action, Life Below Water, Life on Land, Peace and Justice Strong Institutions, Partnerships to Achieve the Goal. (Global Goals, 2018)

From saving lives to helping with personal banking; from raising crop productivity to predicting the location of earthquakes; AI is already playing a key role in realizing these goals. For example, DeepMind Health are striving to 'make a practical difference to patients, nurses and doctors', whilst Microsoft Translator is used in schools across the globe to 'create live subtitles during lessons' to promote inclusivity, and the World Economic Forum is supporting the growth of 'low-carbon, green electricity' through AI with the potential of a huge long-term impact. These are just a few of the areas that are already undergoing immense overhauls from the positive effects of AI.

'As AI becomes more integrated into decision making and our daily lives, it's absolutely critical that we have diversity in who is building and shaping this technology from the beginning. Diverse perspectives yield more innovative, human-centered, and ethical products. And if this powerful tool is only in the hands of a few, it could exacerbate inequities and inequalities that we are trying to fix in society today.' (Tess Posner, AI4ALL)
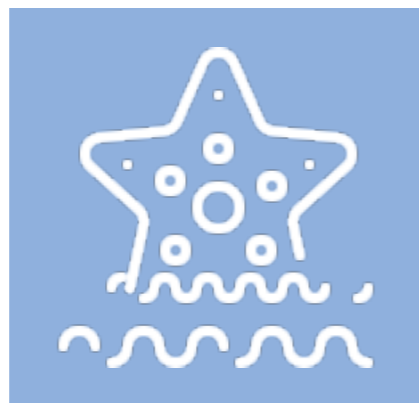
WHAT DO WE MEAN BY AI FOR SOCIAL GOOD?

AI for Social Good is a broad concept but one that generally, in my opinion, refers to the work that is typically undertaken by government and civil society. Applying Artificial Intelligence research for global sustainable development.
Michael Martin, SignalFire

AI for Good means deploying machine learning and data science tools in industries that have a social mission that goes beyond maximizing shareholder value (education, environment, gender equality, etc).
Karry Lu, WeWork

AI for Social Good means recognizing that the technology of artificial intelligence does not exist in a vacuum but is always and already situated in a specific social context and therefore needs to be developed and deployed in a manner that is sensitive to these opportunities and challenges.
David J. Gunkel, Northern Illinois University

There are two aspects to ensuring AI is leveraged for positive impacts. 'One is **policy**, we have seen European nations start to take lead in putting the right measures in place starting with data privacy. This will have to continue throughout. We need to **educate** and **engage** our policy makers. I think that AI like Transformers or Terminators taking over humanity, it is farfetched at this point, but as we start relying more and more on machine learning and deep learning algorithms, we will need to understand **biases** and **social implications** to formulate the right policies. Second, is opening up some of these black back methods. **Interpretability** of models is key. **Artificial Intelligence is not inherently good or bad. It's all about how it's employed.** (Deval Pandya, Shell)

Like any technological innovation, the successful integration of AI into human society is going to require a productive dialogue between research, industry, and public policy. This dialogue, if it is conducted properly, will be and cannot help but be contentious and, at times, polarizing. But this is a good thing. Consider earlier forms of technology innovation, like electrification as it transpired in the US. It was only by way of a rather contentious debate between the industry and state and federal governments that the power grid was developed in such a way as to ensure equal distribution of access in both urban and rural locations. It was not an easy process. But the conflict between competing interests was crucial to formulating a solution that worked. Making this process work for 21st century technology will involve two important components. On the one hand, we need to learn how to value conflict between industry and government as part of a productive process of problem solving and not an impediment to innovation. Cooperation is the outcome of the process, it is not where things begin. It is concern that in many parts of the world, there are all-too-cozy relationships between governments, the supposed representatives of the people, and the industries they are supposed to be regulating. On the other hand, public policy officials and legislators need to start doing their homework. If you look at the questions members of the US Congress put to representatives from Facebook, Twitter, and Google during the recent hearing in Washington, it is clear that many legislators have little or no understanding of the technology. It is difficult to make intelligent laws and proactive public policy decisions, when you do not understand the technology you are trying to regulate. (David J. Gunkel, Northern Illinois University)

Even with policies in place and the most experienced, well-intentioned teams of scientists working on models, human biases are still likely to creep in. Everyone, including developers, are subject to bias, which means that it's essential for diverse teams to be involved in all stages of the process: 'Diversity is really important for identifying and having new perspectives on problems'. (Maithra Raghu, Google Brain)

Just like in the early days of the internet, global protocols need to be in place to maximize the potential of the work researchers are doing. With this in mind, the trend will continue at rapid pace as many of the AI solutions grow and move beyond the traditional tech industry.

TRANSPARENCY & RESPONSIBILITY

AI helps doctors prescribe drugs - who is responsible? A car drives you to work autonomously - who is responsible? A chatbot suggests how you should handle your finances - who is responsible? Most importantly, when any of the above go wrong, who is responsible? This is an area yet to be fully realized, but one where much thought must be given.

'National policy making for new and emerging technologies is never easy. There is the uncertainty inherent in new technologies, all the usual challenges of policy making ranging from politics to personalities and use of resources.' (Babar Bhatti, MutualMind)

Areas of consideration for policy and responsibility:

Education, R&D
- Basic and Applied Research
- Entrepreneurship, Innovation
- Awareness, Training and Talent Development

Vision & Strategy
- Key Priorities
- Guiding Policies
- Coordinated and Compliant Actions

Society & Ethics
- Dangers and Limitations
- Ethics and Fairness, Inequality
- Privacy

AI Policy Facets

Economy, Commerce, Jobs
- Commercial Applications
- Job Creation, Training, Re-skilling

Digital Infrastructure
- Data Governance and Policies
- Cloud Computing
- Machine Learning, Robotics, Automation

Law
- Legislation
- Regulation

Security & Defense
- Privacy & Security Guidelines

Governance & Execution
- Investments, Resource Allocation
- Alliances and Partnerships
- Risk Management
- Program Management

(Babar Bhatti, MutualMind)

AI can easily be used for more nefarious purposes; everything from biased hiring algorithms to spreading misinformation on social media to using computer vision for targeting drone strikes. For now, AI remains a tool in the hands of human beings, the same as sledgehammers or acetylene torches. It's only an extension of our desires and prejudices.
Karry Lu, WeWork

Simply put - responsibility falls on all of us - from Government to commercial/enterprise to non-profit/international development organizations. All of these organizations play a role with helping to keep the public in the loop when it comes to AI and how it can benefit/augment as a society, not supplant us.
Devin Krotman, XPRIZE

There is a risk that a 'common standard' - however it might look - could make ethics a "tick box" exercise for technology companies that really need to engage beyond a set of static rules. Nevertheless, there is clearly a real need for support and guidance too. Groups like the IEEE are driving some great work to lay the foundations across a spectrum of new technologies, and obviously new regulations like the EU GDPR are game-changing in that regard.
Fiona McEvoy, YouTheData.com

3. GLOBAL INITIATIVES

Outside of governments and institutions, companies are forming or creating departments specifically geared towards ensuring current and future generations are sufficiently equipped to tackle the challenges that are arising surrounding AI. These initiatives range from encouraging diverse backgrounds into AI, to solving challenges with the environment, to eliminating diseases and improving patient care with healthcare discoveries.

These are some of the companies with initiatives in place, striving to better the world through the research and applications of AI:

AI4ALL

"AI Will Change the World. Who Will Change AI? AI4ALL is Educating the Next Generation of AI Technologists, Thinkers, and Leaders."
The nonprofit works to increase both diversity and inclusion in AI and create pipelines for underrepresented talent through education and mentorship programs. Their goal is to develop a broad group of individuals advancing AI for humanity's benefit.

GOOGLE AI

"Applying AI to some of the world's biggest challenges. Through research, engineering and initiatives to build the AI ecosystem, we're working to use AI to address societal challenges."
Google runs the AI impact challenge which encourages organizations from across the globe to present their ideas for how they could use AI to help address societal challenges.

XPRIZE FOUNDATION

"Let's change the world - TOGETHER. We make the impossible possible by incentivizing great minds to make a difference."
XPRIZE are dedicated to inspiring and incentivizing breakthroughs to the most critical challenges facing humanity. They address these challenges that cannot be solved by traditional means through incentivized prize competitions.

MICROSOFT AI

"AI for Earth: Helping the planet with data science."
The 5 year $50 million Microsoft initiative supports and partners with environmental groups and researchers to tackle some of the world's pressing challenges through the application of AI, ML and the cloud.

THE IEEE GLOBAL INITIATIVE ON ETHICS OF AUTONOMOUS AND INTELLIGENT SYSTEMS

"An incubation space for new standards and solutions, certifications and codes of conduct, and consensus building for ethical implementation of intelligent technologies."

THE AI INITIATIVE

Created in 2015, the AI Initiative gathers students, researchers, alumni, faculty and experts from Harvard and beyond, interested in understanding the consequences of the rise of Artificial Intelligence. "Its mission is to help shape the global AI policy framework."

AI4ALL, Tess Posner, CEO

2018 was an exciting year for AI. We saw the increasing speed of adoption across many industries, a media fascination with AI applications and stories, and emerging concerns about AI's social implications. As AI becomes integrated into more decision making and our daily lives, it's absolutely critical that we have diversity in who is building and shaping this technology from the beginning. Diverse perspectives yield more innovative, human-centered, and ethical products. And if this powerful tool is only in the hands of a few, it could exacerbate inequities and inequalities that we are trying to fix in society today.

When people of all identities and backgrounds work together to build AI systems, the results better reflect and serve society at large.

Diversity also adds tremendous economic value and drives innovation. Companies that embrace this will tap into a deep well of creative talent that can increase their revenue, improve their products, and drive significant return on investment.

A 2016 Intel study estimated that improving ethnic and gender diversity in the U.S. tech workforce could create $470 to $570 billion in new value for the tech industry, and could add 1.2 to 1.6 % to the GDP.

Today there is a diversity crisis in AI, not just gender, but diversity of race, ethnicity, sexual orientation, and other factors. AI Index 2018 reported that 80% of AI professors at top global universities are male, and the applicant pool for AI jobs in the US is 71% male. Only 11% of science and engineering jobs in the US were held by Black and Latinx people in a 2015 National Science Foundation report, and 7% of bachelor's degrees in CS in the US were earned by women of color. This is contributing to issues like bias appear in AI ranging from facial recognition systems to criminal justice tools to translators.

At AI4ALL, we foster the next generation of leaders through education and mentorship programs and increase access. Our high school students learn AI fundamentals, apply them to solve societal problems, and to do so in deeply innovative ways. 61% of alumni have started their own AI/CS projects and some have done original AI research that has been presented at top academic conferences. We've already seen amazing results from bringing more people into the field: our alumni's research address topics they care about like healthcare, environmental science, and criminal justice reform. It's critical we invest in inclusion early with emerging technologies like AI--not as an afterthought.

80% of AI Professors at top universities are male
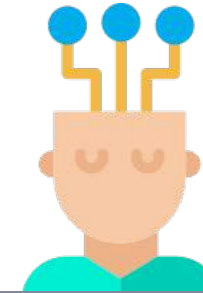
71% of applicants for AI jobs in the USA are male
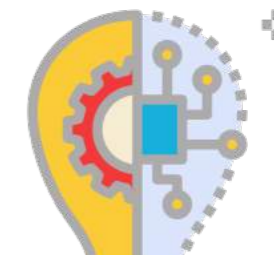
11% of science & engineering jobs are held by black & Latinx people (2015)

7% of bachelor's degrees in computer science in the US are earned by women of color (2015)

IBM Watson™ AI XPRIZE®

The global competition challenges teams to develop and demonstrate how humans can collaborate with powerful AI technologies to **tackle the world's Grand Challenges**. Assisting in encouraging new and innovative thinking, as well as motivating the next generation into AI, the winning team are awarded $5 million.

Focusing on AI for positive social impact, the XPRIZE aims to accelerate adoption of AI technologies and spark creative, innovative, and audacious demonstrations of the technology that are truly scalable to **improve the quality of life for all of society** through new models.

The UN and its agencies are collaborating with XPRIZE to further use the IBM Watson AI XPRIZE as a repository of cases to advance the sustainable development goals that will touch more than 3 billion people on the planet.
- Amir Banifatemi, Prize Lead

At XPRIZE we are dedicated to inspiring and incentivizing breakthroughs to the most critical challenges facing humanity. We address these challenges that cannot be solved by traditional means through our incentivized prize competitions. To encourage people to consider the social impact of the potential of AI, we focus the resources, talent and technology required to enable breakthroughs and accelerate the future, while driving real impact. To date (since 1994) XPRIZE has launched eight prizes worth over $150 million with plans for an additional $200 million in the coming years, all with the goal of transforming society for good with AI.

We are at a point in time where powerful and innovative technologies in the domain of AI can be leveraged to solve or minimize most problems in the world. Artificial Intelligence algorithms can help identify causes of problems, trends in data, and most importantly modeling out potential solutions for these problems.

The IBM Watson AI XPRIZE is a global competition challenging teams to develop and demonstrate how humans can collaborate with powerful AI technologies to tackle the world's Grand Challenges. Where we hope the competition creates impact is by accelerating the adoption of AI technologies and sparking creative ways that these technologies can truly scale.

We are proud to see that the AI for Good/Social Good movement is becoming an ever-popular one.

XPRIZE, and in particular this initiative, was one of the initial thought leaders that really helped to start this way of thinking about AI. Our GM of Innovation & Growth, Amir Banifatemi, involvement in the launch of the IBM Watson AI XPRIZE and ITU's AI for Good Global Summit helped ignite the AI for Good movement. As it continues to gain momentum - we are excited to work with our teams to illustrate to the world the first tangible solutions and products in the AI for Good space.

When tackling social issues, the main problem that first comes to mind is the structuring of data. The data revolution we are experiencing now as a society, where there is a preponderance of data, has in effect lead to the AI revolution we are experiencing at the same time. The more data that is available the stronger and more dynamic the machine learning algorithms become. In short - the AI revolution is only as strong as the data, and the data associated with social issues is not as robust as data from commercial industries. This is quickly changing though, so as the data from international development and social fields becomes more structured then this will help solve these issues.

Devin Krotman, XPRIZE

8.

4. THE CHALLENGES OF AI

> *I believe that AI is yet to solve the global challenges of talent wastage, unequal distribution of resources and access to opportunities, radicalization, misinformation and environment degradation. These are deep topics that require a diverse group of subject matter experts to work together. AI will play a central role but it will be informed by other disciplines.*
>
> (Babar Bhatti, MutualMind)

It would be bold to say all of the problems being identified are soluble, because that evidently won't be the case (for example, AI bias is proving a difficult nut to crack), but it's also important to acknowledge that those working on tech ethics aren't obstructive luddites. For my part, I'm completely in awe of much of the technology we've seen emerge over recent years. I understand and respect its incredible capacity to do good in the world. But I'm also highly cognizant that, left unaddressed, issues like bias, psychological harm, worklessness, data misuse, coercion, and other types of catastrophic error could be extremely damaging to humans and technology alike. (Fiona McEvoy, YouTheData.com)

What are the main problems that arise when attempting to tackle social issues?

The **primary** issue when attempting to solve social issues is that those who may have the solutions do not respect or attempt to learn about already-attempted solutions from those who actually experience the problems. While a variety of social issues may be endemic, an in-depth research and discovery process must be undertaken before subjecting impacted communities to untested technologies/solutions that ultimately, may dispossess them of their agency if they are not adequately involved in the decision-making and implementation process.

The **secondary** challenge is that due to the experimental nature of many AI applications, we as a society do not yet understand the externalities included in the wide-scale deployment of applications. Technologists must understand that, while a full accounting of externalities is near impossible, testing their tech on impacted communities can potentially serve to make problems worse, thus have a net-negative impact on already disadvantaged groups of people.

While AI has been deployed in everything from medicine (X-ray readings, mental health chatbots) to sustainability (reducing the electricity use of servers), we do not yet know the full costs associated with all of these advancements. This is not to say that there is no good that has been done (there is) but we should remain critical and not yet herald a new age of advancement without assessment.
(Michael Martin, SignalFire)
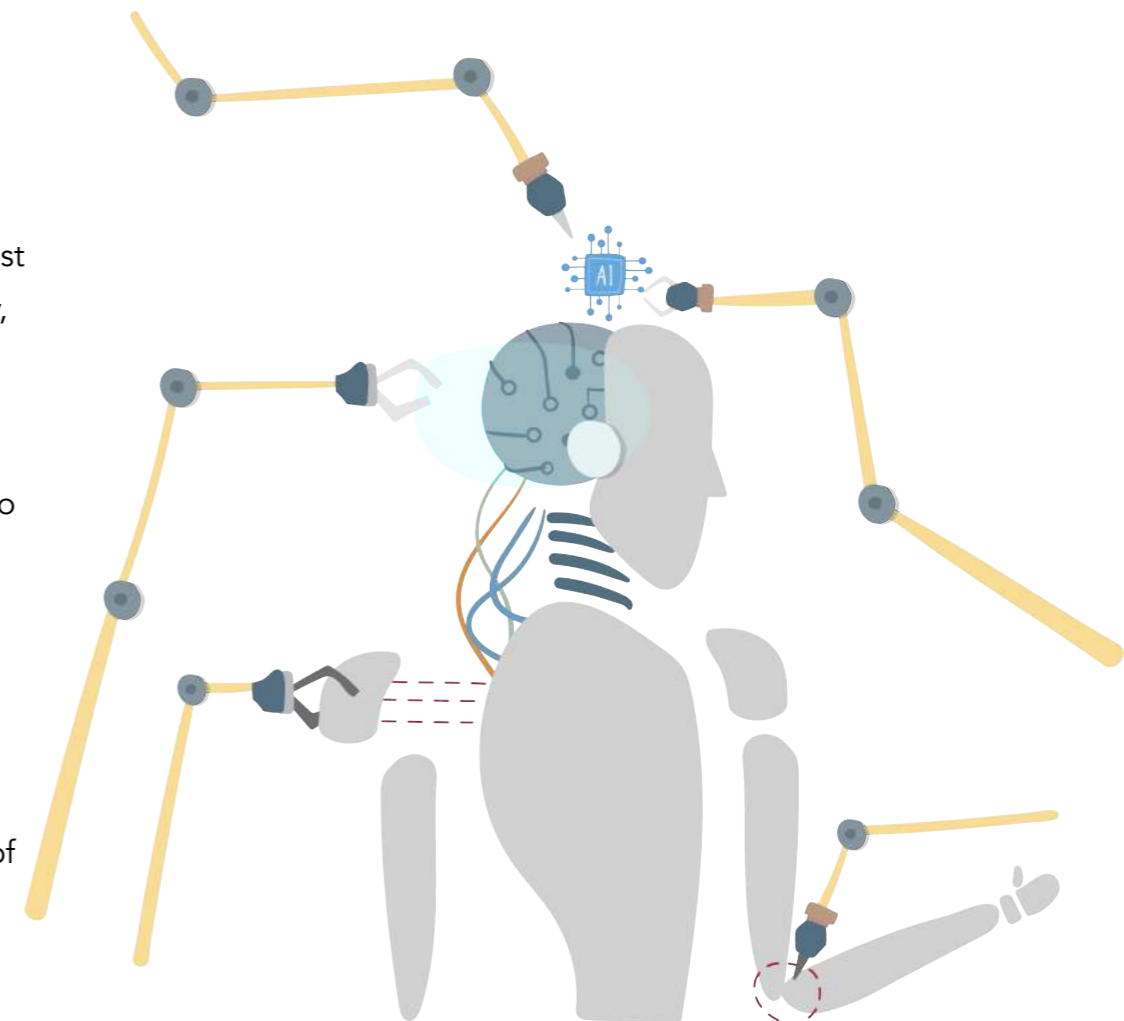
WHAT ARE SOME OF THE GLOBAL CHALLENGES AI IS YET TO SOLVE?

The ownership of artificial intelligence and the current lack of industry oversight and accountability in its development and deployment. If we as a society truly want to realize the benefits of AI for everyone, we must assure that there are adequate mechanisms and policies that support innovation, competition, and where reasonable, shared ownership of AI.

When taking a look at global challenges overall, there are still several areas that AI is still far away from transforming. The ownership of artificial intelligence and the current lack of industry oversight and accountability in its development and deployment are currently providing obstacles. Currently, most of the means by which to produce the most powerful and transformative AI (data and talent) is controlled by a handful of ultra-capitalized market players.

"Hiring the right people is essential, but also a huge challenge for small startups, especially when big companies like Facebook and Google are paying half a million. They have a monopoly that needs to be broken."

(Albert Wang, Qualcomm Ventures)

The challenge of COMPLEXITY and DIVERSITY

(David J. Gunkel, Northern Illinois University)

Social systems are **complex** and exceedingly difficult to model effectively. Human behavior, as we have recently learned with machine learning applications, is often more predictable than had been previously recognized. **This has opened the opportunity for modeling and even manipulating social systems for both good and for ill**. At the same time, however, it is clear that these models are never robust enough to capture everything and that social systems—both large and small—have many more variables and variations in play. Consequently, **identifying and accounting for the full complexity of human social behavior remains a persistent and ongoing challenge.**
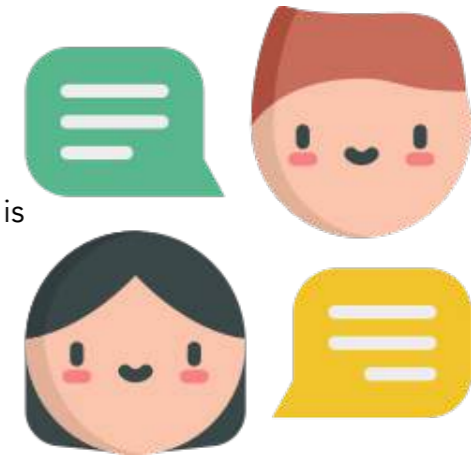
At the same time, and further complicating matters, social systems are **diverse**. They vary across culture, language, and geographical region. This means that approaches to addressing social often do not scale. **A solution that works very well and produces positive results in one region of the world with one particular human population, may not have the same kind of results, when applied elsewhere**. Human history is (for better or worse) full of examples of attempts to scale-up what are assumed to be practical solutions that had worked well in one situation but then have very different results and outcomes in other social contexts. Attention to diversity is crucial, but doing so is often easier said than done. What is needed is a proactive response to difference that is more than just lip-service and posturing.

Responding to both items—complexity and diversity—will require tapping the knowledge and experience of the so-called "soft sciences." There needs to be a more robust interdisciplinary dialogue between what C. P. Snow called the "two worlds." Research and development teams in technology need to involve and learn how to value the insights and contributions of both the social sciences and the humanities. And we have, I think, seen some very positive developments in this direction with AI ethics. This is a good beginning, but it is only a beginning. It portends what I hope will be a more integrated and interdisciplinary approach to addressing the social impact and consequences of tech innovation.

At the Deep Learning Summit In San Francisco this January 2019, RE•WORK ran closed question polls with attendees to identify where we are still facing challenges. Results demonstrate evidence that there is confidence that with the right protocols and concentration in positive areas, the potential for good is high, however there are areas that need attention:

DO YOU CONSIDER YOUR COMPANY TO BE DIVERSE?


- YES
- NO
- COULD BE DOING MORE
- ACTIVELY TRYING
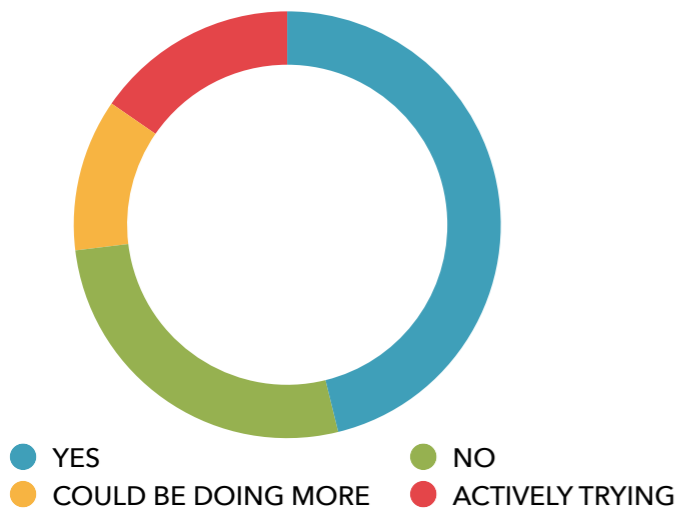
The stereotype of AI companies lacking diversity appears to be improving, with less than 25% thinking their companies are failing here, and 12% calling for improvement.

WILL AI WIDEN OR DECREASE THE WEALTH GAP?


- WIDEN
- DECREASE
- TOO EARLY TO TELL
- NOT SURE
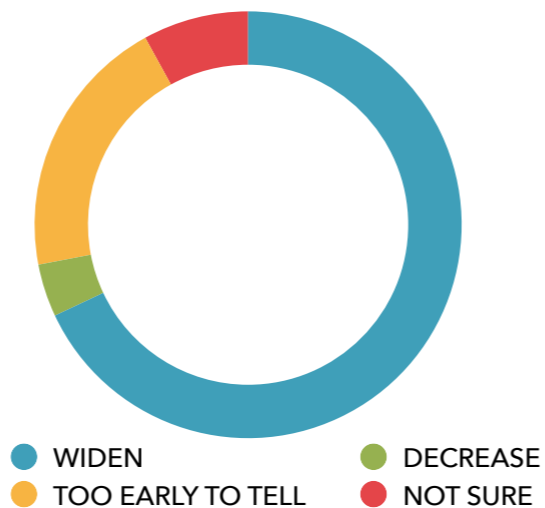
Improving society for good should encompass decreasing the wealth gap. This area needs careful attention and planning to ensure that society as a whole are benefiting from the technology.

WHERE WILL WE SEE AI BENEFIT SOCIETY THE MOST?

- EDUCATION
- SAFETY
- HEALTHCARE
- ACCESSIBILITY
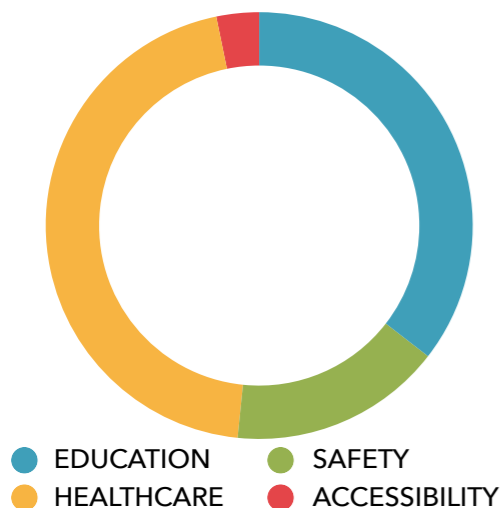
Industries will be transformed at a variety of paces, but it's important to focus on areas, such as accessibility, that may not receive as much attention.

WILL AI SOLVE THE ENVIRONMENTAL CRISIS?

- YES
- YES, BUT YEARS AWAY
- WE'RE CLOSE
- NO
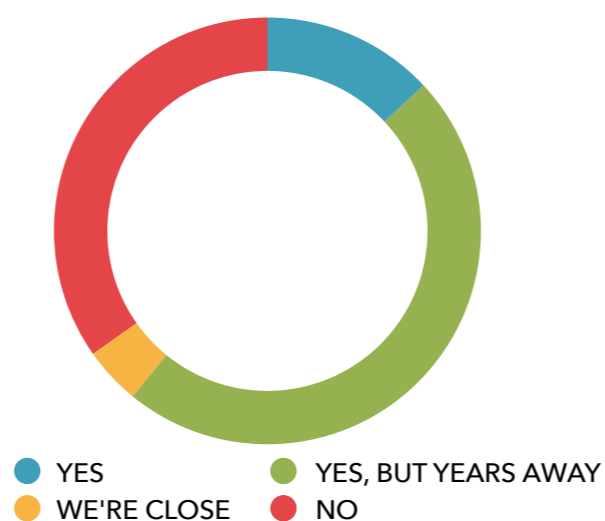
Not all progress will be immediate, and the confidence from 65% of people thinking AI will help here is promising.

ARE THERE INDUSTRIES THAT SHOULDN'T BE USING AI?

"I think using AI to manipulate elections and to manipulate voters is worrying. Cambridge Analytica was set up by Bob Mercer who was a machine learning person, and you've seen that Cambridge Analytica did a lot of damage. We have to take that seriously."
(Geoff Hinton, University of Toronto)

Karry Lu from WeWork shared his expertise on the matter: If you've spent any length of time as a data professional you will probably have noticed the proliferation of adtech/marketing jobs. It's been said in multiple places but I'll echo it here: it's a sad reflection on the human condition and late-stage capitalism that so much potential and effort is being spent optimizing targeting algorithms in an effort to get people to click on things that they probably don't need. It's unclear if AI-powered advertising even benefits the marketers it's supposed to serve, let alone consumers. Particularly, I think the VC-backed culture, the money at stake, and sheer complexity/volume of data (enabled by persistent data collection and online tracking mechanisms) have made it easy for cool-but-useless AI solutions to proliferate and further obscure the signal in the noise. It's all too simple for these companies to fudge their analyses and claim positive ROI when no value for anyone has been generated.

SHOULD THE PUBLIC FEAR AI?

AGI isn't going to happen for quite a while, so automation won't erase jobs in the medium term. "The danger of AI is that it generates an image of impenetrable mathematical formulas that no layperson can ever hope to understand...when usually, AI applications are simply a mixture of basic statistics and nonlinear function approximations. We saw this in the financial crisis of 2008 when the alleged complexity of modern finance created an atmosphere where no one bothered to check the assumptions and formulas underlying the financial instruments that triggered the crisis. I would hate for the public to become intimidated by the hype, and become less vigilant as a result." (Karry Lu, WeWork)

5. BENEFITS OF AI FOR SOCIAL GOOD

Technology is developed and deployed in specific social contexts. It is situated within a particular socio-cultural system. It is developed to respond to specific needs. It is deployed within and for the sake of a particular social formation or population. "AI for social good" seeks to take responsibility for this fact, and is currently making a positive impact in a variety of industries:

FINANCE

Financial services are important for the everyday user across the world. There are many opportunities for us whereby serving as financial institutions we can contribute to a healthier global ecosystem for every step of the socioeconomic ladder. Finance has a negative stereotype - there's an interesting phenomenon in psychology called projective identification (Melanie Klein) - when I project an idea upon you, I can do this so effectively that I can induce you into the behavior I'm projecting. I see this in finance where finance people are told 'you're bad', so some finance people think 'that's the role I'll play' and become okay with being the 'bad guy'. **I hope to affirm people who work in finance and institutions doing financial services that there's so much social impact that can be made, not by donating or having some special category called social good, but by integrating your values in every operation you're engaged in,** and through this you'll see opportunities everywhere - to help a young family get started with their first mortgage, to help an immigrant send money to their families in a different country, to help a small business be able to hire more people. These types of opportunities to help other people are everywhere, we just need to see them and resist the projections. (Mark Webber, IBM)

HEALTHCARE

Neural networks are powerful and expressive concepts, but they're also very complex. Deep Neural Networks have these neurons that are arranged in many layers, and there are many many parameters connecting these together, and the goal of 'deep learning' is to try and find the right values for these parameters. These models are being extremely useful in healthcare. **Healthcare is where you want to have more insight into what your models are actually doing and how they're applied.** (Maithra Raghu, Google Brain)

Medical images contain a wealth of information that is not being leveraged to optimize patient care. It's difficult to precisely track disease in medical images, and even more difficult to predict the course a disease will take. Arterys is using AI powered medical imaging to augment radiologists to improve accuracy, and save both time and money:
"First, we expedite a lot of the tedious work they need to do to set up, interpret and report on imaging findings. Second, **we help them much more accurately and consistently measure anatomy, so they can better track changes in patients over time**. Finally, we are helping the community standardize how disease is measured and categorized, so that there is more consistency across practices." (Daniel Golden, Arterys)

ENVIRONMENT

Deep Learning learning algorithms are being applied daily in ever more challenging domains, however in areas of huge societal importance such as ecology, sustainable resource management and environmental modeling, the analytical methods being used have not always caught up with the recent Deep Learning revolution. By using reinforcement learning, Mark Crowley of the University of Waterloo is working to manage, predict and control forest fires. Reinforcement learning will enable models to 'learn to learn' to mitigate the effects and impacts of large fires. The impact and subsequent cost of wildfires is on the rise, and whilst current simulation methods have high precision but they underestimate the scale and risk. They are also computationally expensive and cannot adapt to changing climate conditions. Mark explained how they're working to overcome this:
"Our 'learning of an agent-based model' approach could also apply to prediction and decision making for other instances of spatially spreading processes such as infectious disease and invasive species. **Sustainability and environmental domains provide a great opportunity for the AI/ML community** to step up and find solutions that will make a real difference to the lives of many people and the health of ecosystems". (Mark Crowley, University of Waterloo)

12.

EDUCATION

#1 - Automating Standard Processes

One teachers' most laborious tasks is marking. AI can mark and grade accurately, giving teachers more time to spend with their students or developing their classes for greater engagement.

#2 - An Improved Collective Experience

ML will allow schools to be compared globally and patterns such as learning ability, productivity peaks, and other metrics can be measured allowing schools to continuously improve.

#3 - Finding Gaps in the Learning Process

Teachers are overloaded - it can be easy for human error to creep in. ML models can locate these gaps, for example, if lots of students got a question wrong on a test. Teachers can then become aware of gaps and can help address them.

#4 - A Greater Level of Feedback

Automated feedback will be generated with precision - this can be either provided directly to students or to the teacher to relay.

#5 - Providing Students with More Support

In a class of 30, providing comprehensive support isn't always possible. **AI allows teachers to provide personalized support that's more helpful**.

#6 - Redefining How Information is Discovered

"With the ability to scour the entire internet for material, as well as to cater the information it's looking for to the student or teacher, AI is able to personalize to the individual's learning ability. Searching for information will have never been more concise." (TeachThought, The Future of Learning)

EQUALITY

We can move towards a more equal society with AI at the helm in several ways. By firstly accepting that technology is not neutral, and being vigilant about where the faults creep in.

Critically, **this means building in ethics at the earliest possible stage** - whether that's taking the care to understand the socio-historical context of the data used to train systems (and how they might create bias), or ensuring there's a robust feedback loop that lets a system receive information about bad decisions and adjust accordingly.

These are just two measures, there are many others being offered up by tech ethicists deeply entrenched in this area. None of this is straightforward. How can we know if a hiring algorithm filtered out the perfect candidate based on some irrelevant factor? Or if an autonomous vehicle could've made a better decision to avoid a crash but for some other competing goal? Or if a vulnerable person has been coerced into buying a product or service by relentless online targeting?

Very broadly, **in order to understand and intercept future bad decisions and influences, we need some level of AI explainability, and for humans to play a key role in continually policing, stress testing, and correcting our systems and their goals**. When anyone is developing anything for mass adoption, the maker should consider how it will affect the eventual users. (Fiona McEvoy, YouTheData.com)

ENERGY

The World Economic Forum states that there is a global demand for clean, cheap, reliable energy – and artificial intelligence is increasingly being used to help meet this need. AI can not only be used to help combat the obstacles of energy shortages, but can help with the growth of low carbon green electricity. Aidan O'Sullivan, Head of University College London's energy and AI research, says using AI to create *"forecasts for electricity demand, generation and weather can lessen the need for backup forms of power which we're currently relying on."* These backups are the likes of coal or diesel generators which are detrimental to the planet, costly, and labour intensive. On a consumer level, AI is already being used to manage energy devices in technologies that we interact with daily, such as mobile phones and laptops, and if we ingrain these technologies into electricity suppliers, AI can also find out how much electricity each appliance is consuming, helping you cut costs and do your bit from home. "*With access to exactly what it costs to run a dishwasher or TV, consumers could rapidly become disenchanted with power-hungry devices.*" Whilst these solutions are plentiful, relying on AI could leave energy networks vulnerable to cyberattacks, and relying on AI allows scope for new risks to emerge. This means that privacy and security will need to become central points of focus from suppliers and their AI teams. (Olivia Gagan, World Economic Forum)

ACCESSIBILITY: Google's Live Transcribe and Sound Amplifier aim to help the hard of hearing
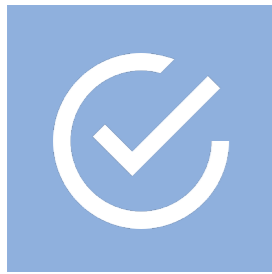
People with disabilities are the largest minority in the world (per the UN), numbering **1 billion** (per the WHO). Like any technology, AI can provide opportunities or barriers for disabled populations. **Barriers may occur when AI technologies replace existing user interfaces without taking into account the diverse needs of users.**

One example would be conversational interfaces replacing text-based services. The interactions can be easier and more natural, but only for users that are facile with speech and able to hear; thereby possibly leaving out large communities of users with disabilities. But if it is designed thoughtfully, AI has the potential to provide significant life improvements for people with disabilities across multiple disability types - visual impairment, hearing loss, motor impairment, neurodiversity, and many more.

"I have personally been deeply involved in AI applications in service of communication."

Communication breakdowns can occur in a number of ways. Deaf people may not be able to understand other speakers. Or speakers may have non-standard speech characteristics that are not readily intelligible to their communication partners. **Automatic speech recognition (ASR) can assist by presenting a transcribed version of a person's speech,** but ASR may perform poorly, especially with non-typical speech.

"I have worked on improving speech recognition technologies for 30 years and do so now at Google Research. More accurate speech recognition has been essential for improving information access and communication opportunities for deaf and hard of hearing users. Significant improvements in speech recognition accuracy support multiple Google products and services to help people with accessibility needs." One example is **Live Transcribe**, which Google announced in February 2019. Live Transcribe provides real-time continuous transcription displayed on your smart phone.

466 MILLION PEOPLE ARE DEAF OR HARD OF HEARING

Estimated by WHO. A crucial technology in empowering communication to this population is ASR. Much improvement has been made here, but there are still gaps.

GOOGLE ANNOUNCED LIVE TRANSCRIBE

A a free Android service that makes real-world conversations more accessible by bringing the power of automatic captioning into everyday, conversational use, powered by Google Cloud.
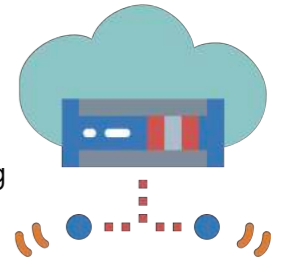
OVER 70 LANGUAGES ARE SUPPORTED

Making it accessible to over 80% of the world's population. It can be launched with a single tap from within any app, directly from the accessibility icon on the system tray.

IMPROVED WITH ON-DEVICE NEURAL NETS

ASR provides us greater accuracy, but Google wanted to reduce the network data consumption that Live Transcribe requires. This meant implementing an on-device neural network-based speech detector, built on previous work with AudioSet.

USERS COME FIRST

Google partnered with Galludet University to ensure users were satisfied, whilst keeping AI at the centre of the project. This included taking into consideration background and environmental noise, transcription confidence, and the best modality of display - the smart phone.

6. CASE STUDY: AI ETHICS
How Voice Interfaces Can Benefit Society

CATHY PEARL: BIOGRAPHY

Cathy Pearl is Head of Conversation Design Outreach at Google, and the author of the O'Reilly book, "Designing Voice User Interfaces". She's been designing and creating Voice User Interfaces (VUIs) for nearly 20 years and is passionate about helping everyone make the best conversational experiences possible. Previously, Cathy was VP of User Experience at Sense.ly, whose virtual nurse avatar helps people stay healthy. She has worked on everything from programming NASA helicopter pilot simulators to designing a conversational iPad app in which Esquire magazine's style columnist tells users what they should wear on a first date. She has an MS in Computer Science from Indiana University and BS in Cognitive Science from UC San Diego.

Much has been made about our current addiction to screens. Many of us have found ourselves trapped in an endless cycle of dopamine hits as we check our phone to see if anyone has liked our post, or replied to a pithy comment.

Some tech companies are creating technical solutions to help us monitor and limit our screen time, such as Google's Digital Wellbeing program. Many articles have been written with tips on how to break our phone addiction. But I believe one of the solutions to this problem is something that's already here: voice user interfaces.I n addition to cutting down on screen time, they also foster communal activities and help bring dignity and independence to those who may not be able to access the Internet in traditional ways.

At our house, we have a rule: no devices at the dinner table. But sometimes, we have questions! A recent sampling from our family includes: "How old is Rosanne Cash?" "When was 'A Wonderful Christmas Time' by Paul McCartney, written?" and "When was the Copper Age?" Rather than pulling out my phone, and getting sucked into my social media feed, I can simply say "Hey Google" and it's as if the Google Assistant briefly joins our dinnertime conversation.
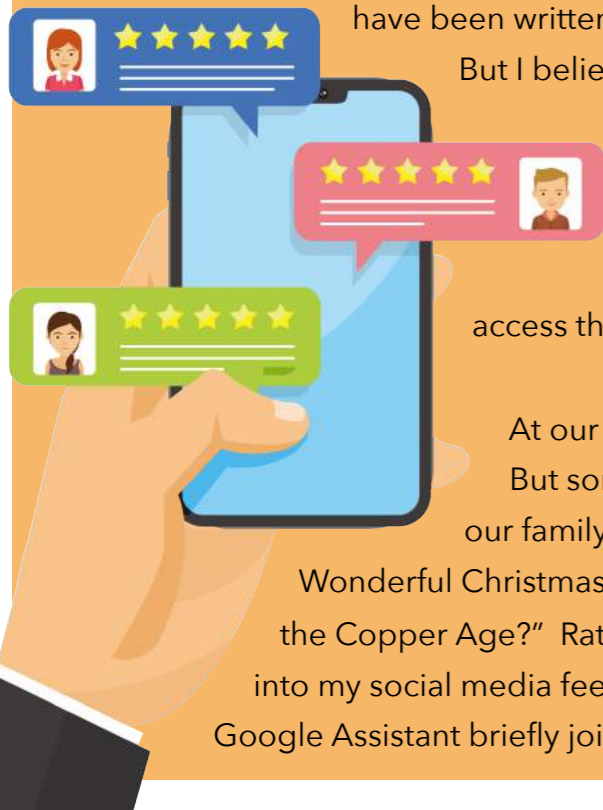
Everyone can hear the question I asked, and everyone can hear the response.

Many people report issues with having their phone in their bedrooms at night, checking that one last email before bed, and being bathed in blue light, which has been found to interfere with sleep. At the same time, many of us rely on our phones to help us through a meditation exercise, set reminders, and as an alarm clock. A voice-only device in the bedroom can still do these basic necessities, but help provide distance.

Voice interfaces can encourage community, as well. In the Spring 2018 NPR and Edison Smart Audio Report, nearly half of users say they use their smart speaker with others in their household "most of the time". In The Future of Voice and the Implications for News, a similar effect was noted:

Many respondents also found there was a social and fun element to these devices. Often they played games together, or asked questions to settle an argument, for example, the year of a film release or the age of a politician. Many appreciated features that introduced whimsy, such as 'tell me a joke'.

Rather than each individual focusing on their screen, and looking down, being able to do these activities while still looking at the world around them makes the time more communal.

Smart speakers aren't just for the young and tech-savvy. Ownership of smart speakers is increasing the most for folks in the age group 55-75. Although many seniors own smart phones or tablets, the number of people over 65 in the US who own smart phones is the smallest demographic, at 46%. Voice interfaces can be another way to stay connected in those cases, as illustrated by a quote from the "Future of Voice" study above.

15.

We also found a thrill for...older users in feeling part of the future. We came across one taxi driver in his 70s who had never been able to master a computer, a smartphone, or a tablet but learnt how to interact with his [smart speaker] within a few days.

Finally, there is the benefit of voice technology to those with disabilities such as visual impairment, difficulty with fine motor control, or memory loss. Because the devices can be connected to thermostats, TVs, fans, lights, etc, access can be given to people who may not be able to operate a remote control, or a light switch, or see the TV listings on the screen.

When designing products, we often consider the "big picture" items, but fail to think of the myriad of small things we do every day that not everyone can do. My 105-year-old grandmother once told me the worst thing about being over 100 was losing the ability to walk. She lives in an assisted living home, and can summon a caretaker at any time, but she doesn't do so for what she considers "minor" requests like adjusting the blinds to let in the best light, or asking what they're having for dinner. A voice interface could bring about these benefits.

One man created a fun "Presidential Quiz" game on a smart speaker for his father, who had recently suffered a stroke, in order to help him improve his cognitive abilities. The company Marvee specializes in voice experiences for those over 65, letting seniors send messages to friends and family members, which can be delivered as texts or emails.

In a 2017 Wall Street Journal article, "The End of Typing: The Next Billion Mobile Users Will Rely on Video and Voice", shows how a man from New Delhi, despite the fact he's not very comfortable with reading and typing, can now use voice to check train schedules, send messages to family, and download movies. Voice allows people to leapfrog screens and get straight to what is needed.

Voice interfaces are not a perfect solution, of course. They are still in early days and have many technical and design issues. And for those with hearing impairments, or with non-standard speech, the technology is currently out of reach. But they are not just a flash-in-the-pan technology; they have real benefits across many parts of society, particularly for those people who are often left behind with the latest gadgets. Voice interfaces will not be the only interface in the future, but they will be a critical one.

BY **2020**, **50%** OF ALL SEARCHES WILL BE **VOICE** SEARCHES

Source: ComScore

72% OF PEOPLE WHO OWN A VOICE-ACTIVATED SPEAKER SAY THEIR DEVICES ARE OFTEN USED AS PART OF THEIR **DAILY ROUTINE**

Source: Think With Google

40% OF ADULTS NOW USE VOICE SEARCH **ONCE PER DAY**

Source: Branded3

7. CASE STUDY: AI FOR THE ENVIRONMENT
Oscar, AI For Zero Waste

HASSAN MURAD: BIOGRAPHY

Hassan Murad is the CEO & Co-Founder of Intuitive AI where they are developing AI to empower a zero waste world. He was driven by the world's massive waste problem to begin creating solutions along with co-founder, Vivek Vyas, and divert waste away from landfills, oceans or incineration. His robotics background, that includes working on autonomous submarines, farm-mapping drones, and self-driving cars at Tesla, has inspired him to develop Oscar, AI for Zero Waste. Hassan has won numerous awards recognizing his leadership in innovation and was recently featured on Forbes as the top 5 Entrepreneurs On The Rise in AI. He holds a B.Sc. in Mechatronics from Simon Fraser University and is currently based in Vancouver, BC.

Intuitive's vision is to create a zero waste world by developing intelligence that helps buildings and spaces in tracking their waste in real-time and engaging users, by nudging them to correctly sort their waste
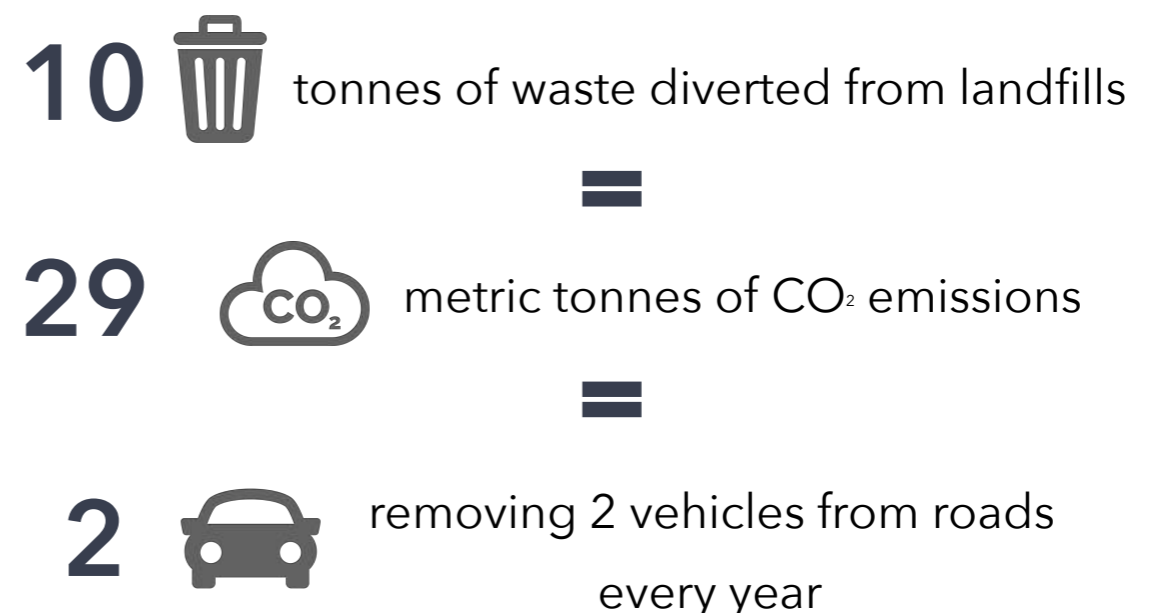
However, this vision has some major challenges:

i)   The ever-increasing volume and complexity of waste as a result of the constant increase in commercial consumerism makes it difficult for the public to recycle properly. For instance, a coffee cup will typically have three components (plastic lid, cardboard sleeve, and paper cup) that needs to be sorted and disposed three different bins.

ii) Differences in recycling guidelines among regions and municipalities which hinder effectiveness of recycling policies as well as the lack of standardization among cities, primarily driven by:

- Local haulers' recycling tech and capabilities
- The current recycling market, i.e, possibility of up-cycling materials

iii) frequent manual waste audits conducted by facilities prove to be costly, lengthy, labor intensive and misleading

According to a study by the World Bank, 98% of the world's waste is sent to landfills, dumped into oceans or being incinerated, even though a high majority of daily consumables are recyclable. This is primarily due to the high level of contaminants found in recyclables, making previously clean material practically unrecyclable and financially unmarketable. In Toronto, for every percentage point decreased in contaminated waste can create up-to $1 million in recycling cost savings every year, which can be attributed to the management and sorting costs incurred by the waste hauling and collection companies.

According to a benchmarking waste study of a university's contamination rates and with early Oscar performance in waste diversion, it is approximated that by nudging users to recycle correctly, each Oscar installed in a facility can reduce up to 10 tonnes of waste from being sent to landfill and thus stopping toxic chemicals from seeping into surrounding water tables.

10 🗑 tonnes of waste diverted from landfills

=

29 ☁CO_2 metric tonnes of CO_2 emissions

=

2 🚗 removing 2 vehicles from roads every year

With a vision to empower a zero waste world, Intuitive is taking its first steps by leading in the development of a state of the art AI system that is robust, lean, and able to efficiently run inference on the edge by integrating embedded systems and sensors.

• Oscar is designed to have an enormous social impact by targeting the issue of public empathy when it comes to recycling at the source and engaging users, by nudging them to separate their waste items at the point of disposal. This is particularly important in places with high traffic flow such as airports, malls, universities etc. It is within such settings that waste disposal is done at a fast pace and without much room for thought.

- Organizations are looking for alternative ways of conducting lengthy, labour intensive, and costly manual waste audits. Buildings will have instant access to Intuitive's dashboard that reports waste intake through real-time waste analytics.
- Oscar will update in real time and educate the public pertaining to their municipalities' latest recycling guidelines. With a retrofittable module that is easily attached to any existing waste bins, there is immense opportunity for expansive and effective public education.

Oscar's AI System leverages Computer Vision and Machine Learning

Oscar recognizes varied items within the hands of users, i.e, a plastic cup, used paper napkins, empty snack chips bag, biodegradable spoon. Through advanced computer vision research and enhanced machine learning algorithms, Oscar is trained to detect partially visible items as well as recognize a variety of body postures of identifiable approaching users.

The main challenge with such algorithms is to facilitate quick edge processing so that Oscar can be as dynamic as possible and make its predictions in real-time. Additionally, data quality is just as important as data quantity to train Oscar effectively; for instance, a coffee cup might look different in a brightly lit food court as compared to a dim office kitchen. Intuitive is building upon its current

technology by training the AI systems with thousands of images of waste items held by users in every possible way, emulating the behaviors at places such as a food court, office kitchens etc. This process is designed to ensure that Oscar will be able to run real time predictions, engaging users in an intuitive and dynamic manner.

Creating AI for social good has the interesting challenge of dealing with social dangers such as detecting faces as people approach the bin. Maintaining the privacy of the public is critical, and it is to this end that Intuitive is continually researching and improving their enhanced machine learning algorithms to anonymize any personally identifiable information. Oscar has been trained to recognize different body movements and poses so that a wide range of users approaching the bin have their faces instantaneously detected and blurred.

Another research opportunity Intuitive is aiming to engage in is how the fill capacity of a waste item such as a coffee cup has the propensity to change the approaching user's body movements as it will alter the way they hold the coffee cup, for instance. Oscar's prediction engine is trained to know what the fill levels of the bins are before the waste item has been disposed. Additionally, there will be research done on how to implement different gestures so that Oscar will be adept at recognizing custodial staff when it's time to empty out the bins.

As a company developing AI for social good, it is important that Intuitive remains social and not cocooned into solitary work. It is to this end that the company aims to partner with sustainability champions that will invite Oscar into their organizations and are keen on implementing a zero waste strategy within their premises. The world is at an opportune moment with cutting edge proprietary technology to solve a global challenge, and Intuitive is diligently working on making Oscar an industry revolutionizing product that will divert waste away from landfills all over the world. With AI deemed the 4th industrial revolution and the waste management industry at a crossroad with the recent global recycling crisis, it seems that waste management is an area where AI will have a positive social impact.

let's meet!
intuitive
ai for zero waste

compost   paper   recycle   landfill
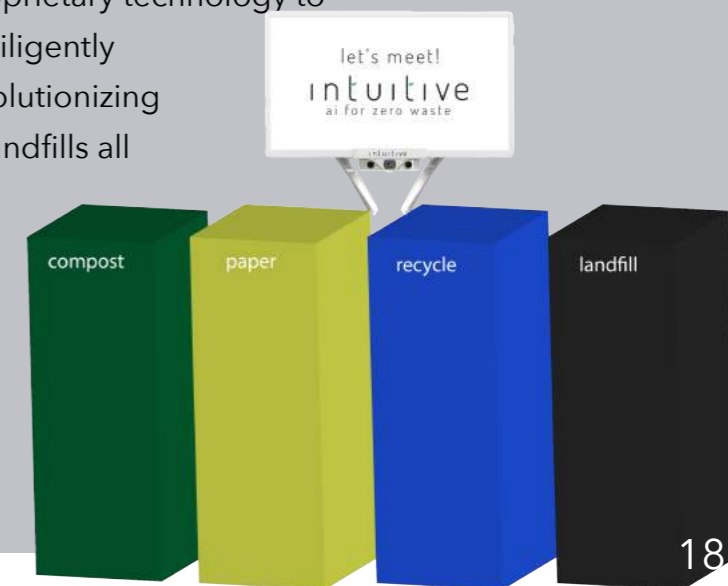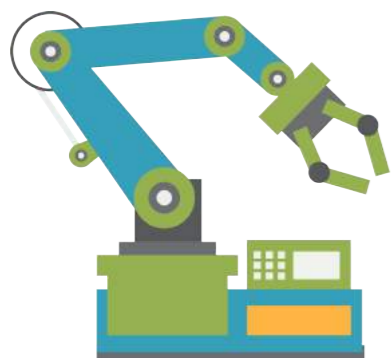
8. CASE STUDY: AI & LABOR
10 Impacts AI Will Have on Jobs and Society to be Considered

MAREK HAVRDA & WILL MILLERSHIP, GoodAI: BIOGRAPHY

The long-term goal at GoodAI is to build general AI brain software that will automate cognitive processes in science, technology, business, and other fields. Our future AI brains will perceive stimuli in the same way that a human does – by seeing, feeling, interacting, and learning – and use this data to generate behavior, perform tasks, and respond to motivations given by human mentors. GoodAI's research has several areas of concentration: **Framework, Roadmaps, School for AI**, **Growing Topology Architectures, AI Roadmap Institute, Brain Simulator**, **Arnold Simulator**, **General AI Challenge.**

Machines are beginning to fill jobs which have traditionally been reserved for humans, such as driving a car or prescribing medical treatment. The paths we will take, or will be taken on, will influence the basics of how societies function with its impacts on social cohesion and democracy. According to Amara's law **we tend to overestimate the effect of technology in the short run and underestimate it in the long run**. Therefore, it seems useful to discuss these issues so that we can prepare better for various eventualities in the future:

1. A PARADIGM SHIFT IN WORK PATTERNS

Today the use of ("narrow") AI systems which carry out defined (structured) tasks is enabling the automation of more tasks than ever. We are seeing an increase in the use of autonomous machines, due to the fact that they cost less, are more efficient, and can ensure a higher quality of work when compared to humans. As well as advances in narrow AI we are seeing increasing investment in AGI. This not only has the potential to be superior to humans at structured tasks, but in a wide variety of domains which involve general problem-solving. Although there is not full agreement in the scientific community on whether AGI may be reached, many experts think it is possible or that we can at least expect significant research advances leading to transformative AI.

400 million to **800 million** people will see their jobs automated by **2030**

2. EMERGENCE OF SUPER-LABOUR

Overall, we may experience the **emergence of "super-labour"**: a labour defined by the super high added value of human activity due to augmentation by AI. Apart from the ability to deploy AI, super-labour will be characterized by creativity and the ability to co-direct and supervise safe exploration of business opportunities together with perseverance in attaining defined goals. Many white-collar jobs today consist of information processing and making recommendations for subsequent action, which will be carried out by AI. These trends will lead to structural changes on labour markets and the speed of change may be accelerating with the possibility of millions of jobs being lost globally in several years.

3. OVERHAUL OF THE WELFARE STATE

Although there will be new jobs created in the private sector, e.g. in recreation, entertainment and arts, this profound change may still bring about a need to rapidly update the role of the state due to the need to **maintain social stability.**

Some historians claim that the welfare state was created to placate workers, in order to prevent the rise of socialism. It is likely that apart from the increased use of technology in the provision of existing public services, we will see a move towards non-state actors involved in welfare, due to the ability to provide better and cheaper services, but more importantly, they will have the ability to adjust faster to these changing conditions.

4. PREVIOUSLY UNPAID JOBS BECOME PAID

This will effect many industries: With robots starting to take up menial tasks it could free up time for care workers to focus on the human aspects of their jobs. The size of classes could be reduced in the schools with better feedback to pupils and more. Patterns like this emerging in childcare as well as in social and healthcare would give carers more time to focus on the important social and emotional aspects of their jobs.

This could accelerate the change of **previously unpaid jobs, like conversing with the elderly or reading to children, becoming paid ones**. Other areas of new or rediscovered public services (and job) creation may include paid **community-building** initiatives and also more traditional public **infrastructure projects** such as parks, green spaces, and leisure infrastructure.

5. REINVENTING FINANCING OF PUBLIC SERVICES

Considering welfare reform we may need to find ways to automatically adjust the tax base in order to generate the needed revenue as traditional sources of taxes. For example, **a higher rate of VAT for less human labour intensive (i.e. more AI intensive) production** would generate both more tax revenue, as well as encourage companies to keep humans in employment. Furthermore, a higher income tax on those profiting most from AI could help speed up the implementation of a new welfare system, possibly through various types of social stipends, **subsidized vouchers or Universal Basic Income (UBI), where certain services (potentially including some of the new paid services described above) are procured directly by citizens**. Other services such as free education, free healthcare, free public transport might be provided directly by public bodies.

6. THE ROLE OF NEW COMPANIES

It may become their responsibility to derive **new models to share AI co-generated wealth, either directly in the form of various grants or in the form of capital gains**. Nevertheless, in the longer-term, if a state has the know-how to build and deploy AGI together with the access to natural resources, the taxation may be reduced even to zero while being able to provide for their citizens.

7. NEW TECHNIQUES FOR MOTIVATION AND RE-TRAINING

It will take time to automate jobs requiring high dexterity of human fingers or natural interaction in an unstructured environment. Completely new jobs will be created, for example, a new generation of shopping assistants or better high-skill advisors such as wine sommeliers or chefs. These advisors will use both the data analysis of consumer preferences as well as in-depth expertise. For all these jobs skilled workers will be needed, at least before AGI allows for their automation. We need to establish motivation mechanisms which will help speed up workforce retraining in order to increase the odds of securing new jobs for those who are under threat of job displacement.

8. OVERHAUL OF OUR EDUCATION SYSTEMS

We need to focus on how humans work and how we put into use knowledge, skills and attitudes acquired. Most of us learn better through experience than by analyzing - this makes us distinct as humans, and could help us develop a competitive advantage in. We need to focus on what a robot cannot do and where we outperform machines. Among key competencies will be mastering the ability to learn new things which allow for continuous life-long learning. Other useful competencies may include so called 4Cs (**critical thinking, collaboration, communication, creativity**) competencies.

9. COMPETITION POLICY

Narrow AI will soon allow the automation of many structured tasks performed by current human employees such as assembly line workers, warehouse operators, accountants or paralegals. With further advances, AI will fare better and better in less structured environments. A company that does not adopt new technology will face higher costs of human labour and will be driven out of business. Automation will allow companies to accumulate vast amounts of new data which will allow them to be successful in the race to improve their products and attract more clients, which in turn will feed in more data and further accelerate the process.

PACE OF TRANSFORMATION 10.

It is important that governments and companies can keep up with the pace at which AI will transform work practices to avoid a situation similar to "Rust Belt syndrome" which saw whole regions suffer long-term economic downturn, urban decay, and population loss. The **mismatch between workforce skill set and the future (automated) labour market** is a crucial area that needs to be addressed. There seems to be only a limited sense of urgency to act, e.g. in a recent poll by Gallup, the majority of Americans expect that AI will have a negative impact on the workforce, but at the same time, the vast majority does not fear losing *their* own job to AI. We need to think about how to provide factual information to the general public and what role governments and companies will have in this effort. In any case, AI can at least help us retrain people more efficiently, among other things, through personalized learning.

The main determining factor will be who has higher ability to adapt to changing conditions. There will be areas where machines can be adopted much faster, e.g. changing size of clothes to make sure they are fully personalized. At the same time, there will be areas where humans will outperform machines when it comes to adaptability. In any case, humans (or at least some of them) will need to master mental capabilities to master the change and thrive. At the level of countries it will be analogous. Those governments who will adopt their policies faster to the new conditions will face less disruptions and their countries will be more competitive.

9. WHAT'S NEXT?

OVERCOMING BOTTLENECKS...

In many industries such as healthcare and finance, the **availability and access to data** is stunting the progress of AI advancements. "Data accessibility remains a significant challenge. *Resolving it will require a willingness, by both private and public sector organizations, to make data available.*"

The expert **talent** in AI needs to develop in order to have enough qualified professionals to carry out the ambitious goals. In encouraging the next generation into AI, we need to ensure that the entrants come from diverse backgrounds, and are instilled with ethical groundwork from the start. Companies and initiatives such as **AI4ALL** will be integral in nurturing this talent.

We need **public and democratic oversight** of the industry as soon as possible. Ethical codes developed by individual companies are an important part of the equation but cannot be the substitute for the public engagement in the development of these technologies that have the power to fundamentally transform our relationship to each other and our institutions. If we as a society truly want to realize the benefits of AI for all, we must assure that there are *adequate mechanisms and policies that support innovation, competition, and where reasonable, shared ownership of AI*. (Michael Martin, SignalFire)
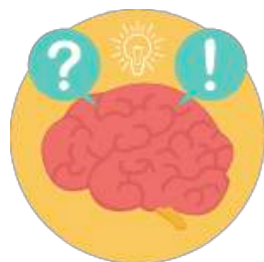
Ensuring that AI **does not inherit human limitations**, faults, or biases is something that is understandably motivated by a desire to engineer systems and devices that are safe, reliable, and trustworthy. It is a laudable goal, but it misses something just as important.

What we typically call "errors" or "faults" produced by AI (especially deep learning neural networks trained on big data) are often useful sources of social insight.

These errors typically reveal hidden patterns in data, especially unacknowledged prejudices or biases that we simply did not nor could not see for ourselves. In other words, machine learning algorithms that reproduce various "intrinsic faults of humans" are not necessarily bad outcomes to be avoided at all costs, especially if these "mistakes" lead us to identify and reexamine often unquestioned assumptions about data, the method of their collection, and their inherent limitations. (David J. Gunkel, Northern Illinois University)

"Outside of critical, high profile areas like system bias and data privacy - we're still trying to anticipate problems that could manifest down the line. It's important to remember that there has been a huge amount of development over an incredibly short period. Left unaddressed, issues like bias, psychological harm, worklessness, data misuse, coercion, and other types of catastrophic error could be extremely damaging to humans and technology alike. When anyone is developing anything for mass adoption, the maker should consider how it will affect the eventual users." (Fiona McEvoy, YouTheData.com)

Safe use and security are essential for progressions to be made in the implementation of AI to solve some of societies biggest challenges.

Without the correct policies and processes in place, their mass production and adoption is not possible. AI must be explainable whilst maintaining a high level of learning performance and accuracy, allowing humans to understand and explain the decisions the machines make. Ultimately, the responsibility of this stems from the correct practices, processes and protocols put in place by both institutions and governments, as well as corporations and technology companies. Significant accentuation must be placed on ensuring that AI can be broken down, interpreted and explained. This, in turn, will create artificial intelligence models that are more transparent and open to investigation. (Ben Dickinson, TeCHTALKS)

9. LEARN MORE

WHITE PAPER:

Should You be Using AI in Your Business?
This paper explores the application of AI in business with research contributions from leading minds in the field including Ankur Handa, Research Scientist, OpenAI, Ian Goodfellow, Senior Research Scientist, Google Brain. Jörg Bornschien, Research Scientist, DeepMind, Maggie Mhanna, Data Scientist, Renault Digital and many more.

WHITE PAPER:

The Ethical Implications of AI
This paper looks at real world examples of AI and takes into consideration the research and applications focusing on bias, responsibility and more. Expert opinions from academics, industry leaders, researchers, CEOs, founders and many more are included to comment on the ethical concerns and solutions across multiple industries including healthcare, finance, sustainability, research and transport.

WHITE PAPER:

Privacy & Security in AI
Companies, governments, policies and all other areas of businesses and research departments must be aware of concerns and regulations. Throughout this paper, we look at some of the key privacy and security concerns of the technology with the goal of understanding the risks involved, and how we can ensure AI is used safely and with both personal and corporate privacy at the helm.

PODCAST:
RE•WORK's podcast showcases leading female minds in AI, Deep Learning and Machine Learning. In each episode, we speak to CEOs, CTOs, Data Scientists, Engineers, Researchers and Industry Professionals to learn about their cutting edge work and advances, as well as their impact on AI and their place in the industry.

BLOG:
With three weekly posts, the blog covers emerging news and trends in AI as well as featuring interviews from event speakers and contributors. RE•WORK's network of technical AI experts contribute guest posts to the blog covering topics such as NLP, computer vision, medical imaging, AI assistants, neural networks and many more.

10. REFERENCES

- United Nations Sustainable Development Goals, *United Nations Department of Global Communications*, January 2019

- DeepMind Health, *About*, [ONLINE] [Accessed 5 February 2019]

- Microsoft, *Microsoft translator for Education*, [ONLINE] [Accessed 5 February 2019]

- World Economic Forum, *Here's How AI Fits into the Future of Energy*, Olivia Gagan, May 2018

- AI4ALL, *About* [ONLINE] [Accessed 5 February 2019]

- Google AI, *About* [ONLINE] [Accessed 5 February 2019]

- XPRIZE Foundation, *About*, [ONLINE] [Accessed 5 February 2019]

- IEEE Standards Association, The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, [ONLINE] [Accessed 5 February 2019]

- The AI Initiative, About, [ONLINE] [Accessed 5 February 2019]

- Gizmodo, *The 'Godfather of Deep Learning' on Why We Need to Ensure AI Doesn't Just Benefit the Rich,* Martin Ford, January 2019

- TeachThought, *10 Roles for Artificial Intelligence in Education*, [ONLINE] [Accessed 5 February 2019]

- AI for Good, *How AI is Making Energy Smarter and Cleaner*, Mary Catherine O'Connor, July 2017

- AI for Good, *How AI Could Smarten Up Our Water System*, Mary Catherine O'Connor, July 2017

- AI for Good, *Why Farmers Are Turning to AI to Boost Yields*, Mary Catherine O'Connor, Sept 2017

- Stanford Social Innovation, *AI as A Force for Good*, Gideon Rosenblatt & Abhishek Gupta, June 2018

- Nesta, *AI for Good: Is it for Real?*, Geoff Mulgan, July 2019

- TechTalks, *Explainable AI: Viewing the world through the eyes of neural networks*, Ben Dickson, Feb 2019

- News Center, *Machine Learning to Help Optimize Traffic and Reduce Pollution*, Julie Chao, Oct 2018

- McKinsey & Company, *Applying Artificial Intelligence for Social Good*, Michael Chui, Martin Harrysson, James Manyika, Roger Roberts, Pieter Nel, Ashely van Heteren, Nov 2018

- Google, *Digital Wellbeing*, [ONLINE] [Accessed 5 February 2019]

- CNBC, *These simple steps will help you stop checking your phone so much*, Jillian D'Onfro, Jan 2018

- Deloitte, *Talking to myself: Voice assistant usage goes mainstream*, Podcast hosted by Kim Barnes, 2018

- Pew Research Center, *Internet and Tech*, [ONLINE] [Accessed 5 February 2019]

- Sean Shadman, *How this Google Home app helped my father after his stroke*, Sean Shadman, Nov 2018

- The Wall Street Journey, *The End of Typing: The Next Billion Mobile Users Will Rely on Video and Voice*, Eric Bellman, Aug 2017

- Huffington Post, *A Sad Number Of Americans Sleep With Their Smartphone In Their Hand*, Alexandra Ma, 2018

- WebMD, *Your Tablet and Smartphone Is Ruining Your Sleep*, Alan Mozes, May 2018

- National Public Media, *The Smart Audio Report, Winter 2018*, [ONLINE] [Accessed 5 February 2019]

- Digital News Report, *The Future of Voice and the Implications for News, Nic Newman*, [ONLINE] [Accessed 5 February 2019]

- Wikipedia, *Roy Amara*, [ONLINE] [Accessed 5 February 2019]

- Wired, *Tug, The Busy Little Robot Nurse, Will See You Now*, Matt Simon, Nov 17

- Wikipedia, *Flexicurity*, [ONLINE] [Accessed 5 February 2019]

- BattleforKids, *Together, we can realize the power and promise of 21st century learning for every student*, [ONLINE] [Accessed 5 February 2019]

- Northeastern University, *Views on the Impact of Artificial Intelligence and Higher Education's Response*, [ONLINE] [Accessed 5 February 2019]