TechnicaCuriosa presents

The Global Competition to Advance Al for Good

147 teams. \$5 million purse. <u>1 inno</u>vative competition.





Technica Curiosa looks at how the IBM Watson AI XPRIZE is incentivizing teams to develop and demonstrate how humans can collaborate with powerful AI technologies to tackle the world's grand challenges.

TechnícaCuríosa





Popular Astronomy



The XPRIZE Foundation believes that you get what you incentivize—and that without a target, you will miss it every time. Leveraging these two core observations, the XPRIZE team designs and manages innovative competitions directed toward solving the world's grand challenges. And now it is turning up amazingly inventive applications of artificial intelligence to do just that. Through the IBM Watson XPRIZE, 59 semifinalist teams are now vying for the competition's \$5 million purse.

A Brief History of Incentive Prizes -

Comparatively speaking, determining latitude was a piece of cake: just note the altitude of the sun at noon and look up the sun's declination for the day on a table, or, if navigating at night, by the position of guiding stars above the horizon. Finding longitude, however, was an entirely different matter. For that, early navigators had to rely on dead reckoning—not so easy on long voyages and when out of sight of land—a condition that sometimes ended in tragedy. Galileo, Isaac Newton, and Edmund Halley all lent their minds to the longitude pro-

blem, but even their sophisticated astronomical methods came up short. The problem was so great that in 1714, England's Parliament offered a £20,000 prize to anyone who could solve it. Thus was born the first "grand challenge"—one whose outcome would forever alter the course of global navigation.

A great many other grand challenges would follow, incentivizing innovators to solve difficult but important problems. Many, though, are unaware that Charles Lindbergh's 1927 solo crossing of the Atlantic was also motivated by a prize—offered on May 22, 1919! The challenge was issued by New York City hotel owner Raymond Orteig, who was inspired to create the prize upon hearing WWI ace Eddie Rickenbacker speak of anticipating the day Self-educated English clockmaker John Harrison (1693-1776) solved the longitude problem with his chronometer—a friction-free timepiece, invulnerable to pitch and roll, temperature, and humidity—that would carry the true time from the home port to any destination.

3

when America and France would be linked by air.

"Gentlemen," Orteig's announcement began, "as a stimulus to the courageous aviators, I desire to offer, through the auspices and regulations of the Aero Club of America, a prize of \$25,000 to the first aviator of any Allied Country crossing the Atlantic in one flight, from Paris to New York or New York to Paris, all other details in your care."

Eight years later—to the very $d\alpha y$ —Lindbergh would claim the purse. The effects of the Orteig prize, however, were far more reaching than the crossing of the Atlantic. In the first place, it stimulated a level of investment that was actually 16 times greater than that of the value of the prize

itself. And once the myriad competitive forces were put in motion, an industry was created. The number of U.S. airline pas-

> sengers skyrocketed from 5,782 to 173,405 in the two years that followed





XPRIZE founder Peter Diamandis

Lindbergh's accomplishment. In 1927 alone, there was a 300 percent increase in applications for pilot's licenses, and a 400 percent increase in licensed aircraft in the United States.

Fast forward 77 years to May 1996, where we meet up with engineer, physician, and entrepreneur Peter Diamandis, who set out to bring the Orteig prize up to date with a \$10 million prize for the first civilian organiza-

tion that could launch a reusable manned spacecraft 100 kilometers into space twice within two weeks. Just as the Orteig prize opened up aviation, Diamandis' "XPRIZE" aimed to spur the development of a new industry—civilian space travel. And that it did.

The competition was won by Burt Rutan, whose SpaceShipOne development was backed by Paul Allen, and later acquired by Sir Richard Branson, who dubbed his new venture Virgin Galactic—the world's first commercial "spaceline." The competition drew more than 26

Burt Rutan's SpaceShipOne



teams from seven nations to commit more than \$100 million in technology investments in pursuit of the prize. And since the prize was won, a half dozen new companies entered the nascent market, with nearly \$1 billion invested. That's some leverage. More important, the prize made possible world-changing breakthroughs that neither government nor industry seemed able to produce.

Incentivizing AI for Good

In June, 2017, the XPRIZE Foundation, together with the United Nations and the International Telecommunications Union (ITU), hosted its inaugural *AI for Good Global Summit*. In his address to the packed hall, XPRIZE CEO Marcus Shingles asked those in attendance if any one of



XPRIZE CEO Marcus Shingles addressing the AI for Good Summit.

them had any confidence that business, industry, or government will solve the world's most pressing problems. Not a single hand was raised.

That lack of confidence is the very reason the XPRIZE Foundation exists. Its aim is to harness the ingenuity, resourcefulness, and willpower of a new type of problem solver—the individual entrepreneur—a mission-driven class of individuals who are willing to take on audacious goals, that if met, can have world-changing impact. To these ends, recent and current prizes include the Google Lunar XPRIZE, the Global Learning XPRIZE, the NRG Cosia Carbon XPRIZE, the Qualcomm Tricorder XPRIZE, and the IBM Watson AI XPRIZE—a contest that seeks to elevate applications of artificial intelligence (AI) to address grand challenges.

Amir Banifatemi leads the IBM Watson AI XPRIZE—and he's a natural for the role. A seasoned technology entrepreneur and investor, he has witnessed firsthand how the world's great innovators successfully execute on ambitious visions. Moreover, his own work is steeped in AI-driven man-and-machine interaction. "Really," he says, "this prize is an opportunity to be



Prize lead Amir Banifatemi at the heart of the action—a very timely opportunity to uncover and motivate the development of so many fascinating applications."

Banifatemi thinks of the XPRIZE Foundation as an engine of innovation. "The mission of XPRIZE," he explains, "is to bring radical breakthroughs for the benefit of humanity. We do that through incentivized prizes and competitions designed around goals that are very audacious, but achievable. So from that perspective, we're pushing the boundaries of competitions in a way that they not only affect the foreseeable future, but can also create new industries. This is where the engine really kicks

in. Industry and government are pushing in a number of different directions, but sometimes they might not be able to advance innovation at a sufficient pace. We accelerate that pace with opportunities that really push things forward."

Competition, it turns out, is a tremendously effective energizer. "Competition awakens," Banifatemi continues, "a sense of achievement and a sense of accomplishment and a sense of possibility—the kinds of challenges and opportunities that high achievers respond to. It's not so much about research or launching a startup—it's about *are you able to solve that problem and win this?*"

The XPRIZE competitions are designed to achieve three primary goals:

Attract mavericks from outside the addressed sectors who take new approaches and think creatively about difficult problems.

2 Create breakthrough outcomes that are real and meaningful through competitions that have measurable finish lines and promote widespread adoption of the resulting innovations.

3 Generate massive financial leverage to drive the new solutions forward.

That last point is particularly compelling. The XPRIZE Foundation has established a philanthropic model in which offering a prize for achieving a specific goal actually stimulates entrepreneurial investment that produces a 10 times or greater return on the prize purse and at least 100X in follow-on investment and social benefit.

"We have observed," Banifatemi explains, "across multiple prizes that there is a factor of between 10 and 40 times the investment in the prizes. Because all the competing teams require the means to compete for the prize, they're supported not only by financial resources but also by a whole ecosystem of supporters and mentors and experts and in-kind donations. It's logical that the pull that we create to bring entrepreneurs and innovators out of the woodwork also attracts many supporters, because the value of the winning is more than just money. And because XPRIZE has no claim on the resulting intellectual property—the IP remains with each of the teams—they can further develop commercially viable routes by virtue of the

competition. So it only makes sense for investors and stakeholders to line up with the teams in order to benefit from any societal or beneficial impact the prize can contribute."

> One of the defining aspects of the IBM Watson AI XPRIZE is its emphasis on AI as a "force for good," particularly in the context of the UN's 17 Sustainable Development Goals, which span education to clean water, health to energy, innovation to economic growth. AI, of course, is being brought to bear in all these areas, but the prize's objective is to accelerate technology innovations to meet these and other grand challenges at a larger, if not exponential, scale by inspiring and incentivizing a world full of potential problem solvers.

Al, of course, is not without its controversies. While the promise of Al's benefits touches nearly every aspect of society, like

most transformative technologies, it also poses certain risks and costs in areas including jobs, the economy, and safety, not to mention the many attendant ethical and legal questions. But also, as we've seen in past waves of technological disruption, the Al-driven wave will create new kinds of jobs, even as it improves services and lowers the cost of many goods, raising living standards for everyone in the bargain. To this end, the XPRIZE

does not consider the Al/human interaction calculus a zero-sum game. Rather, AI will find its best and highest use in its capacity to complement, augment, and enhance human capabilities.

M

đ

6

"As we learn of all the potentials that AI provides," Banifatemi says, "of course there will be polarized views about how intelligent machines will challenge our current way of life. We see this prize as an opportunity for the public, the governments, and other stakeholders to gain a better understanding of what AI can actually do, that it is not about man versus machine, but rather how we The mission of XPRIZE is to bring radical breakthroughs for the benefit of humanity. We do that through incentivized prizes and competitions designed around goals that are very audacious, but achievable."

AI has reached its Inflection Point and with it, its Exponential Potential

If the numbers of startup acquisitions, patents filed, journal articles published, and dramatic proofs of concepts staged are valid signals, AI is full swing into the knee of its growth curve. Indeed, from the new Go champ to Apple's Siri, Google's driverless car to Hanson Robotics' Sophia, it would appear that the world is pregnant with AI.

Hundreds of startups are raising billions of dollars to deploy AI in virtually every area of business, industry, and consumer worlds. The top 100 startups tracked by CB Insights have raised \$3.8B in aggregate funding across 263 deals since 2012, while incumbent companies in every industry are gobbling them up. At the time of this writing, Google has acquired 12 AI startups, with Microsoft, Apple, Amazon, Facebook, Intel, and others at its heels for more.

Then there's the number of Web of Science-indexed journal articles mentioning "deep learning" that have increased dramatically. The trends also reveal the increasingly global nature of research, with the United States no longer leading the world in publication numbers, or even publications receiving at least one citation.

And a search of the Derwent World Patents Index for "deep learning" or "deep neural net" shows a similar rise.





Journal articles mentioning "deep learning" or "deep neural network," by nation.



Analysis of number of patents using term "deep learning" or "deep neural net."



are better together. Because AI can provide superpowers to every human, the question becomes, what kinds of problems are we now able to solve? Part of the design of the prize is to advance and emphasize the aspect of collaboration, to explore and develop innovative scenarios that showcase the collaborative value of AI."

Banifatemi points out many examples of such human-machine collaboration. "In the medical field, to spotlight just one application, consider the work that radiologists do daily where AI can have tremendous impact. Radiologists spend an inordinate amount of time reviewing medical images, looking for the slightest anomalies that might indicate disease. Machines can do that at least as efficiently as, and probably better than, humans. In the 2016 Camelyon Grand Challenge for metastatic cancer detection, for example, the top-performing entry in the competition was an **Al-based computational system** that achieved an error rate of 7.5%. A pathologist reviewing the same set of evaluation images

Because AI can provide superpowers to every human, the question becomes, what kinds of problems are we now able to solve?"

achieved an error rate of 3.5%. Combining the predictions of the AI system with those of the pathologist lowered the error rate down to 0.5%, representing an 85% reduction in error. Combining human capabilities with machine capabilities provides better insight altogether—it's not a 'them or us' proposition."

INTERVIEW: Inside the Workings of the IBM Watson AI XPRIZE –

With the semi-finalist Downselect, 59 teams representing 14 countries are competing in the \$5M IBM Watson AI XPRIZE, a four-year global competition to develop and demonstrate how humans can collaborate with artificial intelligence technologies to tackle some of the world's greatest challenges. This is XPRIZE's first "open" competition, wherein the teams have defined their own goals and chosen their application domains across a variety of fields. We sat down with Amir Banifatemi to learn more about the objectives of the prize and mechanics of managing it.

As one of eight active prizes, how do you see the IBM Watson AI Prize fitting into the larger XPRIZE scheme of things?

The prize was designed about two and a half years ago and launched in June, 2016. We all have observed and witnessed and know that AI is one of the exponential technologies that is redefining everything—how we live, how we work, how we communicate, our environment, how we provide healthcare, how we support our planet, and more generally speaking, redefining how we innovate. So we like to think of AI as the front page of innovation, or to use a metaphor, the user interface of innovation. So given that the XPRIZE is a catalyst for innovation, a prize that puts human and machine cooperation at its core was very important. The competition really serves to highlight of all the domains and challenges that humanity faces. By making it an open competition, we made it inclusive of all the challenges that AI can address, and it invites individuals and teams to create novel solutions that will meet those goals.

What is IBM Watson's role in the prize?

IBM is a sponsor of this prize. Like many large organizations that understand the impact that AI will have on society, IBM is one of the leaders in pushing its usage, and of course, in using computing to transform industries, help professionals do their jobs better, and solve important challenges. To these ends, Watson represents a new era in cognitive computing, where systems understand the world in a way more similar to humans: through senses, learning, and experience. So it's natural for IBM to support the prize and be part of this four-year journey. The myriad problem solving approaches they bring will not only help IBM develop a better understanding of the use cases, but also help align their strategy and marketing and thought leadership with actual practical examples deployed by inventive teams.

I imagine the open nature of the competition could complicate the judging aspect.

It is a challenge, because in a sense, we'll be comparing apples and oranges. But it is important to understand the construction of this prize. It's not to reward innovative technology or sim-

ply humanitarian aspects or ideas. It's about rewarding very clearly how individuals and teams can develop applications of AI to solve a problem, and to implement those solutions in a way that showcases human-machine collaboration. So based on that a number of criteria emerge. Of course they will be judged on their capabilities in AI; they have to come up with novel technologies and deploy applications that make sense. But we also recognize that there are going to be challenges in judging them objectively when we compare, let's say, a cancer

Combining human capabilities with machine capabilities provides better insight altogether—it's not a 'them or us' proposition."

treatment or early detection with water security. Which one is more important? If there is no clean fresh water, then obviously you're going to have a multiplication of diseases and poverty, so we really cannot compare them. But what we will be comparing and judging is the ability of a team to come up with a brilliant application of AI and then demonstrate that they can make an impact. The judges are actually independent from XPRIZE. If they decide that two teams are very close, they may decide that they will share the prize equally—I don't know. There are so many criteria to be taken into account that we don't think the challenge will come from judges evaluating which one will have more impact. The difficulty will be ensuring that the teams are building applications of AI that could have practical applications in a not too distant timeframe.

How do teams qualify for the competition?

The participants have to create a team and they have to be incorporated. Besides that, they

have to demonstrate that they are capable of developing AI solutions in applications that the judges approve as a real challenge. We don't want them working on issues that are so far distant that they may not be achievable. We want them to be audacious, but also realizable. The other challenges for teams involve resources. Can they sustain this effort for the duration of the prize? Can they gain access to expertise, mentorship, and data? Data, of course, is the feedstock of AI. Without data, AI is not going to go very far.

How will you interact with the teams during the course of the prize?

We provide a layer of support to all the teams. We have formed a number of partnerships around XPRIZE, as well as leadership through a number of organizations and institutions that are providing know-how, advice, data sources, and other resources to the teams. We don't get involved directly, but we do invite stakeholders to participate, including those from government. The UN is one of them. There are also NGOs, venture capitalists, and many technology companies that are participating and providing support, which eases some of the challenges. The biggest challenge, though, is simply competing and sustaining the effort for the duration.

How will teams be eliminated as the competition progresses?

One very important aspect of this competition is that not only is it open, but we allow the teams to define the problems they are solving. We also ask them to define their own milestones. So think of it as playing golf against yourself. You have your own handicap and you have to win against yourself. If the milestones are met and the judges approve, then they move forward to the next round. If they don't meet their milestones, then they don't move forward to the next round. This is true for the first two rounds of the competition. So we give them flexibility to define a timeline that they can afford, but at the same time is also ambi-

tious. Starting year three, judges will have more specific guidelines for them in terms of criteria. The AI has to be performant, they have to be doing knowledge sharing and transfer learning to other domains. For example, if they are developing a machine vision capability that is able to detect cancer, can it also detect conditions of interest under water? If so then the judges will credit them and give them additional criteria.

We started with 147 teams, half of which will be eliminated by mid 2018. But we will also have wildcard entrants. It's not to reward innovative technology or simply humanitarian aspects or ideas. It's about rewarding very clearly how individuals and teams can develop applications of AI to solve a problem, and to implement those solutions in a way that showcases human-machine collaboration." Because AI is moving at such a fast pace, we have designed into the prize the opportunity for a team to rejoin the competition if they believe they have what it takes and they have a breakthrough sufficient to reenter the competition one year later and still be able to win. New results will be uncovered, new discoveries will be made, and new opportunities will be created through the life of the prize. Again, the judges will decide. So in January we may have 20, 30, 40, or more new teams that will join the competition, which is very dynamic in nature.

We believe that solutions can come from anyone, anywhere, and that some of the greatest minds of our time remain untapped, ready to be engaged by a world that is in desperate need of help."

What do you believe will be the consequences of this Prize? What will be the ultimate outcomes? And what about those who don't win? How will they be impacted by virtue of having participated?

We believe that this prize will ignite teams to develop and showcase applications of AI to solve challenges, but it will also serve as a model and an opportunity for other similar endeavors, and also contribute to the creation of best practices that hopefully will generate common datasets, shared knowledge, and through the participation of cross-discipline teams to solve problems. Problems can be solved by engineers only, or by policy makers or by economists, but we're interested in how we can put multiple stakeholders and multiple discipline teams together. In that sense, this prize will be a showcase of possibilities. And that's already happening and providing results. We think this is one of the important consequences of it.

Also, participating in this prize will give teams a lot of visibility, as well as opportunities to share their goals and vision. And if they're not already supported, they'll get the support. So the dynamics created by this prize and the conversations around it will hopefully generate more understanding of what AI is capable of, and hopefully advance us from the polarized views, either positive or negative, and yield a much better understanding of what AI can actually do. And while we're cautiously optimistic about the possibilities of AI, we still understand the potential consequences and the limitations.

Lastly, I'd like to highlight is that participation is key. Whether individuals or organizations participate as a team or join as a mentor or supporter or expert or judge, the topic is so important that it will create many conversations and create many opportunities to have a better understanding. So my main message is that participation is important in any way, shape, or form. We believe that solutions can come from anyone, anywhere, and that some of the greatest minds of our time remain untapped, ready to be engaged by a world that is in desperate need of help.

Al for Good

The IBM Watson AI XPRIZE aims to accelerate adoption of AI technologies and spark creative, innovative, and audacious demonstrations of the technology that are truly scal-

able and solve societal grand challenges.

On August 31, 1955, a small group of visionaries gathered at Dartmouth to consider the radical idea of what would come to be called



-Peter Thiel

artificial intelligence. Their **proposed study** would "... proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems

now reserved for humans, and improve themselves."

No doubt these men were inspired by Alan Turing's seminal 1950 paper, "Computing Machinery and Intelligence," which posed the provocative question, "Can machines think?" It would take another 60 years, though, for the trifecta of enabling technologies—algorithms, data, and compute—to mature before dreams of endowing machines with human-like intelligence could come to fruition.

Today, with the advent of AI, we find ourselves upon a new hinge of history. And like all forms of change, this one also comes prepackaged with an ar-



1956: Marvin Minsky with Claude Shannon, Ray Solomonoff and other scientists attending the Dartmouth Summer Research Project on Artificial Intelligence. The 1956 summer workshop, now considered by many, to be the seminal event for artificial intelligence as a field. (Margaret Minsky)

ray of anxieties. Al systems have garnered sensational headline news about their ability to outperform humans when facing off, for example, in games of chess, Jeopardy, Atari, and Go. The stories go beyond mere gamesmanship, though, when they also beg the question, "In



what other ways can AI beat humans?" It's a question that has led to alarm over everything from jobs to civil unrest to myriad other disruptions that collectively paint AI as a dark force that is destined to spin out of control.

From where we stand today, we don't yet know the degree to which such fears will be seen as preposterous or prescient. Indeed, all potential outcomes are possible. And we enter this un-

charted territory without the benefit of historical guidance. We can't look to the effects of the Industrial Revolution as prologue, because the nature, scale, and reach of technology today is exponentially beyond anything imagined by the industrial innovators of the late Victorian era—expect perhaps, by H.G. Wells.

As Ray Kurzweil noted of more elemental forces, "The promise and peril are deeply intertwined. Fire kept us warm and cooked our food and also burned down our houses." Kurzweil also points out that there are strategies to control the peril, ". . . as there have been with biotechnology guidelines."

In one sense, AI, like any other technology, is, *in itself*, benign. But it's benign in the same way that nuclear energy, computers, rockets, or automobiles are benign: the potential for malignancy lies in its application, as well as the extent of its regulation. Nobody questions that AI can

be exploited for evil purposes—and no doubt, it will. Nothing in this world is safe from corruption. But while intrinsically benign it may be, as Elon Musk warned, ". . . with nuclear weapons and AI, we don't want to learn from our mistakes. We want to plan ahead."

Despite these astonishing advances, we are a long way from machines that are as intelligent as humans—or even rats. So far, we've seen only 5% of what AI can do."

—Yann LeCun



It's that very deliberate and proactive planning that is driving the intentionality of XPRIZE Foundation's AI initiatives, which can be summed up in the phrase, "Al for good." And the good is incalculable.

The recent US government report, Preparing for the Future of Artificial Intelligence, makes a compelling case for AI's potential to address society's greatest challenges. "Smart vehicles," it reads, "may save hundreds of thousands of lives every year worldwide, and increase mobility for the elderly and those with disabilities. Smart buildings may save energy and reduce carbon emissions. Precision medicine may extend life and increase quality of life. Smarter government

ey. Al-enhanced education may help teachers give every child an education that opens doors to a secure and fulfilling life. These are just a few of the potential benefits if the

Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years."

-Andrew Ng

technology is developed with an eye to its benefits and with careful consideration of its risks and challenges."

Indeed, AI, primarily in the form of deep learning systems, is being applied to tasks as varied as medical diagnostics, credit scoring, fraud detection, product recommendations, language translation, security, behavioral analysis, robotics. The questions being addressed range from where to plant crops to how to how to conserve water. There is no question that AI can and will impact every aspect of how we live, work, learn, discover, and communicate.

But back to the fundamental issue as to whether AI is a question of man *versus* machine or of man *and* machine. It turns out that the

Learning From Experience

Deep neural networks learn by adjusting the strengths of their connections to better convey input signals through multiple layers to neurons associated with the right general concepts.



When data is fed into a network, each artificial neuron that fires (labeled "1") transmits signals to certain neurons in the next layer, which are likely to fire if multiple signals are received. The process filters out noise and retains only the most relevant features.

Image courtesy of Lucy Reading-Ikkanda/Quanta Magazine.

word "question" is operative. Warren Berger, in his book, *A More Beautiful Question*, suggests that the uniquely human capacity to question is our ace in the hole when it comes to "competing" with AI. "Until Watson acquires the equivalent of human curiosity, creativity, divergent thinking skills, imagination, and judgment," he says, "it will not be able to formulate the kind of original, counterintuitive, and unpredictable questions an innovative thinker—or even just your average four-year-old—can come up with."

Peter Thiel, writing in his book, *From Zero to One*, parses the issue more pragmatically, amplifying the collaborative possibilities of AI. "Men and machines are good at fundamentally

The uniquely human capacity to question is our ace in the hole when it comes to "competing" with AI." different things. People have intentionality—we form plans and make decisions in complicated situations. We're less good at making sense of enormous amounts of data. Computers are exactly the opposite: they excel at efficient data processing, but they struggle to make basic judgments that would be simple for any human. Watson, Deep Blue, and ever-better machine learning algorithms are cool. But the most valuable companies in the future won't ask what problems can be solved with computers alone. Instead, they'll ask: *how can computers help humans solve hard problems?* . . . As we find new ways to use computers, they won't just get better at the kinds of things people already do; they'll help us to do what was previously unimaginable."

And to such ends, solutions are being developed and deployed across the spectrum of AI's "three A's": assisted, augmented, and autonomous, and with an accelerated sense of urgency instilled by the competitive dynamic of the IBM Watson AI XPRIZE—a \$5 million cognitive computing competition challenging teams to "develop and demonstrate how humans can collaborate with powerful AI technologies to tackle the world's grand challenges."

An Introduction to Artificial Intelligence

By Stuart Russell and Peter Norvig

We call ourselves *Homo sapiens*—man the wise—because our intelligence is so important to us. For thousands of years, we have tried to understand *how we think*; that is, how a mere

handful of matter can perceive, understand, predict, and manipulate a world far larger and more complicated than itself. The field of artificial intelligence, or AI, goes further still: it attempts not just to understand but also to *build* intelligent entities.

Al is one of the newest fields in science and engineering. Work started in earnest soon after World War II, and the name itself was coined in 1956. Along with molecular biology, Al is regularly cited as the "field I would most like to be in" by scientists in other disciplines.

Peter Norvig



A student in physics might reasonably feel that

all the good ideas have already been taken by Galileo, Newton, Einstein, and the rest. AI, on the other hand, still has openings for several full-time Einsteins and Edisons.

Al currently encompasses a huge variety of subfields, ranging from the general (learning and perception) to the specific, such as playing chess, proving mathematical theorems, writing poetry, driving a car on a crowded street, and diagnosing diseases. Al is relevant to any intellectual task; it is truly a universal field. Stewart Russell

THINKING HUMANLY

"The exciting new effort to make computers think ... machines with minds, in the full and literal sense." (Haugeland, 1985)

"[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning ..." (Bellman, 1978)

ACTING HUMANLY

"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)

"The study of how to make computers do things at which, at the moment, people are better." (Rich and Knight, 1991)

THINKING RATIONALLY

"The study of mental faculties through the use of computational models." (Charniak and McDermott, 1985)

"The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)

ACTING RATIONALLY

"Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

> "Al... is concerned with intelligent behavior in artifacts." (Nilsson, 1998)

We have claimed that AI is exciting, but we have not said what it is. In the figure above we see eight definitions of AI, laid out along two dimensions. The definitions on top are concerned with *thought processes* and *reasoning*, whereas the ones on the bottom address *behavior*. The definitions on the left measure success in terms of fidelity to *human* performance, whereas the ones on the right measure against an *ideal* performance measure, called rationality. A system is rational if it does the "right thing," given what it knows.

Historically, all four approaches to AI have been followed, each by different people with different methods. A human-centered approach must be in part an empirical science, involving observations and hypotheses about human behavior. A rationalist approach involves a combination of mathematics and engineering. The various groups have both disparaged and helped each other. Let us look at the four approaches in more detail.

Acting humanly: The Turing Test approach

The Turing Test, proposed by Alan Turing (1950), was designed to provide a satisfactory operational definition of intelligence. A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer. We note that programming a computer to pass a rigorously applied test provides plenty to work on. The computer would need to possess the following capabilities:

- **natural language** processing to enable it to communicate successfully in English;
- **knowledge representation** to store what it knows or hears;
- automated reasoning to use the stored information to answer questions and to draw new conclusions;
- **machine learning** to adapt to new circumstances and to detect and extrapolate patterns.

Turing's test deliberately avoided direct physical interaction between the interrogator and the computer, because *physical* simulation of a person is unnecessary for intelligence. However, the so-called **total Turing Test** includes a video signal so that the interrogator can test the subject's perceptual abilities, as well as the opportunity for the interrogator to pass physical objects "through the hatch." To pass the total Turing Test, the computer will need:

- **computer vision** to perceive objects, and
- **robotics** to manipulate objects and move about.

These six disciplines compose most of AI, and Turing deserves credit for designing a test that remains relevant 60 years later. Yet AI researchers have devoted little effort to passing the Turing Test, believing that it is more important to study the underlying principles of intelligence than to duplicate an exemplar. The quest for "artificial flight" succeeded when the Wright brothers and others stopped imitating birds and started using wind tunnels and learning about aerodynamics. Aeronautical engineering texts do not define the goal of their field as making "machines that fly so exactly like pigeons that they can fool even other pigeons."

Thinking humanly: The cognitive modeling approach

If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. We need to get *inside* the actual workings of human minds.

There are three ways to do this: through introspection—trying to catch our own thoughts as they go by; through psychological experiments—observing a person in action; and through brain imaging—observing the brain in action. Once we have a sufficiently precise theory of the mind, it becomes possible to express the theory as a computer program. If the program's input– output behavior matches corresponding human behavior, that is evidence that some of the program's mechanisms could also

If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. We need to get inside the actual workings of human minds." be operating in humans. For example, Allen Newell and Herbert Simon, who developed GPS, the "General Problem Solver" (Newell and Simon, 1961), were not content merely to have their program solve problems correctly. They were more concerned with comparing the trace of its reasoning steps to traces of human subjects solving the same problems. The interdisciplinary field of **cognitive science** brings together computer models from AI and experimental techniques from psychology to construct precise and testable theories of the human mind.

Cognitive science is a fascinating field in itself, worthy of several textbooks and at least one encyclopedia (Wilson and Keil, 1999). Real cognitive science, however, is necessarily based on experimental investigation of actual humans or animals. We will assume, however, that only a computer is available for experimentation.

In the early days of AI there was often confusion between the approaches: an author would argue that an algorithm performs well on a task and that it is therefore a good model of human performance, or vice versa. Modern authors separate the two kinds of claims; this distinction has allowed both AI and cognitive science to develop more rapidly. The two fields continue to fertilize each other, most notably in computer vision, which incorporates neurophysiological evidence into computational models.

Thinking rationally: The "laws of thought" approach

The Greek philosopher Aristotle was one of the first to attempt to codify "right thinking," that is, irrefutable reasoning processes. His **syllogisms** provided patterns for argument structures that always yielded correct conclusions when given correct premises—for example, "Socrates is a man; all men are mortal; therefore, Socrates is mortal." These laws of thought were supposed to govern the operation of the mind; their study initiated the field called logic. Logicians in the 19th century developed a precise notation for statements about all kinds of objects in the world and the relations among them. (Contrast this with ordinary arithmetic notation, which provides only for statements about numbers.) By 1965, programs existed that could, in principle, solve any solvable problem described in logical notation. (Although if no solution exists, the program might loop forever.) The so-called **logicist** tradition within artificial intelligence hopes to build on such programs to create intelligent systems.

There are two main obstacles to this approach. First, it is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain. Second, there is a big difference between solving a problem "in principle" and solving it in practice. Even problems with just a few hundred facts can exhaust the computational resources of any computer unless it has some guidance as to which reasoning steps to try first. Although both of these obstacles apply to any attempt to build computational reasoning systems, they appeared first in the logicist tradition.

Acting rationally: The rational agent approach

An agent is just something that acts (agent comes from the Latin *agere*, to do). Of course, all computer programs do something, but computer agents are expected to do more: operate autonomously, perceive their environment, persist over a prolonged time period, adapt to change, and create and pursue goals. A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.

In the "laws of thought" approach to AI, the emphasis was on correct inferences. Making correct inferences is sometimes part of being a rational agent, because one way to act rationally is to reason logically to the conclusion that a given action will achieve one's goals and then to act on that conclusion. On the other hand, correct inference is not all of rationality; in some situations, there is no provably correct thing to do, but something must still be done. There are also ways of acting rationally that cannot be said to involve inference. For example, recoiling from a hot stove is a reflex action that is usually more successful than a slower action taken after careful deliberation.

All the skills needed for the Turing Test also allow an agent to act rationally. Knowledge representation and reasoning enable agents to reach good decisions. We need to be able to generate comprehensible sentences in natural language to get by in a complex society. We need learning not only for erudition, but also because it improves our ability to generate effective behavior.

The rational-agent approach has two advantages over the other approaches. First, it is more

general than the "laws of thought" approach because correct inference is just one of several possible mechanisms for achieving rationality. Second, it is more amenable to scientific development than are approaches based on human behavior or human thought. The standard of rationality is mathematically well defined and completely general, and can be "unpacked" to generate agent designs that provably achieve it. Human behavior, on the other hand, is well

A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome."

adapted for one specific environment and is defined by, well, the sum total of all the things that humans do. We therefore concentrate on general principles of rational agents and on components for constructing them. We will see, then, that despite the apparent simplicity with which the problem can be stated, an enormous variety of issues come up when we try to solve it.

One important point to keep in mind: We will see before too long that achieving perfect rationality—always doing the right thing—is not feasible in complicated environments. The

computational demands are just too high. However, we will adopt the working hypothesis that perfect rationality is a good starting point for analysis. It simplifies the problem and



provides the appropriate setting for most of the foundational material in the field. Nonetheless, we must also deal with the issue of limited rationality—acting appropriately when there is not enough time to do all the computations one might like.

This essay is adapted from the Introduction to Artificial Intelligence: A Modern Approach, by Stuart Russell and Peter Norvig. Copyright © 2010, 2003, 1995 by Pearson Education, Inc.

Team Vignettes

The following sampling of team profiles highlights the broad diversity of AI applications the contestants are pursuing.



ARTERYS

THE BIG IDEA

We make it far easier to diagnose and track disease in medical images, without taking over the physician's role. Today, image interpretation is manual, tedious, and full of errors and inconsistencies. At the same time, radiologists are experiencing severe burnout and don't have enough time to devote to every image. Our intelligent platform for medical imaging brings speed, far more comprehensive measurements, and increased consistency to image interpretation, while allowing the radiologist to make the ultimate decision, and even edit the output of the different algorithms.

HOW WE'RE ACTUALIZING AI FOR GOOD

We see a fourfold benefit in our specific application of AI to healthcare:

Augmenting clinicians to eliminate large amounts of tedious work, improve disease measurement, and lower costs.

2 Making this technology accessible to hospitals and clinics around the world, without requiring them to invest in a supercomputer.

3 Improving consistency across centers and clinicians to improve diagnoses, leading to better outcomes and reduced waste.

4 Creating high-quality medical software that meets the highest medical device standards.

THE TECHNOLOGIES

Arterys has built a distributed cloud platform on GPUs to provide ultra-fast inference of our algorithms for our customers. Furthermore, this platform complies with patient data privacy requirements in the US, Canada, and all European countries. This platform allows us to collect data from around the world, and learn from each interaction with our users. Our algorithms are convolutional neural networks trained to detect, measure, and track anatomy.

ON HUMAN-AI COLLABORATION

Our platform was designed to augment the radiologist. Our unique, web browser-based user interface allows clinicians to interact with and edit the output of the AI algorithms. That way, clinicians can disagree with predictions or segmentations, edit them, and report their findings. Furthermore, we capture these corrections for future improvement of our algorithms, creating a virtuous cycle.

Learn more about Arterys >>

CHOITEK

THE BIG IDEA

We are not maximizing the full potential of our innovative capacity. In this modern age, there is an unquestioned gender gap in the STEM fields. According to the National Science Foundation, women represent half the total college educated workforce in the United States, but less than 25% of the entire science and engineering workforce, significantly diminishing the potential for technological growth and discovery. In particular, the number of women acquiring computer science degrees has been declining since 1984. This is an alarming trend. Women represent half the internet user base and form a significant portion of the customers of tech companies around the world. The STEM gender gap leaves out a significant amount of the perspectives and the talent needed to create the technology that runs human society as a whole. The world needs a way to effectively inspire and retain more women in STEM fields so that we as a species can build a technology-driven future for everyone. To these ends, we are researching the possibilities of using artificially intelligent social robots to inspire women to pursue careers in STEM fields.

HOW WE'RE ACTUALIZING AI FOR GOOD

Al for good means using Al to teach, inspire, and empower the next generation of artists, engineers, thinkers, and doers of all stripes. The Al-driven technology we're developing will work alongside educators to teach and inspire more young women to enter STEM fields as an artificially intelligent teaching aid that combats negative influences and socio-cultural stereotypes.

THE TECHNOLOGIES

Our Autonomous Support and Positive Inspiration Robot (ASPIR) is a social humanoid robot that comes equipped with an AI system called the Advanced Interactive Artificial Intelligence (AIAI), which is a visual node-based AI editor that allows students to easily create their own "social robot personalities" for the ASPIR robot. A social robot personality integrates several state-of-the-art AI techniques including machine learning, computer vision, voice recognition, and more in a form factor that even children who have little to no robotics or coding experience can understand.

ON HUMAN-AI COLLABORATION

We believe AI can be used as a wonderful instructional aid to multiply the influence and reach of teachers leveraging the best of both worlds: the automation that modern computer technology brings combined with the human touch of real students and teachers.

Learn more about Choitek >>



THE BIG IDEA

We are looking to design a robust AI system that can *dream*. The AI would take in a limited set of information to make predictions of the most probable—as well as the most beneficial—outcome. More importantly, in predicting the most beneficial outcome, the AI would help users understand how to get to that more beneficial outcome. One application is to apply it to ensemble forecasting methods for both severe weather phenomenon as well as for seasonal forecasting. We could also apply it to myriad challenging problems facing our society today.

HOW WE'RE ACTUALIZING AI FOR GOOD

Much of AI work today is focused on commercial optimization: the ability to better bring products or services to the customer—to get that extra .1% of efficiency or accuracy. It's no wonder that big companies like Google, Microsoft, and Facebook are investing heavily in this area. We are more interested in changing the "A" in AI to represent Altruistic. AI = Altruistic Intelligence. There are outcomes that are the most beneficial for all parties involved, but they are often low-probability "long shots." By understanding what is needed for us to get to those long shots, we can tip the scales and influence our systems toward that greater good. We want to utilize AI to show us how to enable that dream that we all have—to benefit all of society, not just a select few.

THE TECHNOLOGIES

Many of the implementations today are black boxes. Consequently, we have difficulty explaining how a neural net can recognize a stop sign and also why, by placing a small sticker on the stop sign, the AI can no longer recognize it. We have a concept currently in patent pending that involves a new way of designing the AI so we not only have insight into how it recognizes objects, but are also inherently creative in the process.

ON HUMAN-AI COLLABORATION

Collaboration is an integral part of the next stage of our development of Altruistic Intelligence. We expect that humans will be able to selectively provide the building blocks necessary to enable the AI to extrapolate both the most likely outcomes and the most beneficial, and provide us a realistic path toward achieving the AI-Human dream.

Learn more about Aerospace >>



THE BIG IDEA

Creativity is considered one of the most challenging and unpredictable capabilities to capture; it involves constructive imagination, dreaming, and abstraction—and all in conjunction with the combination of knowledge and experience. Despite the fact that streaming services drive digital revenues in the music industry now, there remain grand challenges with respect to creativity and productivity. Because the music industry is just that—an industry—the challenges of creativity and productivity go hand in hand. Originality is another matter. Our objective in the context of the XPRIZE is to find a novel Al solution for stimulating creativity in the music field. Enhanced productivity comes as a byproduct to address the time resources required by creativity.

HOW WE'RE ACTUALIZING AI FOR GOOD

In the next era of hyper-intuitive simulation that we propose and forecast, AI for good can extend a means of transparency and fairness when it comes to business dealings. AI could be a collaborative simulation tool to boost peoples' creativity more than ever before, revolutionizing the entire creative content generation space, e.g., movie making, music production, and other fine arts.

THE TECHNOLOGIES

We developed and implemented a novel methodological framework that utilizes evolutionary conceptual blending, machine learning, data visualization, and argumentation for human-computer intuitive interaction into a single Digital Audio Workstation (DAV). These methodologies are combined in a scientifically innovative way. We imagine that if a composer had a creative tool that could stimulate his or her creativity by allowing the interactive exploration of genuinely novel musical ideas, then the problems discussed earlier would have been substantially diminished. Both the interaction and the algorithmic process under the hood will make the system a paradigm of intuitive computing applied in music making, while its adaptive nature will constitute the system a truly personalized music maker.

By using our intuitive interaction that combines data visualization techniques and augmentation-based reasoning technologies, the intuition around musical concepts that the user may have will be similar to the "intuition" of the system, while those "intuitions" will continuously converge as the composer uses the system.

ON HUMAN-AI COLLABORATION

Our concept of human-AI collaboration will be demonstrated through the intuitive, interactive scheme we design. Through the transparent and intuitive interaction we're creating, the system can generate high-conceptual novel music content that reflects the user's choices, and in the process propose a plethora of creative solutions which the user, based on the history of communications with the system, would be likely to accept.

Learn more about Mercury Orbit Music >>



THE BIG IDEA

An important problem for humans is understanding complex dynamical systems from incomplete and imperfectly reliable information, and then making life-critical decisions upon that information. Such dynamical systems include a human body suffering from chronic disease, such as diabetes or cancer; or an aviation turbine engine, displaying unexpected behavior on the wing of an airborne aircraft flying at 10,000 feet. A logical and important interface is then where AI can be used to provide high-confidence insights to assist human executive decision-making for these life-critical tasks, where unassisted human decision-making fails. Important real-world examples of this arise where the existing data to be used for AI training do not contain within them examples of any existing "successful" strategy from which to learn. Furthermore, for real-world implementation the machine logic underpinning AI advice needs to be auditable for safety purposes; in the European Union (EU) this has recently become a controversial legal issue under the EU's new "right to explanation" requirement. The machine logic underpinning the current AI global "gold standard," deep learning neural networks (NN), is notoriously opaque and *prima facie* appears to fail this requirement. Hence there is a major problem within current AI technologies themselves.

HOW WE'RE ACTUALIZING AI FOR GOOD

To us, AI for Good means the use of AI to enhance and transform human interactions with the world around us, and to enhance our understanding of our own cognitive faculties by studying the alternatives. In practical terms it is to assist humans in difficult decision-making, not to replace human decision-making. The two applications we are demonstrating for the AI XPRIZE are (1) the world's first genuinely machine-intelligent artificial pancreas for type-I diabetes to reduce adverse blood glucose events from insulin without reliance on glucagon, and (2) a machine-intelligent tool for evolutionary diagnostics of aviation turbine engines, to deliver unprecedented detail on engine dynamics non-invasively from sensor data, to enhance safety and reduce operating costs and risk.

THE TECHNOLOGIES

We are demonstrating a completely different, new form of AI (technical name: phi-Textured Evolutionary Algorithms or phi-TEA; public name "MachineGenes" in all applications except medicine; "Neuromathix" in medicine). It is able to reconstruct complex dynamical systems from noise-polluted partial information from sensor time-series data, provided the system dynamics has known or guessable mathematical structure underlying its behavior. MachineGenes does this by literally evolving candidate mathematical structures for the system being studied and then comparing predicted behavior with observed data. Evolution rewards good solutions and kills off bad ones. This sounds a lot like Genetic Algorithms (GA)—a well-known approach whereby the specifications of potential solutions to a mathematical problem are encoded as "genes" in "chromosomes," which engage in reproductive breeding and undergo mutation across generations. Chromosomes that represent "good" solutions are rewarded by survival; those for "bad" solutions become extinct. Over a large number of generations, good or excellent solutions literally evolve. Although GA is a potentially powerful algorithm, it has a major problem: a "glass ceiling" in its performance. If the system being studied is both sufficiently complex

and sufficiently underdetermined from available data, then conventional GA fail to achieve improvements past a certain point and stagnate. To us this made no sense: the reason GA were developed in the first place was because evolution in nature achieves sophisticated optimization of extremely complex systems using partial information, so if GA stagnate there's something wrong. Our solution was that GA get the metaphor wrong by being incomplete. Evolution isn't just at the genetic level: for complex organisms such as vertebrates, herd behavior such as mating-based expulsions and hierarchy around waterholes, and secluded survival of more primitive related species such as archaic hominids in Europe, are all part of the natural algorithm which occurs at multiple layers simultaneously. So too are predator-prey dynamics. Physical landscape too is not simply something in which the species evolves to survive; it is actually part of the evolutionary algorithm, hence the biodiversity of the valleys of Papua New Guinea. Hence phi-TEA incorporates the algorithms of GA, but adds layers of additional behavior.

Everyone in the AI field it seems is obsessing with neural networks as the ultimate computational substrate for AI, because it's based on the metaphor of the human brain. We say that the human brain is not the ultimate computational substrate on Earth. The ecosystems of Earth are the ultimate computational substrate on Earth.

And the really interesting thing is, phi-TEA works. More than that: it is able to be used as the building-blocks for more sophisticated AI using differential game theory, which does not require examples of previously-successful strategies in order to design new ones. And due to the way it uses mathematical structure, its decisions are auditable, so it can comply with EU regulations. Furthermore, the use of explicit algebraic structures means phi-TEA requires far less data than NN: for example, we have been able to reconstruct personalized dynamics for type-1 diabetes using 72 hours of individual medical data, instead of months of data.

So much for the algorithms. Our hardware is massively-parallel computing architectures. We are experimenting with Nvidia Kepler cards and Intel Xeon Phi co-processor cards, although the current obsession of hardware manufacturers in optimizing everything for NN is frustrating.

Our core value proposition is the ability for this new form of AI to provide high-confidence auditable analysis of particular life-critical complex systems to human operators, plus suggested strategies to deal with the situation at hand, based on interpretations of the current situation. It is a future-facing form of AI using algebra, geometry, and game theory to perform predictive analysis, as distinct from NN, which are predominantly past-facing forms of AI inasmuch as their advice is predicated on weighting numerical patterns within historical data.

ON HUMAN-AI COLLABORATION

Al is good at bringing huge computational resources to bear to solve a specific problem in ways humans cannot, hence AI can bring a new dimension to analytical capabilities. It is not, however, particularly good at wider issues of context and interdisciplinary thinking. Interdisciplinary thinking includes, for example, combining the technical task of driving a car and balancing the ethical and empathetic question of avoiding humans and animals, because life is precious, versus the question of avoiding potholes, because they ruin the car's suspension. Yes, one can explicitly program an autonomous car to avoid humans and potholes and prefer avoiding humans to avoiding potholes, but when the car implements this programming it is following prescription, not engaging intelligence. A major challenge is imbuing a machine with actual intelligence to deduce priorities among objectives when thrown into situations far outside its training data, rather than following a cookbook of instructions. Human-AI collaboration makes sense, not only because we're human-and hence biased in our own favor—but because (non-sociopathic) humans see more dimensions of meaning and nuance to critical decisions than AI do, and this faculty is most important when the decision-maker is confronted with unexpected circumstance. Empathy and compassion have value and an internal logic of their own that is difficult to program. Perhaps the future of humanity is to be the instinctive, intuitive, and empathetic parts of a larger, distributed "brain"— the cognitive collaboration between humanity and AI.

Learn more about MachineGenes >>

About the Technica Curiosa Tech Entrepreneurship Ecosystem

TechnícaCuríosa

t Technica Curiosa we are dedicated to inspiring and equipping technology entrepreneurs who are ushering in an exciting, world-changing future. From ideation to start-up, to achieving scale and beyond, we are helping entrepreneurs make new and nonobvious connections across emerging technol-

ogies. Along the way, we help them identify trend trajectories and opportunities for lateral developments that culminate in truly disruptive innovations. To these ends, we've architected our iconic titles to cross-pollinate big ideas in interesting and exciting ways that work to empower and stimulate creativity, embolden imagination, and foster mindsets of possibility. In doing so, we expand the field of view—and include you in it.

Our fast-growing ecosystem comprises numerous stakeholder companies and organizations who share our passion for enabling and supporting vital new startup initiatives. Additionally, our platform's thought-leading contributors span tech celebrities to best-selling authors to industry professionals to academia, as well as researchers from Stanford, Caltech, Columbia, MIT, as well as NASA, DAR-

(Eneurshi)

Techn

PA, and the National Labs—all of whom are addressing the important technologies that are revolutionizing industries spanning energy to communications, artificial intelligence to the Internet of Things, genetics to robotics, the commercialization of space to self-driving cars.



To join us, contact John Schroeter at john@technicacuriosa.com

Copyright © 2017, John August Media, LLC. Technica Curiosa, Popular Electronics, Mechanix Illustrated, and Popular Astronomy are trademarks of John August Media, LLC. All rights reserved.

The FULUCE ain't what it SHU to be... at HH MH t's so much

TechnicaCuriosa is the new and exciting hub of a highly connected family of iconic media brands—brands endowed with rich legacies of world-changing, decades-spanning influence.

Because new discoveries and innovation come from connecting non-obvious and disparate dots, we've architected our titles to cross-pollinate big ideas in fascinating ways. In doing so, we expand the field of view—and include you in it.









