

AI Ask for Human Help

This is a demonstration of a situation where an AI algorithm working together with a human can reap the benefits and efficiency of the AI's good decisions, without being locked into its bad ones. [33]

The deep learning analysis has revealed that the extinct hominid is probably a descendant of the Neanderthal and Denisovan populations. [32]

Our team at IBM Research – India collaborated with the IBM MetroPulse team to bring such first-of-a-kind, AI-driven capabilities to MetroPulse, an industry platform that brings together voluminous market, external and client datasets. [31]

For people with hearing loss, it can very difficult to understand and separate voices in noisy environments. This problem may soon be history thanks to a new groundbreaking algorithm that is designed to recognise and separate voices efficiently in unknown sound environments. [30]

While researchers have taken steps to [HYPERLINK "https://doi.org/10.1017/S0140525X00023992"](https://doi.org/10.1017/S0140525X00023992) comprehensively catalogue the preferences of men and women, we still don't know which traits are the most important contributors to a person's attractiveness. [29]

A group of researchers from MIT have already developed an AI robot that can assist in a labour room. [28]

Researchers at Fukuoka University, in Japan, have recently proposed a design methodology for configurable approximate arithmetic circuits. [27]

Researchers at Google have recently developed a new technique for synthesizing a motion blurred image, using a pair of un-blurred images captured in succession. [26]

Constructing a neural network model for each new dataset is the ultimate nightmare for every data scientist. [25]

Algorithmic fairness is increasingly important because as more decisions of greater importance are made by computer programs, the potential for harm grows. [24]

Intel's Gadi Singer believes his most important challenge is his latest: using artificial intelligence (AI) to reshape scientific exploration. [23]

Artificial intelligence is astonishing in its potential. It will be more transformative than the PC and the Internet. Already it is poised to solve some of our biggest challenges. [22]

In the search for extraterrestrial intelligence (SETI), we've often looked for signs of intelligence, technology and communication that are similar to our own. [21]

Call it an a-MAZE-ing development: A U.K.-based team of researchers has developed an artificial intelligence program that can learn to take shortcuts through a labyrinth to reach its goal. In the process, the program developed structures akin to those in the human brain. [20]

And as will be presented today at the 25th annual meeting of the Cognitive Neuroscience networks to enhance their understanding of one of the most elusive intelligence systems, the human brain. [19]

U.S. Army Research Laboratory scientists have discovered a way to leverage emerging brain-like computer architectures for an age-old number-theoretic problem known as integer factorization. [18]

Now researchers at the Department of Energy's Lawrence Berkeley National Laboratory (Berkeley Lab) and UC Berkeley have come up with a novel machine learning method that enables scientists to derive insights from systems of previously intractable complexity in record time. [17]

Quantum computers can be made to utilize effects such as quantum coherence and entanglement to accelerate machine learning. [16]

Neural networks learn how to carry out certain tasks by analyzing large amounts of data displayed to them. [15]

*Who is the better experimentalist, a human or a robot? When it comes to exploring synthetic and crystallization conditions for inorganic gigantic molecules, actively learning machines are clearly ahead, as demonstrated by British Scientists in an experiment with polyoxometalates published in the journal *Angewandte Chemie*. [14]*

Machine learning algorithms are designed to improve as they encounter more data, making them a versatile technology for understanding large sets of photos such as those accessible from Google Images. Elizabeth Holm, professor of materials science and engineering at Carnegie Mellon University, is leveraging this technology to better understand the enormous number of research images accumulated in the field of materials science. [13]

With the help of artificial intelligence, chemists from the University of Basel in Switzerland have computed the characteristics of about two million crystals made up of four chemical elements. The researchers were able to identify 90 previously unknown thermodynamically stable crystals that can be regarded as new materials. [12]

The artificial intelligence system's ability to set itself up quickly every morning and compensate for any overnight fluctuations would make this fragile technology much more useful for field measurements, said co-lead researcher Dr Michael Hush from UNSW ADFA. [11]

Quantum physicist Mario Krenn and his colleagues in the group of Anton Zeilinger from the Faculty of Physics at the University of Vienna and the Austrian Academy of Sciences have developed an algorithm which designs new useful quantum experiments. As the computer does not rely on human intuition, it finds novel unfamiliar solutions. [10]

Researchers at the University of Chicago's Institute for Molecular Engineering and the University of Konstanz have demonstrated the ability to generate a quantum logic operation, or rotation of the qubit, that - surprisingly—is intrinsically resilient to noise as well as to variations in the strength or duration of the control. Their achievement is based on a geometric concept known as the Berry phase and is implemented through entirely optical means within a single electronic spin in diamond. [9]

New research demonstrates that particles at the quantum level can in fact be seen as behaving something like billiard balls rolling along a table, and not merely as the probabilistic smears that the standard interpretation of quantum mechanics suggests. But there's a catch - the tracks the particles follow do not always behave as one would expect from "realistic" trajectories, but often in a fashion that has been termed "surrealistic." [8]

Quantum entanglement—which occurs when two or more particles are correlated in such a way that they can influence each other even across large distances—is not an all-or-nothing phenomenon, but occurs in various degrees. The more a quantum state is entangled with its partner, the better the states will perform in quantum information applications. Unfortunately, quantifying entanglement is a difficult process involving complex optimization problems that give even physicists headaches. [7]

A trio of physicists in Europe has come up with an idea that they believe would allow a person to actually witness entanglement. Valentina Caprara Vivoli, with the University of Geneva, Pavel Sekatski, with the University of Innsbruck and Nicolas Sangouard, with the University of Basel, have together written a paper describing a scenario where a human subject would be able to witness an instance of entanglement—they have uploaded it to the arXiv server for review by others. [6]

The accelerating electrons explain not only the Maxwell Equations and the Special Relativity, but the Heisenberg Uncertainty Relation, the Wave-Particle Duality and the electron's spin also, building the Bridge between the Classical and Quantum Theories.

The Planck Distribution Law of the electromagnetic oscillators explains the electron/proton mass rate and the Weak and Strong Interactions by the diffraction

patterns. The Weak Interaction changes the diffraction patterns by moving the electric charge from one side to the other side of the diffraction pattern, which violates the CP and Time reversal symmetry.

The diffraction patterns and the locality of the self-maintaining electromagnetic potential explains also the Quantum Entanglement, giving it as a natural part of the relativistic quantum theory.

Contents

Preface	7
Artificial intelligence must know when to ask for human help	7
How sure is the algorithm?	7
Keeping a human eye on AI	8
Artificial intelligence applied to the genome identifies an unknown human ancestor	9
Humans had descendants with a species that is unknown to us	10
Deep learning: deciphering the keys to human evolution in ancient DNA	10
An extinct hominid could explain the history of humankind.....	10
AI can help retailers understand the consumer	11
The data layer, which consists of.....	12
The knowledge layer, which consists of	12
The industry intelligence layer, which consists of.....	12
Artificial neural networks make life easier for hearing aid users	13
Noise reduction and speech separation	14
Deep neural networks	14
A computer behind the ear	15
The cocktail party phenomenon.....	15
Why we want to build a machine that can predict a person's attractiveness	15
Could AI take control of human birth?.....	17
Fact or fiction?	18
Emotional support.....	18
A new approximate computing approach using CNNs	19
A new technique for synthesizing motion-blurred images	21
Predicting the accuracy of a neural network prior to training	22
Are computer-aided decisions actually fair?	23
The problem of self-fulfilling predictions	24
The real-world impact of algorithmic bias	24

Overcoming algorithmic bias	25
How is artificial intelligence changing science?	25
Q. How is AI changing science?	26
Q. What's an example?	26
Q. So you know what you're looking for, but you don't know how to find it?	26
Q. Are there other ways AI can change the scientific approach?	26
Q. I'm guessing you could retire if you wanted to. What keeps you going now?	27
The U.S. needs a national strategy on artificial intelligence	27
Can artificial intelligence help find alien intelligence?	29
'Decoding' intelligence	29
Thinking differently	30
Smarter than slime mould?	31
First contact	32
Interdisciplinary futures	33
Scientists make a maze-running artificial intelligence program that learns to take shortcuts	33
Dissecting artificial intelligence to better understand the human brain	36
Army's brain-like computers moving closer to cracking codes	37
Teaching computers to guide science: Machine learning method sees forests and trees....	39
Rise of the quantum thinking machines	41
A Machine Learning Systems That Called Neural Networks Perform Tasks by Analyzing Huge Volumes of Data	42
Active machine learning for the discovery and crystallization of gigantic polyoxometalate molecules	43
Using machine learning to understand materials	44
Artificial intelligence helps in the discovery of new materials	45
Machine learning aids statistical analysis	45
Unknown materials with interesting characteristics	45
Physicists are putting themselves out of a job, using artificial intelligence to run a complex experiment	46
Quantum experiments designed by machines	47
Moving electrons around loops with light: A quantum device based on geometry	47
Quantum geometry	48
A light touch	48
A noisy path	48

Researchers demonstrate 'quantum surrealism'	49
Physicists discover easy way to measure entanglement—on a sphere	50
An idea for allowing the human eye to observe an instance of entanglement.....	52
Quantum entanglement.....	52
The Bridge	53
Accelerating charges	53
Relativistic effect	53
Heisenberg Uncertainty Relation	53
Wave – Particle Duality	54
Atomic model.....	54
The Relativistic Bridge	54
The weak interaction	54
The General Weak Interaction.....	56
Fermions and Bosons	56
Van Der Waals force	56
Electromagnetic inertia and mass.....	56
Electromagnetic Induction	56
Relativistic change of mass	56
The frequency dependence of mass	57
Electron – Proton mass rate	57
Gravity from the point of view of quantum physics	57
The Gravitational force.....	57
The Higgs boson	58
Higgs mechanism and Quantum Gravity	58
What is the Spin?.....	59
The Graviton	59
The Secret of Quantum Entanglement	59
Conclusions.....	60
References	60

Author: George Rajna

Preface

Physicists are continually looking for ways to unify the theory of relativity, which describes largescale phenomena, with quantum theory, which describes small-scale phenomena. In a new proposed experiment in this area, two toaster-sized "nanosatellites" carrying entangled condensates orbit around the Earth, until one of them moves to a different orbit with different gravitational field strength. As a result of the change in gravity, the entanglement between the condensates is predicted to degrade by up to 20%. Experimentally testing the proposal may be possible in the near future. [5]

Quantum entanglement is a physical phenomenon that occurs when pairs or groups of particles are generated or interact in ways such that the quantum state of each particle cannot be described independently – instead, a quantum state may be given for the system as a whole. [4]

I think that we have a simple bridge between the classical and quantum mechanics by understanding the Heisenberg Uncertainty Relations. It makes clear that the particles are not point like but have a dx and dp uncertainty.

Artificial intelligence must know when to ask for human help

Artificial intelligence systems are powerful tools for businesses and governments to process data and respond to changing situations, whether on the stock market or on a battlefield. But there are still some things AI isn't ready for.

We are scholars of computer science working to understand and improve the ways in which algorithms interact with society. AI systems perform best when the goal is clear and there is high-quality data, like when they are asked to distinguish between different faces after learning from many pictures of correctly identified people.

Sometimes AI systems do so well that users and observers are surprised at how perceptive the technology is. However, sometimes success is difficult to measure or defined incorrectly, or the training data does not match the task at hand. In these cases, AI algorithms tend to fail in unpredictable and spectacular ways, though it's not always immediately obvious that something has even gone wrong. As a result, it's important to be wary of the hype and excitement about what AI can do, and not assume the solution it finds is always correct.

When algorithms are at work, there should be a human safety net to prevent harming people. Our research demonstrated that in some situations algorithms can recognize problems in how they're operating, and ask for human help. Specifically, we show, asking for human help can help alleviate algorithmic bias in some settings.

How sure is the algorithm?

Artificial intelligence systems are being used in criminal sentencing, facial-based personality profiling, resume screening, health care enrollment and other difficult tasks where people's lives and well-being are at stake. U.S. government agencies are beginning to ramp up their exploration and use of AI systems, in response to a recent executive order from President Donald Trump.

It's important to remember, though, that AI can cement misconceptions in how a task is addressed, or magnify existing inequalities. This can happen even when no one told the [algorithm](#) explicitly to treat anyone differently.

For instance, many companies have algorithms that try to determine features about a person by their face – say to guess their gender. The systems developed by U.S. companies tend to do significantly [better at categorizing white men](#) than they do women and darker-skinned people; they do worst at dark-skinned women. Systems developed in China, however, tend to [do worse on white faces](#).

The difference is not because one group has faces that are easier to classify than others. Rather, both algorithms are typically trained on a large collection of data that's not as diverse as the overall human population. If the data set is dominated by a particular type of face – white men in the U.S., and Chinese faces in China – then the algorithm will probably do better at analyzing those faces than others.



Biased training data can make systems better, or worse, at recognizing certain kinds of faces.

Credit: [Andrey Popov/Shutterstock.com](#)

No matter how the difference arises, the result is that algorithms can be biased by being more accurate on one group than on another.

Keeping a human eye on AI

For high-stakes situations, the algorithm's confidence in its own result – its estimation of how likely it is that the system came up with the right answer – is just as important as the result itself. The people who receive the output from algorithms need to know how seriously to take the results, rather than assuming that it's correct because it involved a computer.

Only recently have researchers begun to develop ways to identify, much less attempt to fix, [inequalities in algorithms and data](#). Algorithms can be programmed to recognize their own shortcomings – and follow that recognition with a [request for a person to assist with the task](#).

Many types of AI algorithms already calculate an internal [confidence level](#) – a prediction of how well it did at analyzing a particular piece of input. In facial analysis, many AI algorithms [have lower confidence](#) on darker faces and female faces than for white male faces. [It's unclear](#) how much this has been taken into account by law enforcement for high-stakes uses of these algorithms.

The goal is for the AI itself to locate the areas where it is not reaching the same accuracy for different groups. On these inputs, the AI can defer its decision to a human moderator. This technique is especially well-suited for context-heavy tasks like [content moderation](#).

Human content moderators [cannot keep up with](#) the flood of images being posted on social media sites. But AI content moderation is famous for failing to take into account the context behind a post – misidentifying discussions of sexual orientation as [explicit content](#), or identifying the Declaration of Independence as [hate speech](#). This can end up inaccurately censoring one [demographic](#) or [political](#) group over another.

To get the best of both worlds, [our research](#) suggests scoring all content in an automated fashion, using the same AI methods already common today. Then our approach uses newly proposed techniques to automatically locate potential inequalities in the accuracy of the algorithm on different protected groups of people, and to hand over the decisions about certain individuals to a human. As a result, the algorithm can be completely unbiased about those people on which it actually decides. And humans decide on those individuals where algorithmic decision would have inevitably created bias.

This approach does not eliminate bias: It just "concentrates" the potential for bias on a smaller set of decisions, which are then handled by people, using human common sense. The AI can still perform the bulk of the decision-making work.

This is a demonstration of a situation where an AI algorithm working together with a human can reap the benefits and efficiency of the AI's good decisions, without being locked into its bad ones. Humans will then have more time to work on the fuzzy, difficult decisions that are critical to ensuring fairness and equity. [33]

Artificial intelligence applied to the genome identifies an unknown human ancestor

By combining deep learning algorithms and statistical methods, investigators from the Institute of Evolutionary Biology (IBE), the Centro Nacional de Análisis Genómico (CNAG-CRG) of the Centre for Genomic Regulation (CRG) and the Institute of Genomics at the University of Tartu have identified, in the genome of Asiatic individuals, the footprint of a new hominid who cross bred with its ancestors tens of thousands of years ago.

Modern human DNA computational analysis suggests that the [extinct species](#) was a hybrid of Neanderthals and Denisovans and cross bred with Out of Africa modern humans in Asia. This finding

would explain that the hybrid found this summer in the caves of Denisova—the offspring of a Neanderthal mother and a Denisovan father—was not an isolated case, but rather was part of a more general introgression process.

The study, published in *Nature Communications*, uses [deep learning](#) for the first time ever to account for [human evolution](#), paving the way for the application of this technology in other questions in biology, genomics and evolution.

Humans had descendants with a species that is unknown to us

One of the ways of distinguishing between two species is that while both of them may cross breed, they do not generally produce fertile descendants. However, this concept is much more complex when extinct species are involved. In fact, the story told by current human DNA blurs the lines of these limits, preserving fragments of hominids from other species, such as the Neanderthals and the Denisovans, who coexisted with modern humans more than 40,000 years ago in Eurasia.

Now the study investigators have used [deep learning algorithms](#) to identify a new and hitherto-unknown ancestor of humans that would have interbred with modern humans tens of thousands of years ago. "About 80,000 years ago, the so-called Out of Africa occurred, when part of the human population, which already consisted of modern humans, abandoned the African continent and migrated to other continents, giving rise to all the current populations," explained Jaume Bertranpetit, principal investigator at the IBE and head of Department at the UPF. "We know that from that time onwards, [modern humans](#) cross-bred with Neanderthals in all the continents, except Africa, and with the Denisovans in Oceania and probably in South-East Asia, although the evidence of cross-breeding with a third extinct species had not been confirmed with any certainty".

Deep learning: deciphering the keys to human evolution in ancient DNA

Hitherto, the existence of the third ancestor was only a theory that would explain the origin of some fragments of the current human genome (part of the team involved in this study had already posed the existence of the extinct hominid in a previous study). However, deep learning has made it possible to make the transition from DNA to the demographics of ancestral populations.

The problem the investigators had to contend with is that the demographic models they have analysed are much more complex than anything else considered to date and there were no statistical tools available to analyse them. Deep learning "is an algorithm that imitates the way in which the nervous system of mammals works, with different artificial neurons that specialise and learn to detect, in data, patterns that are important for performing a given task", stated Òscar Lao, principal investigator at the CNAG-CRG and an expert in this type of simulation. "We have used this property to get the algorithm to learn to predict human demographics using genomes obtained through hundreds of thousands of simulations. Whenever we run a simulation we are traveling along a possible path in the history of humankind. Of all simulations, deep learning allows us to observe what makes the ancestral puzzle fit together."

An extinct hominid could explain the history of humankind

The deep learning analysis has revealed that the extinct hominid is probably a descendant of the Neanderthal and Denisovan populations. The discovery of a fossil with these characteristics this summer would seem to endorse the study finding, consolidating the hypothesis of this third species or population that coexisted with modern human beings and mated with them. "Our theory

coincides with the hybrid specimen discovered recently in Denisova, although as yet we cannot rule out other possibilities," said Mayukh Mondal, an investigator at the University of Tartu and former investigator at the IBE. [32]

AI can help retailers understand the consumer

Consumer brands and retailers often struggle to fully understand ever-changing customer needs. That is why you mostly find XL sizes in your favorite fashion store and no M sizes. That is why you have to spend hours looking for the style you saw on Instagram and still not find it. That is why the cost of dead inventory to fashion retailers in the US alone is estimated to be a whopping USD 50 billion. And that is part of the reason why the US generated 16 million tons of textile waste in 2014.

This is not because of any lack of intention or effort in the industry; rather, it is extremely difficult to understand consumers at scale. Characterizing consumers with broad brush definitions of age, gender and income is not effective given diverse and ever-changing consumer preferences, and retailers now need to look at much finer market segments—even down to single individuals. Increasingly, consumers are driving trends rather than merchants defining them, and this goes hand in hand with much more experimentation and disruption in the market.

To create and sell the "next big thing" in such a dynamic environment, designers, buyers and merchandisers must use their own creativity but also consider, with unprecedented granularity, how consumer preferences are changing and how different design, merchandising and marketing choices will perform. This is where AI and automation come in.

For example, consider a fashion retail buyer. She is responsible for the financial success of the merchandise she selects in any given season, but it's impossible for her to predict the performance of any design 12 months before the target season, or to identify the best promotional interventions to apply in-season. This is because she has very little visibility into how consumer preferences are changing across her stores over time, and how competing products are performing in the market.

Imagine an AI driven system that could analyze the natural language text of millions of online customer reviews and the images of all products in the market to summarize key relationships between location-specific customer sentiment and product features. For instance, how customers are responding to color block pullovers in Kansas City versus in Buffalo, and what attribute is the likely cause of lower customer sentiment for her color block pullovers versus competing color block pullovers. Such information for her in-market products and new planned products would help her dramatically improve her assortment, pricing and markdown, and marketing plans.

Market sentiment computed for visually similar floral tops across three different brands. The dashboard shows consumer preference for this type of product in different parts of the country, and recommended stock interventions at different stores.

Similarly, consider a sales manager for a yogurt brand. With a system that could analyze cross-brand sales of foods across the country to produce high quality predictions of the demand for the spinach artichoke flavored yogurt produced by the company, the sales manager then could negotiate

product introductions and planograms with retailers. A majority of such negotiations fail today in the absence of such capability.

In fact, IBM's recent study of more than 1,900 retail and consumer product leaders shows that the adoption of intelligent automation in the retail and consumer products industries is projected to balloon from 40 percent of companies today to more than 80 percent in three years.

Our team at IBM Research – India collaborated with the IBM MetroPulse team to bring such first-of-a-kind, AI-driven capabilities to MetroPulse, an industry platform that brings together voluminous market, external and client datasets. The new capabilities use AI and automation to fuse these structured and unstructured datasets around semantic, visual and location contexts and uncovers fine-grained insights about customer preferences hidden in this fused data. These insights will help consumer brands and retailers make smarter choices about product design, inventory planning, demand forecasting and product assortment that are in tune with dynamic consumer preferences.

The platform has three layers, each with deep industry content:

The data layer, which consists of

Market data containing up to date, location-specific signals of consumer preferences, product landscape and brand/retailer behavior. Some examples of datasets we curate include online customer reviews and comments, point of sales data and product images. The analysis of these large data sets can give companies clues as to how consumer preferences are changing across brands, retailers, cultures and regions – at the city or neighborhood level.

Hyper-local third-party data continuously updated, neighborhood-level signals of external factors that influence consumer behavior such as demographics, [weather forecasts](#) and history, local events and visitors.

Private enterprise data containing information about the retailer's own stores, products, merchandising, promotions and sales history. This data is treated with high security and privacy guarantees.

Incorporating such multiple data sets is critical to getting demand sensing and forecasting right, as also noted in Supply Chain Management 2018: In Service of The Customer, Retail Systems Research, Dec 2018 where 60 – 70 percent of respondents see "a lot of value" from considering new data such as sentiment, trade area data and past promotions into demand forecasting.

The knowledge layer, which consists of

A retail-industry specific digital vocabulary – knowledge graphs that capture industry information in the form of entities, properties and relationships. This layer allows the data in the data layer to be interpreted and analyzed in a standard and meaningful way by the AI systems that ultimately generate insights for end users. For example, consider a fashion taxonomy that captures various fashion terms and the relationships between them (e.g. "peplum" is type of "top"), or a grocery ontology that captures various food types, ingredients, flavors and type hierarchy.

The industry intelligence layer, which consists of

A variety of AI algorithms and models that can identify and understand the signals buried in the data, turning them into insights and recommendations exposed in a meaningful and standardized

way via a dashboard and APIs. These insights and recommendations help business leaders, product designers, merchandisers and other enterprise users to understand and optimize for the behaviors and preferences of their target consumer base. Shopper-facing solutions for personalization and cognitive assistance can also consume these APIs for enhanced consumer engagement at the point of sale. These algorithms leverage the latest AI techniques in multimodal AI, explainable AI and forecasting, and tune them to understand industry-specific knowledge and concepts.

Multimodal AI brings together visual perception and natural language processing to extract insights from multiple modalities of data. For instance, it may identify fashion objects in an image and relate them to expressions of customer opinion in accompanying reviews; or a semantic similarity model that understands that apple juice is more similar to lemonade than to apple for a consumer, but close to apple juice in terms of flavor and ingredients.

Explainable AI explains why a model generates a particular output for a given input. As AI models have gotten more complex, it has become nearly impossible for non-data scientists to understand their behavior, making it difficult for them to rely on the predictions from the models. Four techniques bring that explainability and help end users understand the "why." For example, understand the sentiment for different design aspects of summer dresses in the market and explain the contribution of local factors (demographics, weather) and merchandising factors (price, stock, promotion) on the variation of sentiment across the counties of the US.

You can try out these new MetroPulse capabilities with real-world data at the [National Retail Federation \(NRF\)](#) event in New York City in January, 2019, or see [here](#) for further details. [31]

Artificial neural networks make life easier for hearing aid users

For people with hearing loss, it can be very difficult to understand and separate voices in noisy environments. This problem may soon be history thanks to a new groundbreaking algorithm that is designed to recognize and separate voices efficiently in unknown sound environments.

People with normal hearing are usually able to understand each other without effort when communicating in noisy environments. However, for people with hearing loss, it is very challenging to understand and separate voices in noisy environments, and a hearing aid may really help. But there's still some way to go when it comes to general sound processing in hearing aids, explains Morten Kolbæk:

"When the scenario is known in advance, as in certain clinical test setups, existing algorithms can already beat human performance when it comes to recognizing and distinguishing speakers. However, in normal listening situations without any prior knowledge, the human auditory brain remains the best machine."

But this is exactly what Morten Kolbæk has worked on changing with his new [algorithm](#).

"Because of its ability to function in unknown environments with unknown voices, the applicability of this algorithm is so much stronger than what we have seen with previous technology. It's an important step forward when it comes to solving challenging listening situations in everyday life,"

says one of Morten Kolbæk's two supervisors, Jesper Jensen, Senior Researcher at Oticon and Professor at the Centre for Acoustic Signal Processing Research (CASPR) at AAU.

Professor Zheng-Hua Tan, who is also affiliated with CASPR and supervisor of the project, agrees on the major potential of the algorithm within sound research.

"The key to success for this algorithm is its ability to learn from data and then construct powerful statistical models that are able to represent complex listening situations. This leads to solutions that work very well even in new and unknown listening situations," explains Zheng-Hua Tan.

Noise reduction and speech separation

Specifically, Morten Kolbæk's Ph.D. project has dealt with two different but well-known listening scenarios.

The first track sets out to solve the challenges of one-to-one conversations in noisy spaces such as car cabins. Hearing aid users face such challenges on a regular basis.

"To solve them, we have developed algorithms that can amplify the sound of the speaker while reducing noise significantly without any prior knowledge about the listening situation. Current hearing aids are pre-programmed for a number of different situations, but in real life, the environment is constantly changing and requires a hearing aid that is able to read the specific situation instantly," explains Morten Kolbæk.

Demo of a single-microphone speech enhancement and separation system based on deep learning. The system is trained using utterance-level permutation invariant training (uPIT) and the system is speaker independent. That is, the speakers in the demo ...[more](#)

The second track of the project revolves around speech separation. This scenario involves several speakers, and the hearing aid user may be interested in hearing some or all of them. The solution is an algorithm that can separate voices while reducing noise. This track can be considered an extension of the first track, but now with two or more voices.

"You can say that Morten figured out that by tweaking a few things here and there, the algorithm works with several unknown speakers in noisy environments. Both of Morten's research tracks are significant and have attracted a great deal of attention," says Jesper Jensen.

Deep neural networks

The method used in creating the algorithms is called "deep learning," which falls under the machine learning category. More specifically, Morten Kolbæk has worked with deep neural networks, a type of algorithm that you train by feeding it examples of the signals it will encounter in the real world.

"If, for instance, we talk about speech-in-noise, you provide the algorithm with an example of a voice in a noisy environment and one of the voice without any noise. In this way, the algorithm learns how to process the noisy signal in order to achieve a clear voice signal. You feed the network with thousands of examples, and during this process, it will learn how to process a given voice in a realistic environment," Jesper Jensen explains.

"The power of deep learning comes from its hierarchical structure that is capable of transforming noisy or mixed voice signals into clean or separated voices through layer-by-layer processing. The widespread use of deep learning today is due to three major factors: ever-increasing computation power, increasing amount of big data for training algorithms and novel methods for training deep neural networks," says Zheng-Hua Tan.

A computer behind the ear

One thing is to develop the algorithm, another is to make it work in an actual hearing aid. Currently, Morten Kolbæk's algorithm for speech separation only works on a larger scale.

"When it comes to hearing aids, the challenge is always to make the technology work on a small computer behind the ear. And right now, Morten's algorithm requires too much space for this. Even if Morten's algorithm can separate several unknown voices from each other, it isn't able to choose which voice to present to the hearing aid user. So there are some practical issues that we need to solve before we can introduce it in a hearing aid solution. However, the most important thing is that these issues now seem solvable."

The cocktail party phenomenon

People with normal hearing are often capable of focusing on one speaker of interest, even in acoustically difficult situations where other people are speaking simultaneously. Known as the cocktail party phenomenon, the problem has generated a very active research area on how the human brain is able to solve this issue so well. With this Ph.D. project, we're one step closer toward solving this problem, Jesper Jensen explains:

"You sometimes hear that the cocktail party problem has been solved. This is not yet the case. If the environment and voices are completely unknown, which is often the case in the real world, current technology simply cannot match the human brain which works extremely well in unknown environments. But Morten's algorithm is a major step toward getting machines to function and help people with normal hearing and those with hearing loss in such environments," he says. [30]

Why we want to build a machine that can predict a person's attractiveness

It is an age-old question – what makes someone attractive? We often say things like "beauty is in the eye of the beholder" but while this romantic notion may bring comfort to those dealt a poor hand in life, it also gives the impression that the foundations of attractiveness are elusive and unpredictable. It suggests that what each of us sees as an attractive trait – whether physical or psychological – is so variable that everyone must be looking for something different.

While there is variety in what each of us regards as beautiful, cutting through this noise are common and consistent preferences. Psychological traits such as a sense of humour, intelligence and kindness are generally sought after. Similarly, physical attributes such as waist-to-hip ratio (the difference in waist and hip circumference), sex-typical voice pitch (basically, our expectation that men will have deep voices, and women high voices), and facial symmetry are also reliably desirable. Finding someone who could take or leave *some* of these characteristics may be easy, but one would have a hard time finding someone yearning to meet a sour-faced, selfish and dull person who refuses to take a shower.

While researchers have taken steps to [comprehensively catalogue](#) the preferences of men and women, we still don't know which traits are the most important contributors to a person's attractiveness. What we do know is that not all attractive traits are preferred equally. This can be revealed using some basic psychological tasks, such as asking people to design a partner by allocating points to enhance their characteristics (similar to designing a character in a video game).

When given only a small points budget, tough choices have to be made – and some characteristics normally attractive in their own right tend to fade into the background. [One study](#) found that creativity and talents were trumped by the likes of intelligence and social status during the task. Interestingly, basic kindness tends to be one of the top traits when building the ideal long-term partner.

These tasks are great for assessing the individual traits that make up mate preferences. But they do not necessarily capture how people make judgements about the attractiveness of living, breathing human beings. They may tell us that humour is important, for instance, but we balance a range of criteria in assessing attractiveness. A funny personality may seem less appealing in a person who is selfish.

Looking deeper into this, these tasks don't acknowledge the often complicated relationship between characteristics. For example, while the task might allow someone to design a partner who is low in intelligence but high in creativity, these attributes tend to go [hand in hand](#) in the real world.

This leaves us in the position where we know which traits are attractive, and have some idea of what preferences are prioritised over others. But by looking at different traits in isolation we are still missing the complete picture. Perhaps a better way to approach the problem would be to take an objective rating of a person's attractiveness (by asking the public to rate them [on a scale](#), for example) and then figuring out what traits hold the most influence over that number.

Doing this would require taking a large sample of the population and measuring all the psychological and physical traits known to contribute to attractiveness. Then, by adding in objective measures of overall attractiveness – and a dash of [machine learning](#) – creating models capable of learning what traits matter the most.

This is not some science fiction idea – in fact it is something that my lab at Swansea University [is currently crowdfunding](#). Machine learning is a powerful tool that has already accomplished feats such as predicting biological sex with [a 93% accuracy](#) based on brain scans alone. While we won't be scanning brains, we will be measuring dozens of our volunteers' characteristics – including humour, intelligence, impulsivity, facial symmetry, strength, and more.

In the first instance, we will use this information to calculate how these attributes combine to predict how one perceives one's own attractiveness. Then, we will extend this to predict objective judgements of attractiveness such as those made by the public after viewing online profiles of the volunteers.

This resulting model would be able to tell us that, for example, John's rating of "seven out of 10" by the public is primarily driven by his high intelligence, but held back a bit by his lack of muscle mass. It might also tell us that his poor sense of humour would normally work against him, but is

completely overshadowed by his high social status. It would also tell us which traits don't really matter at all – that nobody really cares about John's lack of hair.

After calibration, such a model would also be able to predict the attractiveness of new cases – without the need for public ratings. In other words, it could guess how the public would rate someone's attractiveness, based on a small number of important traits.

Ultimately, this system could even be used to advise people on how they can make themselves more [attractive](#) to a wider range of people. One only has to look at the billions of dollars spent every year on [make up](#) and [cosmetic surgery](#) alone to realise that there is a great public interest in what people can do to enhance their attractiveness.

Some enhancements, such as taking guitar lessons or learning magic tricks for example, may at first glance seem like good methods of self-improvement. However, ultimately, these may pale in comparison to the attractiveness boost experienced by finding a better-paid job or – perhaps controversially – simply trying to be a kinder person. But to know for sure, we need a method to sort the wheat from the chaff. Which is why we want to build a machine that can predict attractiveness. [29]

Could AI take control of human birth?

Instead of looking up at the sky to see whether you need an umbrella, people increasingly ask virtual assistants such as Alexa. And they may be wise to do so. AI methods are powerful – capable of anything from [analysing astrophysical data](#) to detecting tumours or helping to [manage diabetes](#). An algorithm that analyses shopping patterns recently [detected that a teenage girl was pregnant](#), earlier than her father did. So could childbirth be next for AI?

The most popular AI subset is "[machine learning](#)", which allows a machine to learn a task without being explicitly programmed. This is achieved by algorithms that are designed with the ability to discover relationships within large amounts of data.

Just imagine an AI system that can continuously read maternal and foetal movements, breathing patterns and bio-signals such as heart rate or blood pressure – and reliably identify crucial individual patterns in the physiology, emotions and behaviours of both mothers and foetuses – during childbirth. Through learning day by day, it would get more accurate at determining which combination of patterns would lead to which outcome. Could such a system be used to suggest what to do during labour, minute by minute, with excellent levels of accuracy – including whether to go ahead with a [vaginal birth](#) or opt for a caesarean section?

Maybe this could even reduce [unnecessary interventions](#) and [maternal mortality](#), in line with WHO recommendations. If a simple reduction in interventions is the aim, AI could be very promising, in theory.

Some supporters will say that such a system would save lives and taxpayers' money. Others will be horrified at the thought, feeling that it will result in a total loss of human companionship in labour – and of midwifery and obstetric skills and practice.

Fact or fiction?

But how close are we actually to having AI involved in childbirth – and do we know whether it would be beneficial?

A group of researchers from MIT have already developed an AI robot that can assist in a labour room. In their study involving doctors and nurses, they found that the robot's recommendations were accepted 90% of the time and that the number of errors were similar whether the robot was there or not. On that basis [they suggested it should be safe and efficient](#) to use AI during childbirth.

However, this also raises the question – if the technology is no better than human expertise, why would we use it? Especially as humans pick up a range of subtle cues that machines cannot perceive. For example, a clinical trial called INFANT showed that the use of software that was designed to improve the decision making of midwives and obstetricians for women who had continuous electronic foetal monitoring during labour [did not improve clinical outcome](#) when compared to expert judgement.

So, it may be some time before AI is rolled out in maternity units. But we cannot ignore the writing on the wall – the potential for totally AI supported birth is not so very fantastical.

Emotional support

But birth, the beginning of life, is not a transactional enterprise that only requires monitoring and measuring to be both safe and fulfilling. It is an interaction story between the woman, her baby, her partner, labour supporters and healthcare providers. For most women around the world, it is a profoundly important experience that has an [impact on parenting](#) and [self-esteem](#) – lasting way beyond the moment of birth.

It was recently recognised that caring companionship and human emotional and psychological support not only [improve birth health outcomes](#) for both women and infants, but could have long term effects into newborn's adult life. And current versions of AI are not actually that good at understanding human emotions or talking to people.

As we are focusing more and more on the prioritisation of measuring, monitoring, counting and recording in labour over simple human interaction, and as we fall more and more in love with our personal technology devices, there is a risk we could lose sight of what matters for human well-being in a range of areas. In fact, we are making it easier and easier to translate childbirth expertise into an Alexa-style birth assistant interface.

What will happen to women and babies if, as a result, AI one day becomes so smart that it controls us? Maybe, if AI could be used as an aid for skilled and caring [birth](#) professionals and childbearing women – rather than as the ultimate decision maker – it could contribute to the best and safest experience for each woman and her baby.

At the moment though, it seems like no matter that prominent thinkers such as the late [Stephen Hawking](#) and tech-entrepreneur [Elon Musk](#) have repeatedly warned of the potential risks of AI, we are on a headlong rush towards it. While it is impossible to stop a river in spate, it is possible to reroute it, and to use it for benefit rather than just letting it destroy everything in its path. [28]

A new approximate computing approach using CNNs


Researchers at Fukuoka University, in Japan, have recently proposed a design methodology for configurable approximate arithmetic circuits. As part of their study, published on [ResearchGate](#), they applied their method to a prototype system for image processing that relies on deep neural networks.

Approximate computing is a promising computation technique that relies on the ability of many systems to tolerate some loss of quality or optimality in computed results. By reducing the need for precise or entirely deterministic operations, it can achieve high performance with lower energy-consumption.

While many research studies have focused on approximate arithmetic circuits, configurable approximate circuits have only recently become of interest. One of the key challenges in this field is determining parameters for these circuits' configurations, a task that can often be difficult and tedious.

"Our current research focuses on approximate arithmetic circuits," Toshinori Sato, one of the researchers who carried out the study, told TechXplore. "These circuits trade computing accuracy for other design constraints such as power, speed and size. This makes them application specific, rather than for general purpose. We needed an easy design method for application specific circuits that would show the usefulness of approximate arithmetic circuits. We thought that neural networks might help to ease circuit designs."

Past research into configurable approximate circuits has mainly considered circuit structure, without examining target data. This makes it impossible for designers to carry out optimizations based on specific data. The method devised by Sato and his colleagues, on the other hand, is oriented towards processed data, thus considering both applications and their data simultaneously.



Credit: Ishida, Sato & Ukezono.

"We trained our prototype to identify the relationship between circuit design parameters and processed image quality," Sato explained. "After training, the prototype was able to generate an optimal design from configurable approximate arithmetic circuits, when a specific image is provided."

In the design proposed by Sato and his colleagues, the approximate circuit generator considers its target data, as well as some design constraints and user requirements. Their approach also handles parameterized approximate circuits, automatically determining their parameters. This relieves designers from the tedious and time-consuming task of manually determining parameters.

"The most characteristic feature of our method is that it is data-directed," Sato said. "Many studies have targeted a specific domain application, but we believe that only a few consider both applications and their data simultaneously. This feature is important, as it could ultimately facilitate the widespread adoption of approximate [circuits](#)."

The researchers evaluated their method, applying it to a proof-of-concept deep neural network-based prototype for image processing. Their [design](#) achieved promising results, processing images almost as well as traditional approaches with significant improvements in power (33.28 percent), delay (5.67 percent) and area (21.86 percent).

"Our prototype is still at an early stage of development," Sato said. "We now have to improve it so that it can consider more needs of designers, such as power consumption, circuit delay and size. Consequently, we would also like to apply the methodology to other [applications](#) beyond image processing." [27]

A new technique for synthesizing motion-blurred images

Researchers at Google have recently developed a new technique for synthesizing a motion blurred image, using a pair of un-blurred images captured in succession. In their paper, pre-published on *arXiv*, the researchers outlined their approach and evaluated it against several baseline methods.

Motion blur naturally occurs when the objects in a scene or the camera itself are moving as an image is being taken. This results in the moving object/s or the entire image appearing blurred. In some cases, motion blur might be used to indicate the speed of a photographed subject or to separate it from the background.

"Motion blur is a valuable cue in the context of image understanding," Tim Brooks and Jonathan Barron, the researchers who carried out the study, wrote in their paper. "Given a single image containing motion blur, one can estimate the relative direction and magnitude of scene motion that resulted in the observed blur. This motion estimate may be semantically meaningful, or may be used by a de-blurring algorithm to synthesize a sharp image."

Recent research has investigated the use of [deep learning](#) algorithms to remove unwanted motion blur from images or for inferring the motion dynamics of a given scene. To train these algorithms, however, researchers need a substantial amount of data, which is typically generated by synthetically blurring sharp images. Ultimately, the extent to which a deep learning algorithm can effectively remove motion blur in real images greatly depends on the realism of the synthetic data used to train it.

"In this paper, we treat the inverse of this well-studied blur estimation/removal task as a first class problem," Brooks and Barron wrote in their paper. "We present a fast and effective way to synthesize the training data necessary to train a motion de-blurring algorithm, and we quantitatively demonstrate that our technique generalizes from our synthetic training data to real motion-blurred imagery."

The [neural network architecture](#) devised by Brooks and Barron includes a novel "line prediction" layer, which teaches a system to regress from image pairs of consecutively taken images to a motion blurred image that spans the capture time of these two input images. Their model requires a vast amount of training data, so the researchers designed and executed a strategy that uses frame interpolation techniques to generate a large synthetic dataset of motion blurred images, along with their respective inputs.

Brooks and Barron also captured a high quality test set of real motion blurred images synthesized from slow motion videos and then used this to evaluate their model against baseline techniques. Their model achieved very promising results, outperforming existing approaches in both accuracy and speed.

"Our approach is fast, accurate, and uses readily available imagery from videos or 'bursts' as input, and so provides a path for enabling motion blur manipulation in consumer photography applications, and for synthesizing the realistic training data needed by deblurring or motion estimation algorithms," the researchers wrote in their paper.

While experienced photographers and cinematographers often use motion blur as an artistic effect, producing effective motion-blurred photographs can be very challenging. In most cases, these images are the product of a long trial and error process, requiring advanced skills and equipment.

Due to the difficulties in achieving quality motion blur effects, most consumer cameras are designed to take images with as little motion blur as possible. This means that amateur photographers have very little space to experiment with motion [blur](#) in their images.

"By allowing motion blurred images to be synthesized from the conventional un-blurred images that are captured by standard consumer cameras, our technique allows non-experts to create motion blurred images in a post-capture setting," the researchers explained in their paper.

Ultimately, the approach devised by Brooks and Barron could have a number of interesting applications. For instance, it could make artistic [motion blur](#) accessible to casual photographers, while also generating more realistic [motion](#) blurred images to train deep learning algorithms. [26]

Predicting the accuracy of a neural network prior to training

Constructing a neural network model for each new dataset is the ultimate nightmare for every data scientist. What if you could forecast the accuracy of the neural network earlier thanks to accumulated experience and approximation? This was the goal of a recent project at IBM Research and the result is [TAPAS](#) or Train-less Accuracy Predictor for Architecture Search ([click for demo](#)). Its trick is that it can estimate, in fractions of a second, classification performance for unseen input datasets, without training for both image and text classification.

In contrast to previously proposed approaches, TAPAS is not only calibrated on the topological network information, but also on the characterization of the [dataset](#) difficulty, which allows us to re-tune the prediction without any training.

This task was particularly challenging due to the heterogeneity of the datasets used for training neural networks. They can have completely different classes, structures, and sizes, adding to the complexity of coming up with an approximation. When my colleagues and I thought about how to address this, we tried not to think of this as a problem for a computer, but instead to think about how a human would predict the accuracy.

We understood that if you asked a human with some knowledge of deep learning whether a network would be good or bad, that person would naturally have an intuition about it. For example, we would recognize that two types of layers don't mix, or that after one type of layer, there is always another one which follows and improves the accuracy. So we considered whether adding features resembling this human intuitions into a computer could help it do an even better job. And we were correct.

We tested TAPAS on two datasets performed in 400 seconds on a single GPU, and our best discovered networks reached 93.67% accuracy for CIFAR-10 and 81.01% for CIFAR-100, verified by training. These networks perform competitively with other automatically discovered state-of-the-art networks, but needed only a small fraction of the time to solution and computational resources.

Our predictor achieves a performance which exceeds 100 networks per second on a single GPU, thus creating the opportunity to perform large-scale architecture search within a few minutes. We believe this is the first tool which can do predictions based on unseen data.

TAPAS is one of the AI engines in IBM's new breakthrough capability called [NeuNetS](#) as part of [AI OpenScale](#), which can synthesize custom [neural networks](#) in both text and image domains.

In NeuNetS, users will upload their data to the IBM Cloud and then TAPAS can analyze the data and rate it on a scale of 0-1 in terms of complexity of task, 0 meaning hard and 1 being simple. Next TAPAS starts to gather knowledge from its reference library looking for similar datasets based on what the user uploaded. Then based on this, TAPAS can accurately predict how a new [network](#) will perform on the new dataset, very similar to how a human would determine it.

Today's demand for data science skills already exceeds the current supply, becoming a real barrier towards adoption of AI in industry and society. TAPAS is a fundamental milestone towards the demolition of this wall. IBM and the Zurich Research Laboratory are working to make AI technologies as easy to use, as a few clicks on a mouse. This will allow non-expert users to build and deploy AI models in a fraction of the time it takes today—and without sacrificing [accuracy](#). Moreover, these tools will gradually learn over utilization in specialized domains and automatically improve over time, getting better and better. [25]

Are computer-aided decisions actually fair?

Algorithmic fairness is increasingly important because as more decisions of greater importance are made by computer programs, the potential for harm grows. Today, algorithms are already widely used to determine credit scores, which can mean the difference between owning a home and renting one. And they are used in predictive policing, which suggests a likelihood that a crime will be committed, and in scoring how likely a criminal will commit another crime in the future, which influences the severity of sentencing.

That's a problem, says Adam Smith, a professor of computer science at Boston University, because the design of many algorithms is far from transparent.

"A lot of these systems are designed by [private companies](#) and their details are proprietary," says Smith, who is also a data science faculty fellow at the Hariri Institute for Computing. "It's hard to know what they are doing and who is responsible for the decisions they make."

Recently, Smith and a joint team of BU-MIT computer scientists reexamined this problem, hoping to learn what, if anything, can be done to understand and minimize bias from decision-making systems that depend on [computer programs](#).

The BU researchers—Smith, Ran Canetti, a professor of computer science and director of the Hariri Institute's Center for Reliable Information Systems and Cyber Security, and Sarah Scheffler (GRS'21), a computer science doctoral candidate—are working with MIT Ph.D. students Aloni Cohen, Nishanth Dikkala, and Govind Ramnarayan to design systems whose decisions about all subsets of the population are equally accurate.

Their work was recently accepted for publication at the upcoming 2019 Association for Computing Machinery conference on Fairness, Accountability, and Transparency, nicknamed "ACM FAT."

The researchers believe that a system that discriminates against people who have had a hard time establishing a credit history will perpetuate that difficulty, limiting opportunity for a subset of the population and preserving existing inequalities. What that means, they say, is that automated ranking systems can easily become self-fulfilling prophecies, whether they are ranking the likelihood of default on a mortgage or the quality of a university education.

"Automated systems are increasingly complex, and they are often hard to understand for lay people and for the people about whom decisions are being made," Smith says.

The problem of self-fulfilling predictions

"The interaction between the algorithm and human behavior is such that if you create an algorithm and let it run, it can create a different society because humans interact with it," says Canetti. "So you have to be very careful how you design the algorithm."

That problem, the researchers say, will get worse as future algorithms use more outputs from past algorithms as inputs.

"Once you've got the same computer program making lots of decisions, any biases that exist are reproduced many times over on a larger scale," Smith says. "You get the potential for a broad societal shift caused by a computer program."

But how exactly can an algorithm, which is basically a mathematical function, be biased?

Scheffler suggests two ways: "One way is with biased data," she says. "If your algorithm is based on historical data, it will soon learn that a particular institution prefers to accept men over women. Another way is that there are different accuracies on different parts of the population, so maybe an algorithm is really good at figuring out if white people deserve a loan, but it could have high error rate for people who are not white. It could have 90 percent accuracy on one set of the population and 50 percent on another set."

"That's what we are looking at," says Smith. "We're asking 'How is the system making mistakes?' and 'How are these mistakes spread across different parts of the population?'"

The real-world impact of algorithmic bias

In May 2016, reporters from ProPublica, a nonprofit investigative newsroom, examined the accuracy of COMPAS, one of several algorithmic tools used by court systems to predict recidivism, or the likelihood that a criminal defendant will commit another crime. The initial findings were not reassuring.

When ProPublica researchers compared the tool's predicted risk of recidivism with actual recidivism rates over the following two years, they found that, in general, COMPAS got things right 61 percent of the time. They also found that predictions of violent recidivism were correct only 20 percent of the time.

More troubling, they found that black defendants were far more likely than white defendants to be incorrectly deemed more likely to commit crime again, and white defendants were more likely than

black defendants to be incorrectly deemed low risk to recidivate. According to ProPublica's article, this was a clear demonstration of bias by the algorithm.

In response, Northpointe Inc., the creator of COMPAS, published another study that argued that the COMPAS algorithm is in fact fair according to a different statistical measure of bias: calibration. Northpointe's software is widely used, and like many algorithmic tools, its calculations are proprietary, but the company did tell ProPublica that its formula for predicting who will recidivate is derived from answers to 137 questions whose answers come either from defendants or from criminal records.

Northpointe's study found that for each risk score, the fraction of white defendants who received this score and recidivated (out of all white defendants who received this score) roughly equals the fraction of black defendants who received this score and recidivated, out of all black defendants who received this score.

"ProPublica and Northpointe came to different conclusions in their analyses of COMPAS' fairness. However, both of their methods were mathematically sound—the opposition lay in their different definitions of fairness," Scheffler says.

The bottom line is that any imperfect prediction mechanism (either algorithmic or human) will be biased according to at least one of the two approaches: the error-balancing approach used by ProPublica, and the calibration method favored by Northpointe.

Overcoming algorithmic bias

When it came to solving the problem of algorithmic bias, the BU-MIT research team created a method of identifying the subset of the population that the system fails to judge fairly, and sending their review to a different system that is less likely to be biased. That separation guarantees that the method errs in more balanced ways regarding the individuals for whom it does make a decision.

And while the researchers found many situations where that solution appeared to work well, they remain concerned about how the different systems would work together. "There are many different measures of fairness," says Scheffler, "and there are trade-offs between them. So to what extent are the two systems compatible with the notion of fairness we want to achieve?"

"What happens to those people whose decisions would be deferred really influences how we view the system as a whole," says Smith. "At this point, we are still wrapping our heads around what the different solutions would mean."

Still, says Canetti, the research points to a possible way out of the statistical bias conundrum, one that could enable the design of algorithms that minimize the bias. That challenge, he says, will require expertise from many disciplines. [24]

How is artificial intelligence changing science?

Intel's Gadi Singer believes his most important challenge is his latest: using artificial intelligence (AI) to reshape scientific exploration.

In a Q&A timed with the first Intel AI DevCon event, the Intel vice president and architecture general manager for its Artificial Intelligence Products Group discussed his role at the intersection of science—computing's most demanding customer—and AI, how scientists should approach AI and why it is the most dynamic and exciting opportunity he has faced.

Q. How is AI changing science?

Scientific exploration is going through a transition that, in the last 100 years, might only be compared to what happened in the '50s and '60s, moving to data and large data systems. In the '60s, the amount of data being gathered was so large that the frontrunners were not those with the finest instruments, but rather those able to analyze the data that was gathered in any scientific area, whether it was climate, seismology, biology, pharmaceuticals, the exploration of new medicine, and so on.

Today, the data has gone to levels far exceeding the abilities of people to ask particular queries or look for particular insights. The combination of this data deluge with modern computing and deep learning techniques is providing new and many times more disruptive capabilities.

Q. What's an example?

One of them, which uses the basic strength of deep learning, is the identification of very faint patterns within a very noisy dataset, and even in the absence of an exact mathematical model of what you're looking for.

Think about cosmic events happening in a far galaxy, and you're looking for some characteristics of the phenomena to spot them out of a very large dataset. This is an instance of searching without a known equation, where you are able to give examples, and through them, let the deep learning system learn what to look for and ultimately find out a particular pattern.

Q. So you know what you're looking for, but you don't know how to find it?

You can't define the exact mathematical equation or the queries that describe it. The data is too large for trial-and-error and previous big-data analytics techniques do not have enough defined features to successfully search for the pattern.

You know what you're looking for because you tagged several examples of it in your data, and you can generally describe it. Deep learning can help you spot occurrences from such a class within a noisy multidimensional dataset.

Q. Are there other ways AI can change the scientific approach?

Another example is when you do have a mathematical model, like a set of accurate equations. In this case you can use AI to achieve comparable results in 10,000 times less time and computing.

Say you have a new molecular structure and you want to know how it's going to behave in some environment for pharma exploration. There are very good predictive models on how it will behave. The problem is that those models take a tremendous amount of computation and time—it might take you weeks to try just one combination.

More: [Intel AI VP Gadi Singer on One Song to the Tune of Another \(The Next Platform\)](#) | [Intel AI DevCon \(Press Kit\)](#) | [Artificial Intelligence at Intel \(Press Kit\)](#) | [More Intel Explainers](#)

In such a case, you can use a deep learning system to shadow the accurate system of equations. You iteratively feed sample cases to this system of equations, and you get the results days later. The deep learning network learns the relationship between the input and the output, without knowing the equation itself. It just tracks it. It was demonstrated in multiple cases that, after you train the deep learning system with enough examples, it shows excellent ability to predict the result that will be given by the exact model. This translates to an efficiency that could turn hours or days into second.

Granted, sometimes the full computation will be required for ultimate model accuracy. However, that would only be needed for a small subset of cases. The fact that you can generate an accurate result so much faster with a fraction of the power and the time allows you to explore the potential solution space much faster.

In the last couple years, new machine learning methods have emerged for "learning how to learn." These technologies are tackling an almost-endless realm of options—like all the possible mutations in human DNA—and are using exploration and meta-learning techniques to identify the most relevant options to evaluate.

Q. What's the big-picture impact to the scientific method or just the approach that a scientist would take with AI?

Scientists need to partner with AI. They can greatly benefit from mastering the tools of AI, such as [deep learning](#) and others, in order to explore phenomena that are less defined, or when they need faster performance by orders of magnitude to address a large space. Scientists can partner with machine learning to explore and investigate which new possibilities have the best likelihood of breakthroughs and new solutions.

Q. I'm guessing you could retire if you wanted to. What keeps you going now?

Well, I'm having a great time. AI at Intel today is about solving the most exciting and most challenging problems the industry and science are facing. This is an area that moves faster than anything I've seen in my 35 years at Intel, by far.

The other aspect is that I'm looking at it as a change that is brewing in the interaction between humans and machines. I want to be part of the effort of creating this new link. When I talk about partnership of science and AI, or autonomous vehicles and other areas, there's a role here for a broader thinking than just how to give the fastest processor for the task. This newly forged interaction between people and AI is another fascinating part of this space. [23]

The U.S. needs a national strategy on artificial intelligence

China, India, Japan, France and the European Union are crafting bold plans for artificial intelligence (AI). They see AI as a means to economic growth and social progress. Meanwhile, the U.S. disbanded its AI taskforce in 2016. Without an AI strategy of its own, the world's technology leader risks falling behind.

The U.S. technology sector has long been a driver of global economic growth. From the PC to the Internet, the greatest advancements of the past 50 years were spawned in the U.S. This country's unique approach to limited regulation combined with public-private partnerships creates an environment for innovation generally unmatched in the free world. A national AI strategy can build on this history of economic and technological leadership.

Artificial intelligence is astonishing in its potential. It will be more transformative than the PC and the Internet. Already it is poised to solve some of our biggest challenges. As examples, AI has been used to more precisely detect and diagnose cancer, treat depression, improve crop yields, save energy, increase supply chain efficiencies and protect our financial systems. These are remarkable successes for such a young technology.

Governments can and should help build on these successes. A national strategy for AI will provide the necessary guideposts that enable [industry](#) and academia to innovate. When the regulatory environment is known and understood, businesses and government researchers can maximize their impact by pursuing the same goals.

In this context, it will also be important to address concerns about AI's impact on individuals. Privacy, cybersecurity, ethics and potential employment impact are all worthy of careful analysis. Governments and industry can and should work together to better understand these concerns before any new regulation is enacted.

As evidenced by their AI plans, governments around the world see AI as a catalyst for [economic growth](#) and a means to improve the lives of their citizens. They are prioritizing research and development and the establishment of a strong and diverse ecosystem to bring AI to fruition.

China's plan, for example, includes measurable objectives and detailed direction on specific areas of focus. This is backed by significant public-private funding commitments as well as industry-government alignment on direction.

The EU's strategy provides deliberate direction to avoid regulation while investing in R&D. It offers a clear focus on greater investment, preparation for socio-economic changes, and formation of an ethical and legal framework. Japan, India, France and others are adopting similar strategies.

Industry has partnered with many of these governments to develop their plans; we stand ready to work with the U.S. government in the same way. A good model for success is the semiconductor sector, where industry and the U.S. government partnered in the early 1980s to build the vast ecosystem that is considered the North Star for technology success today. AI can be history's greatest economic engine. Governments can – and should – help make this real.

Before disbandment, the U.S. Artificial Intelligence Research and Development Taskforce defined seven strategic objectives and two specific recommendations for AI. This report called on the government to develop a more detailed AI R&D plan, and study the creation of an AI R&D workforce. These recommendations can be the starting point for a definitive U.S. national strategy for AI.

This is not a call for a swarm of new laws and regulations. Rather, a U.S. national strategy can provide the structure for researchers and industry to follow as they develop [artificial intelligence](#). Such direction provides operating certainty that lessens risk.

A national [strategy](#) therefore should aim to foster innovation across the industry and academia, and prepare society for changes to come. It can also provide operating clarity that lessens business risk. Two areas of focus should be prioritized: government funding of R&D to augment the great work being done by industry and the availability of [government](#) data for innovators to use in developing artificial intelligence capabilities. AI needs data to learn, and there are ways to do this without compromising privacy and security.

AI is too big for one company – or one country – to realize alone. The transformative potential of AI has been likened to electricity and the steam engine. Ensuring a role for the U.S. in this global revolution is critical to not just the U.S. economic engine but that of our entire world. [22]

Can artificial intelligence help find alien intelligence?

In the search for extraterrestrial intelligence (SETI), we've often looked for signs of intelligence, technology and communication that are similar to our own.

But as astronomer and SETI trailblazer Jill Tarter [points out](#), that approach means searching for detectable [technosignatures](#), like radio transmissions, not searching for [intelligence](#).

Now scientists are considering whether artificial intelligence (AI) could help us search for alien intelligence in ways we haven't even thought of yet.

'Decoding' intelligence

As we think about [extraterrestrial intelligence](#) it's helpful to remember [humans are not the only intelligent life](#) on Earth.

[Chimpanzees have culture and use tools](#), spiders [process information with webs](#), [cetaceans have dialects](#), crows [understand analogies](#) and [beavers are great engineers](#). Non-human intelligence, language, culture and technology are all around us.

Alien intelligence could look like an octopus, an ant, a dolphin or [a machine](#) —or be radically different from anything on Earth.

We often imagine extraterrestrial life relative to our ideas about difference, but those ideas aren't even universal on Earth and are [unlikely to be universal across interstellar space](#).

If some of us have only recently recognized non-human intelligence on Earth, what could we be missing when we imagine extraterrestrial life?

In early 2018, astronomers, neuroscientists, anthropologists, AI researchers, historians and others gathered for a "[Decoding Alien Intelligence](#)" workshop at the SETI Institute in Silicon Valley. Astrobiologist Nathalie Cabrol organized the workshop around her 2016 paper "[Alien](#)

minds," where she calls for a new SETI road map and a long-term vision for "the search for life as we do not know it."

In her paper, Cabrol asks how SETI can move past "looking for other versions of ourselves" and think "outside of our own brains" to imagine truly different extraterrestrial intelligence.

Thinking differently

Silicon Valley is famous for valuing "disruptive" thinking and this culture intersects with SETI research. Ever since the U.S. government stopped funding SETI in the mid-1990s, Silicon Valley ideas, technology and funding have been increasingly important.



A capuchin (*Sapajus libidinosus*) using a stone tool (T. Falótico). An octopus (*Amphioctopus marginatus*) carrying shells as shelter (N. Hobgood). (Wikimedia/Tiago Falótico, Nick Hobgood), CC BY-NC-SA

For example, the SETI Institute's Allen Telescope Array is named after Microsoft co-founder Paul Allen, who contributed over US\$25 million to the project. And, in 2015, technology investor Yuri Milner announced Breakthrough Listen, a 10-year US\$100 million SETI initiative.

Now, the SETI Institute, NASA, Intel, IBM and other partners are tackling space science problems through an AI research and development program called the Frontier Development Lab.

Lucianne Walkowicz, the [Astrobiology Chair at the Library of Congress](#), described one AI-based method as "signal agnostic searching" at [Breakthrough Discuss](#) in 2017.

Walkowicz explained that this means using machine learning methods to look at any set of data without predetermined categories and instead let that data cluster into their "natural categories." The software then lets us know what stands out as outliers. These outliers could then be the target of additional investigations.

It turns out that SETI researchers think AI might be useful in their work because they believe machine learning is good at spotting difference.

But its success depends on how we —and the AI we create —conceptualize the idea of difference.

Smarter than slime mould?

Thinking outside our brains also means thinking [outside our scientific, social and cultural systems](#). But how can we do that?

AI has been used to look for [simulations of what researchers imagine alien radio signals might look like](#), but now SETI researchers hope it can find things we aren't yet looking for.

Graham Mackintosh, an AI consultant at the SETI Institute workshop, said extraterrestrials might be doing things we can't even imagine, using technologies so different we don't even think to look for them. AI, he proposed, might be able to do that advanced thinking for us.

We may not be able to make ourselves smarter, but perhaps, Mackintosh suggested, we can make machines that are smarter for us.

In a keynote at this year's Breakthrough Discuss conference, astrophysicist Martin Rees shared a similar hope, that AI could lead to "intelligence which surpasses humans as much as we intellectually surpass slime mould."



Parts of the *Armillaria ostoyae* organism include the mushrooms, the black rhizomorphs and the white mycelial felts. Credit: USDA/Forest Service/Pacific Northwest Region

First contact

If we met extraterrestrial slime mould, what could we assume about its intelligence? One challenge of SETI is that we don't know the limits of life or intelligence, so we need to be open to all possible forms of difference.

We might find intelligence in forms that Euro-American science has historically disregarded: Microbial communities, insects or other complex systems like the symbiotic plant-fungus relationships in mycorrhizal networks that learn from experience.

Intelligence might appear in atmospheres or geology at a planetary scale, or as astrophysical phenomena. What appears to be a background process in the universe, or just part of what we think of as nature, could turn out to be intelligence.

Consider that the largest living thing on Earth may be an *Armillaria ostoyae* fungus in Eastern Oregon's Blue Mountains, which extends to 10 square kilometres and is between 2,000 and 9,000 years old.

While this fungus may not be what most people think of as intelligence, it reminds us to think about the unexpected when searching for life and intelligence, and of what we might be missing right under our feet.

Thinking differently about intelligence means understanding that anything we encounter could be first contact with intelligent life. This might include our first encounter with artificial general intelligence (AGI), also called Strong AI, something closer to the sentient computer HAL 9000 from *2001: A Space Odyssey* or Data from *Star Trek: The Next Generation*.

As we work with machine learning to expand the SETI search, we also need social sciences to understand how our ideas shape the future of AI—and how AI will shape the future of our ideas.

Interdisciplinary futures

To avoid a human-centred point of view in SETI we need to consider how we encode ideas about difference into AI and how that shapes the outcomes. This is vital for finding and recognizing intelligence as we don't yet know it.

Some of the methods used in anthropology can help us identify ideas about difference that we've naturalized—concepts so familiar they seem invisible, like the divides many still see between nature and culture or biology and technology, for example.

Recent research on algorithms reveals how our naturalized ideas [shape the technology we create and how we use it](#). And Microsoft's [infamous AI chat bot Tay](#) reminds us the AI we create can easily reflect the worst of those ideas.

We may never entirely stop building bias into search engines and search strategies for SETI, or coding it into AI. But through collaborations between scientists and social scientists we can think critically about how we conceptualize difference.

A critical, interdisciplinary approach will help us understand how our ideas about difference impact lives, research and possibilities for the future both here on Earth and beyond. [21]

Scientists make a maze-running artificial intelligence program that learns to take shortcuts

Call it an a-MAZE-ing development: A U.K.-based team of researchers has developed an artificial intelligence program that can learn to take shortcuts through a labyrinth to reach its goal. In the process, the program developed structures akin to those in the human brain.

The emergence of these computational "[grid cells](#)," described in the journal *Nature*, could help scientists design better navigational software for future robots and even offer a new window through which to probe the mysteries of the mammalian brain.

In recent years, AI researchers have developed and fine-tuned deep-learning networks—layered programs that can come up with novel solutions to achieve their assigned goal. For example, a

deep-learning network can be told which face to identify in a series of different photos, and through several rounds of training, can tune its algorithms until it spots the right face virtually every time.

These networks are inspired by the brain, but they don't work quite like them, said Francesco Savelli, a neuroscientist at Johns Hopkins University who was not involved in the paper. So far, AI systems don't come close to emulating the brain's architecture, the diversity of real neurons, the complexity of individual neurons or even the rules by which they learn.

"Most of the learning is thought to occur with the strengthening and weakening of these synapses," Savelli said in an interview, referring to the connections between neurons. "And that's true of these AI systems too—but exactly how you do it, and the rules that govern that kind of learning, might be very different in the brain and in these systems."

Regardless, AI has been really useful for a number of functions, from facial recognition to deciphering handwriting and translating languages, Savelli said. But higher-level activities—such as navigating a complex environment—have proved far more challenging.

One aspect of navigation that our brains seem to perform without conscious effort is path integration. Mammals use this process to recalculate their position after every step they take by accounting for the distance they've traveled and the direction they're facing. It's thought to be key to the brain's ability to produce a map of its surroundings.

Interview with Caswell Barry about grid cells. Credit: DeepMind

Among the neurons associated with these "cognitive maps": [place cells](#), which light up when their owner is in some particular spot in the environment; [head-direction cells](#), which tell their owner what direction they're facing; and grid cells, which appear to respond to an imaginary hexagonal grid mapped over the surrounding terrain. Every time a person steps on a "node" in this grid, the neuron fires.

"Grid cells are thought to endow the cognitive map with geometric properties that help in planning and following trajectories," Savelli and fellow Johns Hopkins neuroscientist James Knierim wrote in a commentary on the paper. The discovery of grid cells earned three scientists the 2014 Nobel Prize in physiology or medicine.

Humans and other animals seem to have very little trouble moving through space because all of these highly specialized neurons work together to tell us where we are and where we're going.

Scientists at DeepMind, which is owned by Google and University College London, wondered whether they could develop a program that could also perform path integration. So they trained the network with simulations of paths used by rodents looking for food. They also gave it data for a rodent's movement and speed as well as feedback from simulated place cells and head-direction cells.

During this training, the researchers noticed something strange: The simulated rodent appeared to develop patterns of activity that looked remarkably like grid cells—even though grid cells had not been part of their training system.

"The emergence of grid-like units is an impressive example of deep learning doing what it does best: inventing an original, often unpredicted internal representation to help solve a task," Savelli and Knierim wrote.

Interview with Matt Botvinick about neuroscience and AI. Credit: DeepMind

Grid cells appear to be so useful for path integration that this faux-rodent came up with a solution eerily similar to a real rodent brain. The researchers then wondered: Could grid cells also be useful in another crucial aspect of mammal navigation?

That aspect, called vector-based navigation, is basically the ability to calculate the straight-shot, "as the crow flies" distance to a goal even if you originally took a longer, less-direct route. That's a useful skill for finding shortcuts to your destination, Savelli pointed out.

To test this, researchers challenged the grid-cell-enabled faux-rodent to solve a maze, but blocked off most of the doorways so the program would have to take the long route to its goal. They also modified the program so it was rewarded for actions that brought it closer to the goal. They trained the network on a given maze and then opened shortcuts to see what happened.

Sure enough, the simulated rodent with grid cells quickly found and used the shortcuts, even though those pathways were new and unknown. And it performed far better than a faux-rodent whose start point and goal point were tracked only by place cells and head-direction cells. It even beat out a "human expert," the study authors said.

The findings eventually could prove useful for robots making their way through unknown territory, Savelli said. And from a neuroscientific perspective, they could help researchers better understand how these neurons do their job in the mammalian brain.

Of course, this program was highly simplified compared to its biological counterpart, Savelli pointed out. In the simulated rodent, the "place cells" didn't change—even though place cells and grid cells influence each other in complex ways in real brains.

"By developing the network such that the place-cell layer can be modulated by grid-like inputs, we could begin to unpack this relationship," Savelli and Knierim wrote.

Developing this AI program further could help scientists start to understand all the complex relationships that come into play in living neural systems, they added.

But whether they want to hone the technology or use it to understand biology, scientists will have to get a better handle on their own deep-learning programs, whose solutions to problems are often hard to decipher even if they consistently get results, scientists said.

"Making deep-learning systems more intelligible to human reasoning is an exciting challenge for the future," Savelli and Knierim wrote. [20]

Dissecting artificial intelligence to better understand the human brain

In the natural world, intelligence takes many forms. It could be a bat using echolocation to expertly navigate in the dark, or an octopus quickly adapting its behavior to survive in the deep ocean.

Likewise, in the computer science world, multiple forms of artificial intelligence are emerging - different networks each trained to excel in a different task. And as will be presented today at the 25th annual meeting of the Cognitive Neuroscience Society (CNS), cognitive neuroscientists increasingly are using those emerging artificial networks to enhance their understanding of one of the most elusive intelligence systems, the human brain.

"The fundamental questions cognitive neuroscientists and computer scientists seek to answer are similar," says Aude Oliva of MIT. "They have a complex system made of components - for one, it's called neurons and for the other, it's called units - and we are doing experiments to try to determine what those components calculate."

In Oliva's work, which she is presenting at the CNS symposium, neuroscientists are learning much about the role of contextual clues in human image recognition. By using "artificial neurons" - essentially lines of code, software - with neural network models, they can parse out the various elements that go into recognizing a specific place or object.

"The brain is a deep and complex neural network," says Nikolaus Kriegeskorte of Columbia University, who is chairing the symposium. "Neural network models are brain-inspired models that are now state-of-the-art in many artificial intelligence applications, such as computer vision."

In one recent study of more than 10 million images, Oliva and colleagues taught an artificial network to recognize 350 different places, such as a kitchen, bedroom, park, living room, etc. They expected the network to learn objects such as a bed associated with a bedroom. What they didn't expect was that the network would learn to recognize people and animals, for example dogs at parks and cats in living rooms.

The machine intelligence programs learn very quickly when given lots of data, which is what enables them to parse contextual learning at such a fine level, Oliva says. While it is not possible to dissect human neurons at such a level, the computer model performing a similar task is entirely transparent. The artificial neural networks serve as "mini-brains that can be studied, changed, evaluated, compared against responses given by human neural networks, so the cognitive neuroscientists have some sort of sketch of how a real brain may function."

Indeed, Kriegeskorte says that these models have helped neuroscientists understand how people can recognize the objects around them in the blink of an eye. "This involves millions of signals emanating from the retina, that sweep through a sequence of layers of neurons, extracting semantic information, for example that we're looking at a street scene with several people and a dog," he says. "Current neural network models can perform this kind of task using only computations that biological neurons can perform. Moreover, these neural network models can predict to some extent how a neuron deep in the brain will respond to any image."

Using computer science to understand the human brain is a relatively new field that is expanding rapidly thanks to advancements in computing speed and power, along with neuroscience imaging

tools. The artificial networks cannot yet replicate human visual abilities, Kriegeskorte says, but by modeling the human brain, they are furthering understanding of both cognition and [artificial intelligence](#). "It's a uniquely exciting time to be working at the intersection of neuroscience, cognitive science, and AI," he says.

Indeed, Oliva says; "Human cognitive and computational neuroscience is a fast-growing area of research, and knowledge about how the [human brain](#) is able to see, hear, feel, think, remember, and predict is mandatory to develop better diagnostic tools, to repair the [brain](#), and to make sure it develops well." [19]

Army's brain-like computers moving closer to cracking codes

U.S. Army Research Laboratory scientists have discovered a way to leverage emerging brain-like computer architectures for an age-old number-theoretic problem known as integer factorization.

By mimicking the brain functions of mammals in computing, Army scientists are opening up a new solution space that moves away from traditional computing architectures and towards devices that are able to operate within extreme size-, weight-, and power-constrained environments.

"With more computing power in the battlefield, we can process information and solve computationally-hard problems quicker," said Dr. John V. "Vinnie" Monaco, an ARL computer scientist. "Programming the type of devices that fit these criteria, for example, brain-inspired computers, is challenging, and cracking crypto codes is just one application that shows we know how to do this."

The problem itself can be stated in simple terms. Take a composite integer N and express it as the product of its prime components. Most people have completed this task at some point in grade school, often an exercise in elementary arithmetic. For example, 55 can be expressed as $5 \cdot 11$ and 63 as $3 \cdot 3 \cdot 7$. What many didn't realize is they were performing a task that if completed quickly enough for large numbers, could break much of the modern day internet.

Public key encryption is a method of secure communication used widely today, based on the RSA algorithm developed by Rivest, Shamir, and Adleman in 1978. The security of the RSA algorithm relies on the difficulty of factoring a large composite integer N , the public key, which is distributed by the receiver to anyone who wants to send an encrypted message. If N can be factored into its prime components, then the private key, needed to decrypt the message, can be recovered. However, the difficulty in factoring large integers quickly becomes apparent.

As the size of N increases by a single digit, the time it would take to factor N by trying all possible combinations of prime factors is approximately doubled. This means that if a number with ten digits takes 1 minute to factor, a number with twenty digits will take about 17 hours and a number with 30 digits about two years, an exponential growth in effort. This difficulty underlies the security of the RSA algorithm.

Challenging this, Monaco and his colleague Dr. Manuel Vindiola, of the lab's Computational Sciences Division, demonstrated how brain-like computers lend a speedup to the currently best known algorithms for factoring integers.

The team of researchers have devised a way to factor large composite integers by harnessing the massive parallelism of novel computer architectures that mimic the functioning of the mammalian brain. So called [neuromorphic computers](#) operate under vastly different principles than conventional computers, such as laptops and mobile devices, all based on an architecture described by John von Neumann in 1945.

In the von Neumann architecture, memory is separate from the central processing unit, or CPU, which must read and write to memory over a bus. This bus has a limited bandwidth, and much of the time, the CPU is waiting to access memory, often referred to as the von Neumann bottleneck.

Neuromorphic computers, on the other hand, do not suffer from a von Neumann bottleneck. There is no CPU, memory, or bus. Instead, they incorporate many individual computation units, much like neurons in the brain.

These units are connected by physical or simulated pathways for passing data around, analogous to synaptic connections between neurons. Many neuromorphic devices operate based on the physical response properties of the underlying material, such as graphene lasers or magnetic tunnel junctions. Because of this, these devices consume orders of magnitude less energy than their von Neumann counterparts and can operate on a molecular time scale. As such, any algorithm capable of running on these devices stands to benefit from their capabilities.

The speedup acquired by the ARL researchers is due to the formulation of a method for integer factorization with the help of a neuromorphic co-processor. The current fastest algorithms for factoring integers consist primarily of two stages, sieving and a matrix reduction, and the sieving stage comprises most of the computational effort.

Sieving involves searching for many integers that satisfy a certain property called B-smooth, integers that don't contain a prime factor greater than B. Monaco and Vindiola were able to construct a neural network that discovers B-smooth numbers quicker and with greater accuracy than on a von Neumann architecture. Their algorithm leverages the massive parallelism of brain-inspired computers and the innate ability of individual neurons to perform arithmetic operations, such as addition. As neuromorphic architectures continue to increase in size and speed, not limited by Moore's Law, their ability to tackle larger integer factorization problems also grows. In their work, it's estimated that 1024-bit keys could be broken in about a year, a task once thought to be out of reach. For comparison, the current record, a 232 decimal digit number (RSA-768) took about 2,000 years of computing time over the course of several years.

From a broader perspective, this discovery pushes us to question how a shift in computing paradigm might affect some of our most basic security assumptions. As emerging devices shift to incorporate massive parallelism and harness material physics to compute, the computational hardness underlying some security protocols may be challenged in ways not previously imagined. This work also opens the door to new research areas of emerging computer architectures, in terms of

algorithm design and function representation, alongside low-power machine learning and artificial intelligence applications.

"Encrypted messages in warfare often have an expiration date, when their contents become un-actionable," Monaco said. "There is an urgency to decrypt enemy communications, especially those at the field level, since these expire the quickest, compared to communication at higher echelons. In field conditions, power and connectivity are extremely limited. This is a strong motivating factor for using a brain-inspired computer for such a task where conventional computers are not practical."

[18]

Teaching computers to guide science: Machine learning method sees forests and trees

While it may be the era of supercomputers and "big data," without smart methods to mine all that data, it's only so much digital detritus. Now researchers at the Department of Energy's Lawrence Berkeley National Laboratory (Berkeley Lab) and UC Berkeley have come up with a novel machine learning method that enables scientists to derive insights from systems of previously intractable complexity in record time.

In a paper published recently in the *Proceedings of the National Academy of Sciences (PNAS)*, the researchers describe a technique called "iterative Random Forests," which they say could have a transformative effect on any area of science or engineering with complex systems, including biology, precision medicine, materials science, environmental science, and manufacturing, to name a few.

"Take a human cell, for example. There are 10^{170} possible molecular interactions in a single cell. That creates considerable computing challenges in searching for relationships," said Ben Brown, head of Berkeley Lab's Molecular Ecosystems Biology Department. "Our method enables the identification of interactions of high order at the same computational cost as main effects - even when those interactions are local with weak marginal effects."

Brown and Bin Yu of UC Berkeley are lead senior authors of "Iterative Random Forests to Discover Predictive and Stable High-Order Interactions." The co-first authors are Sumanta Basu (formerly a joint postdoc of Brown and Yu and now an assistant professor at Cornell University) and Karl Kumbier (a Ph.D. student of Yu in the UC Berkeley Statistics Department). The paper is the culmination of three years of work that the authors believe will transform the way science is done. "With our method we can gain radically richer information than we've ever been able to gain from a learning machine," Brown said.

The needs of machine learning in science are different from that of industry, where machine learning has been used for things like playing chess, making self-driving cars, and predicting the stock market.

"The machine learning developed by industry is great if you want to do high-frequency trading on the [stock market](#)," Brown said. "You don't care why you're able to predict the stock will go up or down. You just want to know that you can make the predictions."

But in science, questions surrounding why a process behaves in certain ways are critical. Understanding "why" allows scientists to model or even engineer processes to improve or attain a desired outcome. As a result, machine learning for science needs to peer inside the black box and understand why and how computers reached the conclusions they reached. A long-term goal is to use this kind of information to model or engineer systems to obtain desired outcomes.

In highly complex systems - whether it's a single cell, the human body, or even an entire ecosystem - there are a large number of variables interacting in nonlinear ways. That makes it difficult if not impossible to build a model that can determine cause and effect. "Unfortunately, in biology, you come across interactions of order 30, 40, 60 all the time," Brown said. "It's completely intractable with traditional approaches to statistical learning."

The method developed by the team led by Brown and Yu, iterative Random Forests (iRF), builds on an algorithm called random forests, a popular and effective predictive modeling tool, translating the internal states of the black box learner into a human-interpretable form. Their approach allows researchers to search for complex interactions by decoupling the order, or size, of interactions from the computational cost of identification.

"There is no difference in the computational cost of detecting an interaction of order 30 versus an interaction of order two," Brown said. "And that's a sea change."

In the PNAS paper, the scientists demonstrated their method on two genomics problems, the role of gene enhancers in the fruit fly embryo and alternative splicing in a human-derived cell line. In both cases, using iRF confirmed previous findings while also uncovering previously unidentified higher-order interactions for follow-up study.

Brown said they're now using their method for designing phased array laser systems and optimizing sustainable agriculture systems.

"We believe this is a different paradigm for doing science," said Yu, a professor in the departments of Statistics and Electrical Engineering & Computer Science at UC Berkeley. "We do prediction, but we introduce stability on top of prediction in iRF to more reliably learn the underlying structure in the predictors."

"This enables us to learn how to engineer systems for goal-oriented optimization and more accurately targeted simulations and follow-up experiments," Brown added.

In a [PNAS commentary](#) on the technique, Danielle Denisko and Michael Hoffman of the University of Toronto wrote: "iRF holds much promise as a new and effective way of detecting interactions in a variety of settings, and its use will help us ensure no branch or leaf is ever left unturned." [17]

Rise of the quantum thinking machines

Quantum computers can be made to utilize effects such as quantum coherence and entanglement to accelerate machine learning.

Although we typically view information as being an abstract or virtual entity, information, of course, must be stored in a physical medium. Information processing devices such as computers and phones are therefore fundamentally governed by the laws of physics. In this way, the fundamental physical limits of an agent's ability to learn are governed by the laws of physics. The best known theory of physics is quantum theory, which ultimately must be used to determine the absolute physical limits of a machine's ability to learn.

A quantum algorithm is a stepwise procedure performed on a quantum computer to solve a problem such as searching a database. Quantum machine learning software makes use of quantum algorithms to process information in ways that classical computers cannot. These quantum effects open up exciting new avenues which can, in principle, outperform the best known classical algorithms when solving certain machine learning problems. This is known as quantum enhanced machine learning.

Machine learning methods use mathematical algorithms to search for certain patterns in large data sets. Machine learning is widely used in biotechnology, pharmaceuticals, particle physics and many other fields. Thanks to the ability to adapt to new data, machine learning greatly exceeds the ability of people. Despite this, machine learning cannot cope with certain difficult tasks.

Quantum enhancement is predicted to be possible for a host of machine learning tasks, ranging from optimization to quantum enhanced deep learning.

In the new paper published in Nature, a group of scientists led by Skoltech Associate Professor Jacob Biamonte produced a feasibility analysis outlining what steps can be taken for practical quantum enhanced machine learning.

The prospects of using quantum computers to accelerate machine learning has generated recent excitement due to the increasing capabilities of quantum computers. This includes a commercially available 2000 spin quantum accelerated annealing by the Canada-based company D-Wave Systems Inc. and a 16 qubit universal quantum processor by IBM which is accessible via a (currently free) cloud service.

The availability of these devices has led to increased interest from the machine learning community. The interest comes as a bit of a shock to the traditional quantum physics community, in which researchers have thought that the primary applications of quantum computers would be using quantum computers to simulate chemical physics, which can be used in the pharmaceutical industry for drug discovery. However, certain quantum systems can be mapped to certain machine learning models, particularly deep learning models. Quantum machine learning can be used to work in tandem with these existing methods for quantum chemical emulation, leading to even greater capabilities for a new era of quantum technology.

"Early on, the team burned the midnight oil over Skype, debating what the field even was—our synthesis will hopefully solidify topical importance. We submitted our draft to Nature, going forward subject to significant changes. All in all, we ended up writing three versions over eight months with nothing more than the title in common," said lead study author Biamonte. [16]

A Machine Learning Systems That Called Neural Networks Perform Tasks by Analyzing Huge Volumes of Data

Neural networks learn how to carry out certain tasks by analyzing large amounts of data displayed to them. These machine learning systems continually learn and readjust to be able to carry out the task set out before them. Understanding how neural networks work helps researchers to develop better applications and uses for them.

At the 2017 Conference on Empirical Methods on Natural Language Processing earlier this month, MIT researchers demonstrated a new general-purpose technique for making sense of neural networks that are able to carry out natural language processing tasks where they attempt to extract data written in normal text opposed to something of a structured language like database-query language.

The new technique works great in any system that reads the text as input and produces symbols as the output. One such example of this can be seen in an automatic translator. It works without the need to access any underlying software too. Tommi Jaakkola is Professor of Electrical Engineering and Computer Science at MIT and one of the authors on the paper. He says, "I can't just do a simple randomization. And what you are predicting is now a more complex object, like a sentence, so what does it mean to give an explanation?"

As part of the research, Jaakkola, and colleague David Alvarez-Melis, an MIT graduate student in electrical engineering and computer science and first author on the paper, used a black-box neural net in which to generate test sentences to feed black-box neural nets. The duo began by teaching the network to compress and decompress natural sentences. As the training continues the encoder and decoder get evaluated simultaneously depending on how closely the decoder's output matches up with the encoder's input.

Neural nets work on probabilities. For example, an object-recognition system could be fed an image of a cat, and it would process that image as it saying 75 percent probability of being a cat, while still having a 25 percent probability that it's a dog. Along with that same line, Jaakkola and Alvarez-Melis' sentence compressing network has alternative words for each of those in a decoded sentence along with the probability that each is correct. So, once the system has generated a list of closely related sentences they're then fed to a black-box natural language processor. This then allows the researchers to analyze and determine which inputs have an effect on which outputs.

During the research, the pair applied this technique to three different types of a natural language processing system. The first one inferred the way in which words were pronounced; the second was a set of translators, and the third was a simple computer dialogue system which tried to provide adequate responses to questions or remarks. In looking at the results, it was clear and pretty obvious that the translation systems had strong dependencies on individual words of both the input and output sentences. A little more surprising, however, was the identification of gender

biases in the texts on which the machine translation systems were trained. The dialogue system was too small to take advantage of the training set.

“The other experiment we do is in flawed systems,” says Alvarez-Melis. “If you have a black-box model that is not doing a good job, can you first use this kind of approach to identify problems? A motivating application of this kind of interpretability is to fix systems, to improve systems, by understanding what they’re getting wrong and why.” [15]

Active machine learning for the discovery and crystallization of gigantic polyoxometalate molecules

Who is the better experimentalist, a human or a robot? When it comes to exploring synthetic and crystallization conditions for inorganic gigantic molecules, actively learning machines are clearly ahead, as demonstrated by British Scientists in an experiment with polyoxometalates published in the journal *Angewandte Chemie*.

Polyoxometalates form through self-assembly of a large number of metal atoms bridged by oxygen atoms. Potential uses include catalysis, electronics, and medicine. Insights into the self-organization processes could also be of use in developing functional chemical systems like "molecular machines".

Polyoxometalates offer a nearly unlimited variety of structures. However, it is not easy to find new ones, because the aggregation of complex inorganic molecules to gigantic molecules is a process that is difficult to predict. It is necessary to find conditions under which the building blocks aggregate and then also crystallize, so that they can be characterized.

A team led by Leroy Cronin at the University of Glasgow (UK) has now developed a new approach to define the range of suitable conditions for the synthesis and crystallization of polyoxometalates. It is based on recent advances in machine learning, known as active learning. They allowed their trained machine to compete against the intuition of experienced experimenters. The test example was $\text{Na}(6)[\text{Mo}(120)\text{Ce}(6)\text{O}(366)\text{H}(12)(\text{H}(2)\text{O})(78)] \cdot 200 \text{H}(2)\text{O}$, a new, ring-shaped polyoxometalate cluster that was recently discovered by the researchers' automated chemical robot.

In the experiment, the relative quantities of the three necessary reagent solutions were to be varied while the protocol was otherwise prescribed. The starting point was a set of data from successful and unsuccessful crystallization experiments. The aim was to plan ten experiments and then use the results from these to proceed to the next set of ten experiments - a total of one hundred crystallization attempts.

Although the flesh-and-blood experimenters were able to produce more successful crystallizations, the far more "adventurous" machine algorithm was superior on balance because it covered a significantly broader domain of the "crystallization space". The quality of the prediction of whether an experiment would lead to crystallization was improved significantly more by the machine than the human experimenters. A series of 100 purely random experiments resulted in no improvement. In addition, the machine discovered a range of conditions that led to crystals which would not have been expected based on pure intuition. This "unbiased" automated method makes the discovery of

novel compounds more probably than reliance on human intuition. The researchers are now looking for ways to make especially efficient "teams" of man and machine. [14]

Using machine learning to understand materials

Whether you realize it or not, machine learning is making your online experience more efficient. The technology, designed by computer scientists, is used to better understand, analyze, and categorize data. When you tag your friend on Facebook, clear your spam filter, or click on a suggested YouTube video, you're benefitting from machine learning algorithms.

Machine learning algorithms are designed to improve as they encounter more data, making them a versatile technology for understanding large sets of photos such as those accessible from Google Images. Elizabeth Holm, professor of materials science and engineering at Carnegie Mellon University, is leveraging this technology to better understand the enormous number of research images accumulated in the field of materials science. This unique application is an interdisciplinary approach to machine learning that hasn't been explored before.

"Just like you might search for cute cat pictures on the internet, or Facebook recognizes the faces of your friends, we are creating a system that allows a computer to automatically understand the visual data of materials science," explains Holm.

The field of materials science usually relies on human experts to identify research images by hand. Using machine learning algorithms, Holm and her group have created a system that automatically recognizes and categorizes microstructural images of materials. Her goal is to make it more efficient for materials scientists to search, sort, classify, and identify important information in their visual data.

"In materials science, one of our fundamental data is pictures," explains Holm. "Images contain information that we recognize, even when we find it difficult to quantify numerically."

Holm's machine learning system has several different applications within the materials science field including research, industry, publishing, and academia. For example, the system could be used to create a visual search of a scientific journal archives so that a researcher could find out whether a similar image had ever been published. Similarly, the system can be used to automatically search and categorize image archives in industries or research labs. "Big companies can have archives of 600,000 or more research images. No one wants to look through those, but they want to use that data to better understand their products," explains Holm. "This system has the power to unlock those archives."

Holm and her group have been working on this research for about three years and are continuing to grow the project, especially as it relates to the metal 3-D printing field. For example, they are beginning to compile a database of experimental and simulated metal powder micrographs in order to better understand what types of raw materials are best suited for 3-D printing processes.

Holm published an article about this research in the December 2015 issue of Computational Materials Science titled "A computer vision approach for automated analysis and classification of microstructural image data." [13]

Artificial intelligence helps in the discovery of new materials

With the help of artificial intelligence, chemists from the University of Basel in Switzerland have computed the characteristics of about two million crystals made up of four chemical elements. The researchers were able to identify 90 previously unknown thermodynamically stable crystals that can be regarded as new materials.

They report on their findings in the scientific journal *Physical Review Letters*.

Elpasolite is a glassy, transparent, shiny and soft mineral with a cubic crystal structure. First discovered in El Paso County (Colorado, USA), it can also be found in the Rocky Mountains, Virginia and the Apennines (Italy). In experimental databases, elpasolite is one of the most frequently found quaternary crystals (crystals made up of four chemical elements). Depending on its composition, it can be a metallic conductor, a semi-conductor or an insulator, and may also emit light when exposed to radiation.

These characteristics make elpasolite an interesting candidate for use in scintillators (certain aspects of which can already be demonstrated) and other applications. Its chemical complexity means that, mathematically speaking, it is practically impossible to use quantum mechanics to predict every theoretically viable combination of the four elements in the structure of elpasolite.

Machine learning aids statistical analysis

Thanks to modern artificial intelligence, Felix Faber, a doctoral student in Prof. Anatole von Lilienfeld's group at the University of Basel's Department of Chemistry, has now succeeded in solving this material design problem. First, using quantum mechanics, he generated predictions for thousands of elpasolite crystals with randomly determined chemical compositions. He then used the results to train statistical machine learning models (ML models). The improved algorithmic strategy achieved a predictive accuracy equivalent to that of standard quantum mechanical approaches.

ML models have the advantage of being several orders of magnitude quicker than corresponding quantum mechanical calculations. Within a day, the ML model was able to predict the formation energy – an indicator of chemical stability – of all two million elpasolite crystals that theoretically can be obtained from the main group elements of the periodic table. In contrast, performance of the calculations by quantum mechanical means would have taken a supercomputer more than 20 million hours.

Unknown materials with interesting characteristics

An analysis of the characteristics computed by the model offers new insights into this class of materials. The researchers were able to detect basic trends in formation energy and identify 90 previously unknown crystals that should be thermodynamically stable, according to quantum mechanical predictions.

On the basis of these potential characteristics, elpasolite has been entered into the Materials Project material database, which plays a key role in the Materials Genome Initiative. The initiative was launched by the US government in 2011 with the aim of using computational support to accelerate the discovery and the experimental synthesis of interesting new materials.

Some of the newly discovered elpasolite crystals display exotic electronic characteristics and unusual compositions. "The combination of artificial intelligence, big data, quantum mechanics and supercomputing opens up promising new avenues for deepening our understanding of materials and discovering new ones that we would not consider if we relied solely on human intuition," says study director von Lilienfeld. [12]

Physicists are putting themselves out of a job, using artificial intelligence to run a complex experiment

The experiment, developed by physicists from The Australian National University (ANU) and UNSW ADFA, created an extremely cold gas trapped in a laser beam, known as a Bose-Einstein condensate, replicating the experiment that won the 2001 Nobel Prize.

"I didn't expect the machine could learn to do the experiment itself, from scratch, in under an hour," said co-lead researcher Paul Wigley from the ANU Research School of Physics and Engineering.

"A simple computer program would have taken longer than the age of the Universe to run through all the combinations and work this out."

Bose-Einstein condensates are some of the coldest places in the Universe, far colder than outer space, typically less than a billionth of a degree above absolute zero.

They could be used for mineral exploration or navigation systems as they are extremely sensitive to external disturbances, which allows them to make very precise measurements such as tiny changes in the Earth's magnetic field or gravity.

The artificial intelligence system's ability to set itself up quickly every morning and compensate for any overnight fluctuations would make this fragile technology much more useful for field measurements, said co-lead researcher Dr Michael Hush from UNSW ADFA.

"You could make a working device to measure gravity that you could take in the back of a car, and the artificial intelligence would recalibrate and fix itself no matter what," he said.

"It's cheaper than taking a physicist everywhere with you."

The team cooled the gas to around 1 microkelvin, and then handed control of the three laser beams over to the artificial intelligence to cool the trapped gas down to nanokelvin.

Researchers were surprised by the methods the system came up with to ramp down the power of the lasers.

"It did things a person wouldn't guess, such as changing one laser's power up and down, and compensating with another," said Mr Wigley.

"It may be able to come up with complicated ways humans haven't thought of to get experiments colder and make measurements more precise.

The new technique will lead to bigger and better experiments, said Dr Hush.

"Next we plan to employ the artificial intelligence to build an even larger Bose-Einstein condensate faster than we've seen ever before," he said.

The research is published in the Nature group journal Scientific Reports. [11]

Quantum experiments designed by machines

The idea was developed when the physicists wanted to create new quantum states in the laboratory, but were unable to conceive of methods to do so. "After many unsuccessful attempts to come up with an experimental implementation, we came to the conclusion that our intuition about these phenomena seems to be wrong. We realized that in the end we were just trying random arrangements of quantum building blocks. And that is what a computer can do as well - but thousands of times faster", explains Mario Krenn, PhD student in Anton Zeilinger's group and first author research.

After a few hours of calculation, their algorithm - which they call Melvin - found the recipe to the question they were unable to solve, and its structure surprised them. Zeilinger says: "Suppose I want build an experiment realizing a specific quantum state I am interested in. Then humans intuitively consider setups reflecting the symmetries of the state. Yet Melvin found out that the most simple realization can be asymmetric and therefore counterintuitive. A human would probably never come up with that solution."

The physicists applied the idea to several other questions and got dozens of new and surprising answers. "The solutions are difficult to understand, but we were able to extract some new experimental tricks we have not thought of before. Some of these computer-designed experiments are being built at the moment in our laboratories", says Krenn.

Melvin not only tries random arrangements of experimental components, but also learns from previous successful attempts, which significantly speeds up the discovery rate for more complex solutions. In the future, the authors want to apply their algorithm to even more general questions in quantum physics, and hope it helps to investigate new phenomena in laboratories. [10]

Moving electrons around loops with light: A quantum device based on geometry

Researchers at the University of Chicago's Institute for Molecular Engineering and the University of Konstanz have demonstrated the ability to generate a quantum logic operation, or rotation of the qubit, that - surprisingly—is intrinsically resilient to noise as well as to variations in the strength or duration of the control. Their achievement is based on a geometric concept known as the Berry phase and is implemented through entirely optical means within a single electronic spin in diamond.

Their findings were published online Feb. 15, 2016, in Nature Photonics and will appear in the March print issue. "We tend to view quantum operations as very fragile and susceptible to noise, especially when compared to conventional electronics," remarked David Awschalom, the Liew Family Professor of Molecular Engineering and senior scientist at Argonne National Laboratory,

who led the research. "In contrast, our approach shows incredible resilience to external influences and fulfills a key requirement for any practical quantum technology."

Quantum geometry

When a quantum mechanical object, such as an electron, is cycled along some loop, it retains a memory of the path that it travelled, the Berry phase. To better understand this concept, the Foucault pendulum, a common staple of science museums helps to give some intuition. A pendulum, like those in a grandfather clock, typically oscillates back and forth within a fixed plane. However, a Foucault pendulum oscillates along a plane that gradually rotates over the course of a day due to Earth's rotation, and in turn knocks over a series of pins encircling the pendulum.

The number of knocked-over pins is a direct measure of the total angular shift of the pendulum's oscillation plane, its acquired geometric phase. Essentially, this shift is directly related to the location of the pendulum on Earth's surface as the rotation of Earth transports the pendulum along a specific closed path, its circle of latitude. While this angular shift depends on the particular path traveled, Awschalom said, it remarkably does not depend on the rotational speed of Earth or the oscillation frequency of the pendulum.

"Likewise, the Berry phase is a similar path-dependent rotation of the internal state of a quantum system, and it shows promise in quantum information processing as a robust means to manipulate qubit states," he said.

A light touch

In this experiment, the researchers manipulated the Berry phase of a quantum state within a nitrogen-vacancy (NV) center, an atomic-scale defect in diamond. Over the past decade and a half, its electronic spin state has garnered great interest as a potential qubit. In their experiments, the team members developed a method with which to draw paths for this defect's spin by varying the applied laser light. To demonstrate Berry phase, they traced loops similar to that of a tangerine slice within the quantum space of all of the potential combinations of spin states.

"Essentially, the area of the tangerine slice's peel that we drew dictated the amount of Berry phase that we were able to accumulate," said Christopher Yale, a postdoctoral scholar in Awschalom's laboratory, and one of the co-lead authors of the project.

This approach using laser light to fully control the path of the electronic spin is in contrast to more common techniques that control the NV center spin, through the application of microwave fields. Such an approach may one day be useful in developing photonic networks of these defects, linked and controlled entirely by light, as a way to both process and transmit quantum information.

A noisy path

A key feature of Berry phase that makes it a robust quantum logic operation is its resilience to noise sources. To test the robustness of their Berry phase operations, the researchers intentionally added noise to the laser light controlling the path. As a result, the spin state would travel along its intended path in an erratic fashion.

However, as long as the total area of the path remained the same, so did the Berry phase that they measured.

"In particular, we found the Berry phase to be insensitive to fluctuations in the intensity of the laser. Noise like this is normally a bane for quantum control," said Brian Zhou, a postdoctoral scholar in the group, and co-lead author.

"Imagine you're hiking along the shore of a lake, and even though you continually leave the path to go take pictures, you eventually finish hiking around the lake," said F. Joseph Heremans, co-lead author, and now a staff scientist at Argonne National Laboratory. "You've still hiked the entire loop regardless of the bizarre path you took, and so the area enclosed remains virtually the same."

These optically controlled Berry phases within diamond suggest a route toward robust and fault-tolerant quantum information processing, noted Guido Burkard, professor of physics at the University of Konstanz and theory collaborator on the project.

"Though its technological applications are still nascent, Berry phases have a rich underlying mathematical framework that makes them a fascinating area of study," Burkard said. [9]

Researchers demonstrate 'quantum surrealism'

In a new version of an old experiment, CIFAR Senior Fellow Aephraim Steinberg (University of Toronto) and colleagues tracked the trajectories of photons as the particles traced a path through one of two slits and onto a screen. But the researchers went further, and observed the "nonlocal" influence of another photon that the first photon had been entangled with.

The results counter a long-standing criticism of an interpretation of quantum mechanics called the De Broglie-Bohm theory. Detractors of this interpretation had faulted it for failing to explain the behaviour of entangled photons realistically. For Steinberg, the results are important because they give us a way of visualizing quantum mechanics that's just as valid as the standard interpretation, and perhaps more intuitive.

"I'm less interested in focusing on the philosophical question of what's 'really' out there. I think the fruitful question is more down to earth. Rather than thinking about different metaphysical interpretations, I would phrase it in terms of having different pictures. Different pictures can be useful. They can help shape better intuitions."

At stake is what is "really" happening at the quantum level. The uncertainty principle tells us that we can never know both a particle's position and momentum with complete certainty. And when we do interact with a quantum system, for instance by measuring it, we disturb the system. So if we fire a photon at a screen and want to know where it will hit, we'll never know for sure exactly where it will hit or what path it will take to get there.

The standard interpretation of quantum mechanics holds that this uncertainty means that there is no "real" trajectory between the light source and the screen. The best we can do is to calculate a "wave function" that shows the odds of the photon being in any one place at any time, but won't tell us where it is until we make a measurement.

Yet another interpretation, called the De Broglie-Bohm theory, says that the photons do have real trajectories that are guided by a "pilot wave" that accompanies the particle. The wave is still probabilistic, but the particle takes a real trajectory from source to target. It doesn't simply "collapse" into a particular location once it's measured.


In 2011 Steinberg and his colleagues showed that they could follow trajectories for photons by subjecting many identical particles to measurements so weak that the particles were barely disturbed, and then averaging out the information. This method showed trajectories that looked similar to classical ones - say, those of balls flying through the air.

But critics had pointed out a problem with this viewpoint. Quantum mechanics also tells us that two particles can be entangled, so that a measurement of one particle affects the other. The critics complained that in some cases, a measurement of one particle would lead to an incorrect prediction of the trajectory of the entangled particle. They coined the term "surreal trajectories" to describe them.

In the most recent experiment, Steinberg and colleagues showed that the surrealism was a consequence of non-locality - the fact that the particles were able to influence one another instantaneously at a distance. In fact, the "incorrect" predictions of trajectories by the entangled photon were actually a consequence of where in their course the entangled particles were measured. Considering both particles together, the measurements made sense and were consistent with real trajectories.

Steinberg points out that both the standard interpretation of quantum mechanics and the De Broglie-Bohm interpretation are consistent with experimental evidence, and are mathematically equivalent. But it is helpful in some circumstances to visualize real trajectories, rather than wave function collapses, he says. [8]

Physicists discover easy way to measure entanglement—on a sphere



Entanglement on a sphere: This Bloch sphere shows entanglement for the one-root state ρ and its radial state ρ_c . The color on the sphere corresponds to the value of the entanglement, which is determined by the distance from the root state z , the point at which there is no entanglement. The

closer to z , the less the entanglement (red); the further from z , the greater the entanglement (blue). Credit: Regula and Adesso. ©2016 American Physical Society

Now in a new paper to be published in *Physical Review Letters*, mathematical physicists Bartosz Regula and Gerardo Adesso at The University of Nottingham have greatly simplified the problem of measuring entanglement.

To do this, the scientists turned the difficult analytical problem into an easy geometrical one. They showed that, in many cases, the amount of entanglement between states corresponds to the distance between two points on a Bloch sphere, which is basically a normal 3D sphere that physicists use to model quantum states.

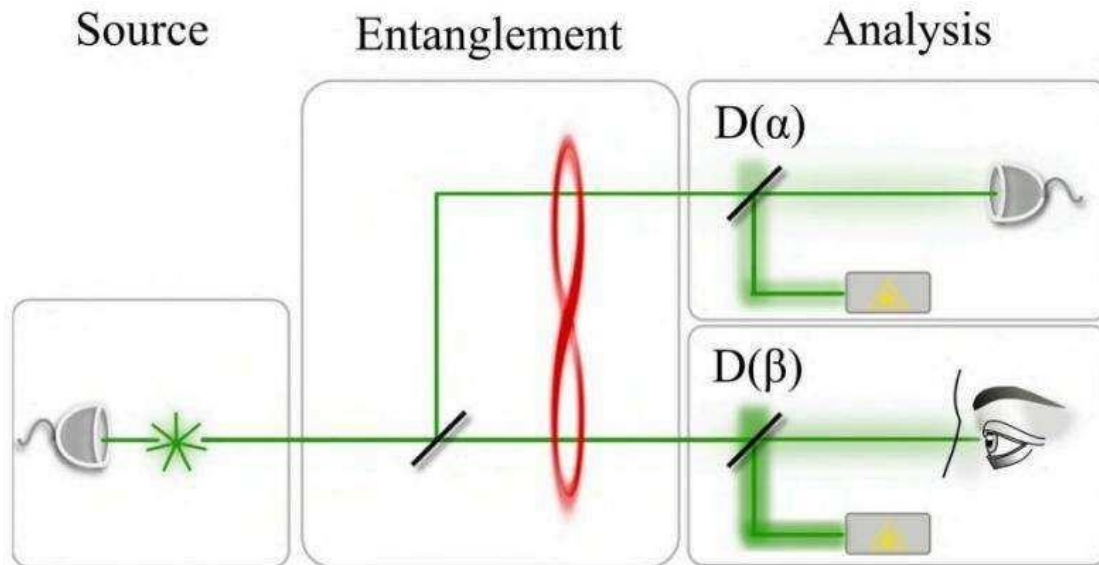
As the scientists explain, the traditionally difficult part of the math problem is that it requires finding the optimal decomposition of mixed states into pure states. The geometrical approach completely eliminates this requirement by reducing the many possible ways that states could decompose down to a single point on the sphere at which there is zero entanglement. The approach requires that there be only one such point, or "root," of zero entanglement, prompting the physicists to describe the method as "one root to rule them all."

The scientists explain that the "one root" property is common among quantum states and can be easily verified, transforming a formidable math problem into one that is trivially easy. They demonstrated that the new approach works for many types of two-, three- and four-qubit entangled states.

"This method reveals an intriguing and previously unexplored connection between the quantum features of a state and classical geometry, allowing all one-root states to enjoy a convenient visual representation which considerably simplifies the study and understanding of their properties," the researchers explained.

The simple way of measuring a state's entanglement could have applications in many technological areas, such as quantum cryptography, computation, and communication. It could also provide insight into understanding the foundations of thermodynamics, condensed matter physics, and biology. [7]

An idea for allowing the human eye to observe an instance of entanglement



Scheme of the proposal for detecting entanglement with the human eye. Credit: arXiv:1602.01907

Entanglement, is of course, where two quantum particles are intrinsically linked to the extent that they actually share the same existence, even though they can be separated and moved apart. The idea was first proposed nearly a century ago, and it has not only been proven, but researchers routinely cause it to occur, but, to date, not one single person has ever actually seen it happen—they only know it happens by conducting a series of experiments. It is not clear if anyone has ever actually tried to see it happen, but in this new effort, the research trio claim to have found a way to make it happen—if only someone else will carry out the experiment on a willing volunteer.

The idea involves using a beam splitter and two beams of light—an initial beam of coherent photons fired at the beam splitter and a secondary beam of coherent photons that interferes with the photons in the first beam causing a change of phase, forcing the light to be reflected rather than transmitted. In such a scenario, the secondary beam would not need to be as intense as the first, and could in fact be just a single coherent photon—if it were entangled, it could be used to allow a person to see the more powerful beam while still preserving the entanglement of the original photon.

The researchers suggest the technology to carry out such an experiment exists today, but also acknowledge that it would take a special person to volunteer for such an assignment because to prove that they had seen entanglement taking place would involve shooting a large number of photons in series, into a person's eye, whereby the resolute volunteer would announce whether they had seen the light on the order of thousands of times. [6]

Quantum entanglement

Measurements of physical properties such as position, momentum, spin, polarization, etc.

performed on entangled particles are found to be appropriately correlated. For example, if a pair of particles is generated in such a way that their total spin is known to be zero, and one particle is found to have clockwise spin on a certain axis, then the spin of the other particle, measured on the same axis, will be found to be counterclockwise. Because of the nature of quantum measurement, however, this behavior gives rise to effects that can appear paradoxical: any measurement of a property of a particle can be seen as acting on that particle (e.g. by collapsing a number of superimposed states); and in the case of entangled particles, such action must be on the entangled system as a whole. It thus appears that one particle of an entangled pair "knows" what measurement has been performed on the other, and with what outcome, even though there is no known means for such information to be communicated between the particles, which at the time of measurement may be separated by arbitrarily large distances. [4]

The Bridge

The accelerating electrons explain not only the Maxwell Equations and the Special Relativity, but the Heisenberg Uncertainty Relation, the wave particle duality and the electron's spin also, building the bridge between the Classical and Quantum Theories. [1]

Accelerating charges

The moving charges are self maintain the electromagnetic field locally, causing their movement and this is the result of their acceleration under the force of this field. In the classical physics the charges will distributed along the electric current so that the electric potential lowering along the current, by linearly increasing the way they take every next time period because this accelerated motion. The same thing happens on the atomic scale giving a dp impulse difference and a dx way difference between the different part of the not point like particles.

Relativistic effect

Another bridge between the classical and quantum mechanics in the realm of relativity is that the charge distribution is lowering in the reference frame of the accelerating charges linearly: $ds/dt = at$ (time coordinate), but in the reference frame of the current it is parabolic: $s = a/2 t^2$ (geometric coordinate).

Heisenberg Uncertainty Relation

In the atomic scale the Heisenberg uncertainty relation gives the same result, since the moving electron in the atom accelerating in the electric field of the proton, causing a charge distribution on Δx position difference and with a Δp momentum difference such a way that they product is about the half Planck reduced constant. For the proton this Δx much less in the nucleon, than in the orbit of the electron in the atom, the Δp is much higher because of the greater proton mass.

This means that the electron and proton are not point like particles, but has a real charge distribution.

Wave – Particle Duality

The accelerating electrons explains the wave – particle duality of the electrons and photons, since the elementary charges are distributed on Δx position with Δp impulse and creating a wave packet of the electron. The photon gives the electromagnetic particle of the mediating force of the electrons electromagnetic field with the same distribution of wavelengths.

Atomic model

The constantly accelerating electron in the Hydrogen atom is moving on the equipotential line of the proton and its kinetic and potential energy will be constant. Its energy will change only when it is changing its way to another equipotential line with another value of potential energy or getting free with enough kinetic energy. This means that the Rutherford-Bohr atomic model is right and only that changing acceleration of the electric charge causes radiation, not the steady acceleration. The steady acceleration of the charges only creates a centric parabolic steady electric field around the charge, the magnetic field. This gives the magnetic moment of the atoms, summing up the proton and electron magnetic moments caused by their circular motions and spins.

The Relativistic Bridge

Commonly accepted idea that the relativistic effect on the particle physics is the fermions' spin - another unresolved problem in the classical concepts. If the electric charges can move only with accelerated motions in the self maintaining electromagnetic field, once upon a time they would reach the velocity of the electromagnetic field. The resolution of this problem is the spinning particle, constantly accelerating and not reaching the velocity of light because the acceleration is radial. One origin of the Quantum Physics is the Planck Distribution Law of the electromagnetic oscillators, giving equal intensity for 2 different wavelengths on any temperature. Any of these two wavelengths will give equal intensity diffraction patterns, building different asymmetric constructions, for example proton - electron structures (atoms), molecules, etc. Since the particles are centers of diffraction patterns they also have particle – wave duality as the electromagnetic waves have. [2]

The weak interaction

The weak interaction transforms an electric charge in the diffraction pattern from one side to the other side, causing an electric dipole momentum change, which violates the CP and time reversal symmetry. The Electroweak Interaction shows that the Weak Interaction is basically electromagnetic in nature. The arrow of time shows the entropy grows by changing the temperature dependent diffraction patterns of the electromagnetic oscillators.

Another important issue of the quark model is when one quark changes its flavor such that a linear oscillation transforms into plane oscillation or vice versa, changing the charge value with 1 or -1.

This kind of change in the oscillation mode requires not only parity change, but also charge and time changes (CPT symmetry) resulting a right handed anti-neutrino or a left handed neutrino.

The right handed anti-neutrino and the left handed neutrino exist only because changing back the quark flavor could happen only in reverse, because they are different geometrical constructions, the u is 2 dimensional and positively charged and the d is 1 dimensional and negatively charged. It needs also a time reversal, because anti particle (anti neutrino) is involved.

The neutrino is a $1/2$ spin creator particle to make equal the spins of the weak interaction, for example neutron decay to 2 fermions, every particle is fermions with $1/2$ spin. The weak interaction changes the entropy since more or less particles will give more or less freedom of movement. The entropy change is a result of temperature change and breaks the equality of oscillator diffraction intensity of the Maxwell–Boltzmann statistics. This way it changes the time coordinate measure and makes possible a different time dilation as of the special relativity.

The limit of the velocity of particles as the speed of light appropriate only for electrical charged particles, since the accelerated charges are self maintaining locally the accelerating electric force. The neutrinos are CP symmetry breaking particles compensated by time in the CPT symmetry, that is the time coordinate not works as in the electromagnetic interactions, consequently the speed of neutrinos is not limited by the speed of light.

The weak interaction T-asymmetry is in conjunction with the T-asymmetry of the second law of thermodynamics, meaning that locally lowering entropy (on extremely high temperature) causes the weak interaction, for example the Hydrogen fusion.

Probably because it is a spin creating movement changing linear oscillation to 2 dimensional oscillation by changing d to u quark and creating anti neutrino going back in time relative to the proton and electron created from the neutron, it seems that the anti neutrino fastest then the velocity of the photons created also in this weak interaction?

A quark flavor changing shows that it is a reflection changes movement and the CP- and T-symmetry breaking!!! This flavor changing oscillation could prove that it could be also on higher level such as atoms, molecules, probably big biological significant molecules and responsible on the aging of the life.

Important to mention that the weak interaction is always contains particles and antiparticles, where the neutrinos (antineutrinos) present the opposite side. It means by Feynman's interpretation that these particles present the backward time and probably because this they seem to move faster than the speed of light in the reference frame of the other side.

Finally since the weak interaction is an electric dipole change with $1/2$ spin creating; it is limited by the velocity of the electromagnetic wave, so the neutrino's velocity cannot exceed the velocity of light.

The General Weak Interaction

The Weak Interactions T-asymmetry is in conjunction with the T-asymmetry of the Second Law of Thermodynamics, meaning that locally lowering entropy (on extremely high temperature) causes for example the Hydrogen fusion. The arrow of time by the Second Law of Thermodynamics shows the increasing entropy and decreasing information by the Weak Interaction, changing the temperature dependent diffraction patterns. A good example of this is the neutron decay, creating more particles with less known information about them.

The neutrino oscillation of the Weak Interaction shows that it is a general electric dipole change and it is possible to any other temperature dependent entropy and information changing diffraction pattern of atoms, molecules and even complicated biological living structures.

We can generalize the weak interaction on all of the decaying matter constructions, even on the biological too. This gives the limited lifetime for the biological constructions also by the arrow of time. There should be a new research space of the Quantum Information Science the 'general neutrino oscillation' for the greater than subatomic matter structures as an electric dipole change. There is also connection between statistical physics and evolutionary biology, since the arrow of time is working in the biological evolution also.

The Fluctuation Theorem says that there is a probability that entropy will flow in a direction opposite to that dictated by the Second Law of Thermodynamics. In this case the Information is growing that is the matter formulas are emerging from the chaos. So the Weak Interaction has two directions, samples for one direction is the Neutron decay, and Hydrogen fusion is the opposite direction.

Fermions and Bosons

The fermions are the diffraction patterns of the bosons such a way that they are both sides of the same thing.

Van Der Waals force

Named after the Dutch scientist Johannes Diderik van der Waals – who first proposed it in 1873 to explain the behaviour of gases – it is a very weak force that only becomes relevant when atoms and molecules are very close together. Fluctuations in the electronic cloud of an atom mean that it will have an instantaneous dipole moment. This can induce a dipole moment in a nearby atom, the result being an attractive dipole–dipole interaction.

Electromagnetic inertia and mass

Electromagnetic Induction

Since the magnetic induction creates a negative electric field as a result of the changing acceleration, it works as an electromagnetic inertia, causing an electromagnetic mass. [1]

Relativistic change of mass

The increasing mass of the electric charges the result of the increasing inductive electric force acting against the accelerating force. The decreasing mass of the decreasing acceleration is the result of the inductive electric force acting against the decreasing force. This is the relativistic mass

change explanation, especially importantly explaining the mass reduction in case of velocity decrease.

The frequency dependence of mass

Since $E = h\nu$ and $E = mc^2$, $m = h\nu/c^2$ that is the m depends only on the ν frequency. It means that the mass of the proton and electron are electromagnetic and the result of the electromagnetic induction, caused by the changing acceleration of the spinning and moving charge! It could be that the m_0 inertial mass is the result of the spin, since this is the only accelerating motion of the electric charge. Since the accelerating motion has different frequency for the electron in the atom and the proton, they masses are different, also as the wavelengths on both sides of the diffraction pattern, giving equal intensity of radiation.

Electron – Proton mass rate

The Planck distribution law explains the different frequencies of the proton and electron, giving equal intensity to different lambda wavelengths! Also since the particles are diffraction patterns they have some closeness to each other – can be seen as a gravitational force. [2]

There is an asymmetry between the mass of the electric charges, for example proton and electron, can understood by the asymmetrical Planck Distribution Law. This temperature dependent energy distribution is asymmetric around the maximum intensity, where the annihilation of matter and antimatter is a high probability event. The asymmetric sides are creating different frequencies of electromagnetic radiations being in the same intensity level and compensating each other. One of these compensating ratios is the electron – proton mass ratio. The lower energy side has no compensating intensity level, it is the dark energy and the corresponding matter is the dark matter.

Gravity from the point of view of quantum physics

The Gravitational force

The gravitational attractive force is basically a magnetic force.

The same electric charges can attract one another by the magnetic force if they are moving parallel in the same direction. Since the electrically neutral matter is composed of negative and positive charges they need 2 photons to mediate this attractive force, one per charges. The Bing Bang caused parallel moving of the matter gives this magnetic force, experienced as gravitational force.

Since graviton is a tensor field, it has spin = 2, could be 2 photons with spin = 1 together.

You can think about photons as virtual electron – positron pairs, obtaining the necessary virtual mass for gravity.

The mass as seen before a result of the diffraction, for example the proton – electron mass rate $M_p=1840 M_e$. In order to move one of these diffraction maximum (electron or proton) we need to intervene into the diffraction pattern with a force appropriate to the intensity of this diffraction maximum, means its intensity or mass.

The Big Bang caused acceleration created radial currents of the matter, and since the matter is composed of negative and positive charges, these currents are creating magnetic field and attracting forces between the parallel moving electric currents. This is the gravitational force experienced by the matter, and also the mass is result of the electromagnetic forces between the charged particles. The positive and negative charged currents attracts each other or by the magnetic forces or by the much stronger electrostatic forces!?

The gravitational force attracting the matter, causing concentration of the matter in a small space and leaving much space with low matter concentration: dark matter and energy.

There is an asymmetry between the mass of the electric charges, for example proton and electron, can understood by the asymmetrical Planck Distribution Law. This temperature dependent energy distribution is asymmetric around the maximum intensity, where the annihilation of matter and antimatter is a high probability event. The asymmetric sides are creating different frequencies of electromagnetic radiations being in the same intensity level and compensating each other. One of these compensating ratios is the electron – proton mass ratio. The lower energy side has no compensating intensity level, it is the dark energy and the corresponding matter is the dark matter.

The Higgs boson

By March 2013, the particle had been proven to behave, interact and decay in many of the expected ways predicted by the Standard Model, and was also tentatively confirmed to have + parity and zero spin, two fundamental criteria of a Higgs boson, making it also the first known scalar particle to be discovered in nature, although a number of other properties were not fully proven and some partial results do not yet precisely match those expected; in some cases data is also still awaited or being analyzed.

Since the Higgs boson is necessary to the W and Z bosons, the dipole change of the Weak interaction and the change in the magnetic effect caused gravitation must be conducted. The Wien law is also important to explain the Weak interaction, since it describes the T_{\max} change and the diffraction patterns change. [2]

Higgs mechanism and Quantum Gravity

The magnetic induction creates a negative electric field, causing an electromagnetic inertia. Probably it is the mysterious Higgs field giving mass to the charged particles? We can think about the photon as an electron-positron pair, they have mass. The neutral particles are built from negative and positive charges, for example the neutron, decaying to proton and electron. The wave – particle duality makes sure that the particles are oscillating and creating magnetic induction as an inertial mass, explaining also the relativistic mass change. Higher frequency creates stronger magnetic induction, smaller frequency results lesser magnetic induction. It seems to me that the magnetic induction is the secret of the Higgs field.

In particle physics, the Higgs mechanism is a kind of mass generation mechanism, a process that gives mass to elementary particles. According to this theory, particles gain mass by interacting with the Higgs field that permeates all space. More precisely, the Higgs mechanism endows gauge bosons in a gauge theory with mass through absorption of Nambu–Goldstone bosons arising in spontaneous symmetry breaking.

The simplest implementation of the mechanism adds an extra Higgs field to the gauge theory. The spontaneous symmetry breaking of the underlying local symmetry triggers conversion of components of this Higgs field to Goldstone bosons which interact with (at least some of) the other fields in the theory, so as to produce mass terms for (at least some of) the gauge bosons. This mechanism may also leave behind elementary scalar (spin-0) particles, known as Higgs bosons.

In the Standard Model, the phrase "Higgs mechanism" refers specifically to the generation of masses for the W^\pm , and Z weak gauge bosons through electroweak symmetry breaking. The Large Hadron Collider at CERN announced results consistent with the Higgs particle on July 4, 2012 but stressed that further testing is needed to confirm the Standard Model.

What is the Spin?

So we know already that the new particle has spin zero or spin two and we could tell which one if we could detect the polarizations of the photons produced. Unfortunately this is difficult and neither ATLAS nor CMS are able to measure polarizations. The only direct and sure way to confirm that the particle is indeed a scalar is to plot the angular distribution of the photons in the rest frame of the centre of mass. A spin zero particles like the Higgs carries no directional information away from the original collision so the distribution will be even in all directions. This test will be possible when a much larger number of events have been observed. In the mean time we can settle for less certain indirect indicators.

The Graviton

In physics, the graviton is a hypothetical elementary particle that mediates the force of gravitation in the framework of quantum field theory. If it exists, the graviton is expected to be massless (because the gravitational force appears to have unlimited range) and must be a spin-2 boson. The spin follows from the fact that the source of gravitation is the stress-energy tensor, a second-rank tensor (compared to electromagnetism's spin-1 photon, the source of which is the four-current, a first-rank tensor). Additionally, it can be shown that any massless spin-2 field would give rise to a force indistinguishable from gravitation, because a massless spin-2 field must couple to (interact with) the stress-energy tensor in the same way that the gravitational field does. This result suggests that, if a massless spin-2 particle is discovered, it must be the graviton, so that the only experimental verification needed for the graviton may simply be the discovery of a massless spin-2 particle. [3]

The Secret of Quantum Entanglement

The Secret of Quantum Entanglement that the particles are diffraction patterns of the electromagnetic waves and this way their quantum states every time is the result of the quantum state of the intermediate electromagnetic waves. [2] When one of the entangled particles wave function is collapses by measurement, the intermediate photon also collapses and transforms its

state to the second entangled particle giving it the continuity of this entanglement. Since the accelerated charges are self-maintaining their potential locally causing their acceleration, it seems that they entanglement is a spooky action at a distance.

Conclusions

The accelerated charges self-maintaining potential shows the locality of the relativity, working on the quantum level also.

The Secret of Quantum Entanglement that the particles are diffraction patterns of the electromagnetic waves and this way their quantum states every time is the result of the quantum state of the intermediate electromagnetic waves.

One of the most important conclusions is that the electric charges are moving in an accelerated way and even if their velocity is constant, they have an intrinsic acceleration anyway, the so called spin, since they need at least an intrinsic acceleration to make possible they movement .

The bridge between the classical and quantum theory is based on this intrinsic acceleration of the spin, explaining also the Heisenberg Uncertainty Principle. The particle – wave duality of the electric charges and the photon makes certain that they are both sides of the same thing. Basing the gravitational force on the accelerating Universe caused magnetic force and the Planck Distribution Law of the electromagnetic waves caused diffraction gives us the basis to build a Unified Theory of the physical interactions.

References

- [1] The Magnetic field of the Electric current and the Magnetic induction
[http://academia.edu/3833335/The Magnetic field of the Electric current](http://academia.edu/3833335/The_Magnetic_field_of_the_Electric_current)
- [2] 3 Dimensional String Theory [http://academia.edu/3834454/3 Dimensional String Theory](http://academia.edu/3834454/3_Dimensional_String_Theory)
- [3] Graviton Production By Two Photon and Electron-Photon Processes In Kaluza-Klein Theories With Large Extra Dimensions <http://arxiv.org/abs/hep-ph/9909392>
- [4] Quantum Entanglement
[http://en.wikipedia.org/wiki/Quantum entanglement](http://en.wikipedia.org/wiki/Quantum_entanglement)
- [5] Space-based experiment could test gravity's effects on quantum entanglement
<http://phys.org/news/2014-05-space-based-gravity-effects-quantum-entanglement.html>
- [6] An idea for allowing the human eye to observe an instance of entanglement
<http://phys.org/news/2016-02-idea-human-eye-instance-entanglement.html>
- [7] Physicists discover easy way to measure entanglement—on a sphere
<http://phys.org/news/2016-02-physicists-easy-entanglementon-sphere.html>

- [8] Researchers demonstrate 'quantum surrealism' <http://phys.org/news/2016-02-quantum-surrealism.html>
- [9] Moving electrons around loops with light: A quantum device based on geometry <http://phys.org/news/2016-02-electrons-loops-quantum-device-based.html>
- [10] Quantum experiments designed by machines <http://phys.org/news/2016-02-quantum-machines.html>
- [11] Physicists are putting themselves out of a job, using artificial intelligence to run a complex experiment <http://phys.org/news/2016-05-physicists-job-artificial-intelligence-complex.html>
- [12] Artificial intelligence helps in the discovery of new materials <http://phys.org/news/2016-09-artificial-intelligence-discovery-materials.html>
- [13] Using machine learning to understand materials <https://techxplore.com/news/2016-09-machine-materials.html>
- [14] Active machine learning for the discovery and crystallization of gigantic polyoxometalate molecules <https://phys.org/news/2017-08-machine-discovery-crystallization-gigantic-polyoxometalate.html>
- [15] A Machine Learning Systems That Called Neural Networks Perform Tasks by Analyzing Huge Volumes of Data <http://trendintech.com/2017/09/14/a-machine-learning-systems-that-called-neural-networksperform-tasks-by-analyzing-huge-volumes-of-data/>
- [16] Rise of the quantum thinking machines <https://phys.org/news/2017-09-quantum-machines.html>
- [17] Teaching computers to guide science: Machine learning method sees forests and trees <https://phys.org/news/2018-03-science-machine-method-forests-trees.html>
- [18] Army's brain-like computers moving closer to cracking codes <https://phys.org/news/2018-03-army-brain-like-closer-codes.html>
- [19] Dissecting artificial intelligence to better understand the human brain <https://medicalxpress.com/news/2018-03-artificial-intelligence-human-brain.html>

[20] Scientists make a maze-running artificial intelligence program that learns to take shortcuts

<https://phys.org/news/2018-05-scientists-maze-running-artificial-intelligence-shortcuts.html>

[21] Can artificial intelligence help find alien intelligence?

<https://phys.org/news/2018-05-artificial-intelligence-alien.html>

[22] The U.S. needs a national strategy on artificial intelligence

<https://phys.org/news/2018-05-national-strategy-artificial-intelligence.html>

[23] How is artificial intelligence changing science?

<https://phys.org/news/2018-05-artificial-intelligence-science.html>

[24] Are computer-aided decisions actually fair?

<https://phys.org/news/2018-12-computer-aided-decisions-fair.html>

[25] Predicting the accuracy of a neural network prior to training

<https://phys.org/news/2018-12-accuracy-neural-network-prior.html>

[26] A new technique for synthesizing motion-blurred images

<https://techxplore.com/news/2018-12-technique-motion-blurred-images.html>

[27] A new approximate computing approach using CNNs

<https://techxplore.com/news/2018-12-approximate-approach-cnns.html>

[28] Could AI take control of human birth?

<https://techxplore.com/news/2018-12-ai-human-birth.html>

[29] Why we want to build a machine that can predict a person's attractiveness

<https://phys.org/news/2019-01-machine-person.html>

[30] Artificial neural networks make life easier for hearing aid users

<https://phys.org/news/2019-01-artificial-neural-networks-life-easier.html>

[31] AI can help retailers understand the consumer

<https://phys.org/news/2019-01-ai-retailers-consumer.html>

[32] Artificial intelligence applied to the genome identifies an unknown human ancestor

<https://phys.org/news/2019-01-artificial-intelligence-genome-unknown-human.html>

[33] Artificial intelligence must know when to ask for human help

<https://phys.org/news/2019-03-artificial-intelligence-human.html>