

DISTRIBUTED ARTIFICIAL INTELLIGENCE


Blockchain as an Operating Platform for AI

Don Tapscott and Anjan Vinod

Blockchain Research Institute and Blockchain at Berkeley

March 2019





Realizing the new promise of the digital economy

In 1994, Don Tapscott coined the phrase, “the digital economy,” with his book of that title. It discussed how the Web and the Internet of information would bring important changes in business and society. Today the Internet of value creates profound new possibilities.

In 2017, Don and Alex Tapscott launched the Blockchain Research Institute to help realize the new promise of the digital economy. We research the strategic implications of blockchain technology and produce practical insights to contribute global blockchain knowledge and help our members navigate this revolution.

Our findings, conclusions, and recommendations are initially proprietary to our members and ultimately released to the public in support of our mission. To find out more, please visit www.blockchainresearchinstitute.org.



Blockchain Research Institute, 2020

Except where otherwise noted, this work is copyrighted 2020 by the Blockchain Research Institute and licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Public License. To view a copy of this license, send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA, or visit creativecommons.org/licenses/by-nc-nd/4.0/legalcode.

This document represents the views of its author(s), not necessarily those of Blockchain Research Institute or the Tapscott Group. This material is for informational purposes only; it is neither investment advice nor managerial consulting. Use of this material does not create or constitute any kind of business relationship with the Blockchain Research Institute or the Tapscott Group, and neither the Blockchain Research Institute nor the Tapscott Group is liable for the actions of persons or organizations relying on this material.

Users of this material may copy and distribute it as is under the terms of this Creative Commons license and cite it in their work. This document may contain material (photographs, figures, and tables) used with a third party’s permission or under a different Creative Commons license; and users should cite those elements separately. Otherwise, we suggest the following citation:

Don Tapscott and Anjan Vinod, “Distributed Artificial Intelligence: Blockchain as an Operating Platform for AI,” Blockchain Research Institute, 22 March 2019; adapted 10 Jan. 2020.

To request permission for remixing, transforming, building upon the material, or distributing any derivative of this material for any purpose, please contact the Blockchain Research Institute, www.blockchainresearchinstitute.org/contact-us, and put “Permission request” in subject line. Thank you for your interest!



Contents

Message from the Executive Director	3
Idea in brief	4
What is artificial intelligence?	4
Training data	6
AI and centralization	6
Blockchains and artificial intelligence	8
Immutability	9
Blockchain and democratization	11
Smart contracts	13
Blockchain governance	14
Decentralized artificial intelligence	14
Distributed AI versus decentralized AI	14
DAI and blockchain's role	15
Autonomous agents as AI representatives?	17
Looking to the future: DAO	18
Challenges with blockchains	19
Scalability	20
Governance	21
Conclusion and recommendations	21
About the authors	23
About the Brightline Initiative	24
About the Blockchain Research Institute	25
Notes	26



Message from the Executive Director

Technologies are offering the transformational means for organizations to change how they work and deliver value to customers, at a much larger scale and greater depth than ever before. Artificial intelligence (AI) is diffusing, and its enterprise use is expanding, making it one of the broadest digital revolutions ever. While we argue that AI is the most impactful technology today, blockchain has the potential to be the most impactful technology of tomorrow. Together, these technologies can help us take the next step in enterprise capabilities.

In this research project, Don Tapscott of the Blockchain Research Institute and Anjan Vinod of ParaFi Capital examine the intersection of AI and blockchain technologies. Leading organizations are already using AI and other expert systems to capture and communicate what they're learning as they manage large and complex initiatives. With blockchain platforms, they can share data securely, validate those data, automate some types of project transactions, and allow AI agents to adjust activity according to ongoing risk-based probability assessments.

Blockchain technology has the unique ability to improve the reliability of AI training data and secure AI implementations through smart contracts and on-chain governance. This project examines AI, the different kinds of AI (centralized, decentralized, and distributed), the utility of autonomous agents, and scalability and governance challenges.

It's must-read research for leaders looking to increase their knowledge and understanding of AI and blockchain technology as complements. We hope it will help leaders to leverage these technologies in overcoming challenges and improving their success in implementing strategy and transforming their businesses.

Brightline is a Project Management Institute (PMI) initiative together with leading global organizations, dedicated to helping executives bridge the expensive and unproductive gap between strategy design and delivery. We are excited to share this research on blockchain technology as a result of an exclusive partnership with the Blockchain Research Institute, the leading think tank on the study of the strategic implications of blockchain in organizations, industries, and our society.

A handwritten signature in black ink that reads 'Ricardo Viana Vargas'.

 RICARDO VIANA VARGAS
*Executive Director
Brightline Initiative*



Seventy-two percent of business leaders see AI as being fundamental for the future.

Idea in brief

- » *Artificial intelligence (AI)*, although heavily used as a buzzword, is emerging as one of the biggest technological revolutions ever. The ability for computing systems to absorb and interpret large data sets in a meaningful way evolved in the 20th century and has advanced rapidly for about the last 10 years. The biggest disruptions are yet to come.
- » With a market size projected to reach \$70 billion by 2020, the opportunity to discern actionable insights from data using AI will transform a variety of industries from transportation to healthcare, finance, manufacturing, retail, and so many more.¹
- » As adept as the human mind is at finding opportunities, systems now have the ability to pore through thousands of times more data in a way that ultimately will augment our decision-making skills. Seventy-two percent of business leaders see AI as being fundamental for the future.²
- » As powerful as AI is becoming, blockchain technology has the unique ability to push the AI revolution forward through a decentralized, trustless set of features. In this paper, we look at what AI is exactly, and what are some of its most powerful aspects.
- » We look at centralized AI solutions, the solutions blockchain technology could provide, and the issues posed by these emerging fields. We then analyze what the power of decentralized AI could unveil.
- » Blockchain has the potential to be the operating platform for AI, whether for existing centralized entities or for emerging uses of decentralized AI.

What is artificial intelligence?

At an extremely high level, artificial intelligence is the design of intelligent systems capable of learning and simulating human cognition. The field of AI encompasses a wide variety of techniques including machine learning, natural language processing, neural nets, deep learning, and many others.




AI-based enterprise applications are projected to grow from \$1.6 billion in annual revenue to \$31.2 billion by 2025.

We can divide AI into two classes: the first is symbolic AI, and the second is *machine learning* (ML), also known as *connectionism*. The former uses mathematical logic to formalize the complex tasks that AI machines will perform. For tasks too difficult to formalize, such as pattern recognition—say, a set of rules to distinguish a picture of a cat from that of a dog—machine learning is particularly powerful. It is a set of techniques that model cognitive processes directly from experiences. Given a large set of examples, machine learning algorithms can perform complex tasks (e.g., market predictions) by identifying patterns in the data set.

We are surrounded by tangible AI applications daily. In fact, 77 percent of consumers use an AI-powered service or device.³ AI-based enterprise applications are projected to grow from \$1.6 billion in annual revenue to \$31.2 billion by 2025. This estimated 52 percent compounded annual growth rate is fueled by such applications as image recognition, recruitment, data analytics, and cybersecurity. With investment in AI start-ups increasing sixfold since 2000, the amount of attention, money, and development in the space is unprecedented.⁴ Enterprise adoption of AI is growing rapidly as well, with a reported 61 percent of businesses implementing AI in 2018.⁵

Recommendations on Netflix and targeted ads on Facebook are both supported by AI-powered systems in some manner. IBM has deployed Watson, its AI platform and *Jeopardy!* game show champion, for a variety of companies.⁶ Korean Air leveraged Watson to analyze years' worth of maintenance records for its fleet, cutting down the maintenance analysis time by 90 percent.⁷ Amazon leverages machine learning to optimize the thousands of robots it uses for fulfillment and to map voice commands to real-time execution in Alexa.⁸

There are various AI-powered applications in the market today. How does AI actually work? How are computers becoming "intelligent"?



Board Electronics Computer Electrical Engineering by Michael Schwarzenberger (blickpixel), 2014, used under Pixabay license, as of 10 March 2019.



Many of our most ambitious societal objectives rely on the progress of machine learning. AI is here to stay, and while the revolution is just beginning, blockchain technology may help accelerate its growth.

Training data

The answer lies in the computing systems' ability to interpret and extract meaning from large data sets. These data sets, sometimes called *training data*, serve as an input into a system. The system interprets the data thoroughly, ideally to provide a decision or outcome that achieves some goal. The power lies in *machine learning* (ML), a subset of AI, which trains a system to *learn* from its experiences. Programming a computer to complete a task step by step has become impossible for complex tasks such as image recognition or market predictions. However, ML pushes the boundary by ultimately allowing a system to identify its mistake and retry one of many actions to achieve its goal. The system notes the mistake and can reference past errors in future decision-making processes.

This incredible technology has resulted in Terminatoresque theories of the future, even while technology giants are investing billions in research to drive AI development. Google's acquisition of Nest for \$3.2 billion brought a powerful suite of ML and products driven by the Internet of Things (IoT) to the Google portfolio.⁹ Amazon's \$775 million purchase of Kiva Systems integrated a suite of robust robots into the company's core warehouse operations.¹⁰

Many of our most ambitious societal objectives—from self-driving cars to cures for diseases to space exploration—rely and will continue to rely on the progress of machine learning. AI is here to stay, and while the revolution is just beginning, blockchain technology may help accelerate its growth.

AI and centralization

"Data is the new oil" is a common expression that reveals the value of data.¹¹ However, we need our AI algorithms to be processing data—extracting, refining, distributing, and monetizing it—in ways that respect human rights and create more opportunities for people, not just for corporations.¹² ML requires a massive amount of data for these training models to accurately develop and learn over time. Developers feed these data sets into systems to interpret and adjust their decision-making skills.

One key issue arises: how can we trust the data fed into these models?

One key issue arises: how can we trust the data fed into these models? How can we better understand how a particular system or algorithm is actually making its decision? For example, if an automotive manufacturer discovers a major fault in a testing simulation in its self-driving program, the company needs to turn to an auditable, trustless source of the data used for the simulation. This provides an intuitive way to verify the data used.

Additionally, a handful of large technology companies control the vast majority of data used for AI development. Google and Facebook collect gigabytes' worth of data per user, millions of data points that include every message exchanged through Messenger, location history, music preferences, and more.¹³



Acquiring these data is impossible for smaller companies, allowing a few large players to serve as gatekeepers to AI development. Public data sets, for the most part, are very limited and may have covenants for commercial use. There are multiple public data sets available (Table 1), but these data can be unstructured and not as insightful as the more proprietary data sets.

How can we democratize these data so that developers all around the world can easily gain access and use it to build AI data sets?

Some start-ups, including Spil.ly, needed hundreds of thousands of images for their image recognition development. Unable to pay for the input data required, Spil.ly was forced to create synthetic data from computer-generated graphics to feed its AI models.¹⁴

How can we democratize these data so that developers all around the world can easily gain access and use it to build AI data sets? Could a decentralized marketplace help democratize the data required for AI training in a trustless manner?

Table 1: Public data sets

Source of data	Type of data
CERN (Conseil Européen pour la Recherche Nucléaire) Open Data Portal	Particle physics experiments
eBay Market Data Insights	Online sales and auctions
US Federal Bureau of Investigation	Uniform crime reporting at national, state, and county levels
<i>Financial Times</i>	Global market data including stock price indexes, commodities, and foreign exchange
FiveThirtyEight	Polling on public opinion of politics and sports
Glassdoor	Job openings, salaries, company reviews
Google Trends	International internet search activity and trending news stories
International Monetary Fund	International finances, debt rates, foreign exchange reserves, commodity prices, and investments
Microsoft Machine Reading Comprehension (MARCO)	Reading comprehension and question answering
United Nations Comtrade	International trade
US Bureau of Justice	Adults on state and federal probation and parole, jail inmate populations, jail capacity, arrest-related deaths, capital punishment
World Bank	Population demographics and economic and development indicators


Source of data: Bernard Marr, "Big Data and AI: 30 Amazing (and Free) Public Data Sources for 2018," Forbes Magazine, Forbes Media, 26 Feb. 2018. www.forbes.com/sites/bernardmarr/2018/02/26/big-data-and-ai-30-amazing-and-free-public-data-sources-for-2018/#6e0d8b7e5f8a, accessed 6 March 2019.



Blockchains and artificial intelligence

The space is evolving from low-throughput decentralized systems to industrial-strength high-velocity differentiators.

Blockchain, at its core, is a distributed database stored by parties in a decentralized network. Over the past two years, we have seen an explosion in attention and resources in the blockchain space with hundreds of billions' worth of cryptocurrency purchased and traded. We saw an *initial coin offering* (ICO) boom in 2017 in which ICOs raised over \$4 billion, promising disruptive changes in practically every single existing industry.¹⁵ Bancor, a network for token exchange, raised \$153 million in three hours, even though it had no significant consumer adoption or use. There has been little connection between the prices of tokens as digital assets associated with different blockchains, and the real value of the underlying technology as a decentralized structure for storing data. The space is evolving from low-throughput decentralized systems to industrial-strength high-velocity differentiators.



Board Computer Circuit Background Technology by axonite, 2008, used under Pixabay license, as of 10 March 2019.

There has been no better time than now to begin understanding the powerful abilities of the blockchain. Satoshi Nakamoto's Bitcoin white paper emphasizes the importance of placing value in "cryptographic proof instead of trust."¹⁶ This seemingly simple idea unveiled a completely new realm of interactions. Rather than placing our trust in individuals and enterprises prone to human error, let us trust the objectivity of cryptography and mathematics. The Bitcoin white paper introduced a new way of thinking that AI may serve to benefit from one day.

Two of blockchain's most powerful characteristics are its immutability and decentralization. Both of these present opportunities to improve AI's centralized state.



Immutability

AI training models not only require vast amounts of data, but the data used needs to have a clear audit trail to ensure quality. Blockchains have the ability to store data in an immutable manner, to avoid any deletion or alteration.

AI training models not only require vast amounts of data, but the data used needs to have a clear audit trail to ensure quality.

For example, on the Bitcoin platform, a valid transaction is broadcast and confirmed by a set of “miners” that know neither the other miners nor the sender. Once a transaction is confirmed and added to the ledger, it cannot be altered or reversed without an immense amount of money and computing power. This helps provide a level of confidence in the transaction and, looking higher level, can translate into secure data points available for analysis.

KenSci, a prediction platform powered by artificial intelligence and machine learning, is focused on healthcare data analytics. Leveraging large sets of medical records and training sets, the company relies on the accuracy of the data input into the models, because patients’ lives are on the line. As co-founder Ankur Teredesai underscored, “We’re talking about real patients, real lives” when running these AI models to provide healthcare recommendations to doctors and insurance companies.¹⁷

However, AI has the potential to reinforce existing cultural and social biases embedded in the data sets that it relies on. For example, Google faced backlash for offensive tags related to image recognition in its photos app, and so did Microsoft when its Twitter chatbot used unverified and inappropriate input data to determine responses.


As companies undertake tests on a particular system, they will be able to point to a specific data point found at a particular timestamp on the blockchain. For example, a multinational clothing retailer may want to understand how its products are selling across various channels such as brick-and-mortar and e-commerce. It could input point-of-sale data from each channel into an inventory-management training model that could automatically adjust forecasting and inventory based on consumer trends. Over time, this model could help predict shifts in consumer tastes.

With a shared blockchain, data from each channel could be hashed and uploaded to a blockchain so that the retailer could be confident that the data used in its model had not been altered.

Retailers such as Walmart and Nike have been investing in the data analysis tools required for increased sales and customer engagement. Companies are looking at customer sentiment and inquiries on Twitter during and after their shopping experience.¹⁸

The data from these various channels must be as accurate as possible. Multiple parties and companies are selling the retailer’s products; the data are not stored in one place. With a shared blockchain, data from each channel could be hashed and uploaded to a blockchain so that the retailer could be confident that the data used in its model had not been altered. All parties to transactions would need to reach consensus on any changes to the data, and so no single player could rig the data in its favor.





Board Circuits Control Center Trace CPU PC by Gerd Altmann (geralt), 2013, used under Pixabay license, as of 10 March 2019.

Requiring consensus would serve as a governance tool: it would not be completely decentralized, yet would prevent a single entity from having control over sensitive data.

Current solutions involve external parties aggregating and selling data to interested companies; however, these data cannot truly be audited for accuracy, and the aggregators bear none of the risk associated with its use. So there is moral risk in this market. Companies relying on this information to make key decisions, do so at their own risk.

"Deep learning loves data. The larger and more credible the data set, the more intelligent AI's interpretation and predictions based on that data will be."

 BOB TAPSCOTT
Consultant
The Tapscott Group

Blockchain technology may help provide a high-quality, ever-growing stream of transparent data to which AI models could refer regularly over time. Once hashed, the data would be available in a secure manner, only to parties involved. There would be no central entity or aggregator of the retail data. Bob Tapscott, blockchain and AI author and strategy consultant, explained how computers' ability to forget nothing is incredibly important for AI:

Deep learning loves data. The larger and more credible the data set, the more intelligent AI's interpretation and predictions based on that data will be. With AI, 20 years of actual data and the subsequent human decision-making can be reviewed, learned, refined, and critiqued overnight.¹⁹

For example, given a large enough data set of MRI data and the actual patient outcomes, an AI routine was able to detect colorectal cancer with 86 percent accuracy; and AI routines have the potential to predict precancers earlier than human beings can.²⁰ Where there are massive amounts of data, AI routines can readily see subtle correlations in millions of images we cannot.



It is all about the data. If we run a remote text-based diagnostic shop (be it for sick people or broken machines) and have the text for 100,000 queries, responses, and subsequent satisfaction scores, with current AI technology, we can train a computer to perform those tasks to get consistently higher satisfaction scores than human beings. Then the computer can simultaneously translate the results into 20 different languages, depending on customer preferences. Of course, processing natural language can still be quite challenging, depending on the complexities of the written requests.

Note that these models are improved with the accuracy of data. With vacuum-cleaner repair, the right answers and the most pleasing answers are likely the same thing. With medical patients, these answers may not be the same. The more data we use to train a model, the more accurate the model will be.

“Based on 100,000 previous observations, when something goes wrong the computer can filter out anomalies to the data,” Tapscott said.²¹ In other words, the computer could learn from mistakes and apply its learning to future decision-making. Tapscott touched not only on the importance of having data but also on the ability to *verify* the data inputted.

When data are put on a blockchain, we can identify the exact point of entry or faulty input and then correct it, for a decision at hand and for future decisions. The decision-making process improves through each iteration. AI, supplemented by blockchain, will allow us to understand a problem or pattern as no human could.²²

AI, supplemented by blockchain, will allow us to understand a problem or pattern as no human could.

Blockchain and democratization

Our identities are increasingly becoming linked to data that may not be accurate or kept securely (Table 2, next page). With current centralized data-capturing systems such as Amazon, Facebook, and YouTube, we have less and less control over our data. As a result, our digital identities have too often been mismanaged and compromised, sometimes without our knowledge. Corporations should not be able to sell our data without providing us with some financial incentive.

A key advantage of blockchain technologies is their potential for ownership of personal data, ensuring increased privacy and security. Blockchain also has the potential to democratize the availability of data. A handful of corporations control the vast amount of data that AI models require. The real insight lies in data such as customer buying patterns or app usage, yet only a few corporations can leverage this information. Imagine if a variety of firms and developers of all sizes could gain permission to access large sets of data to push innovation forward.

Blockchain can facilitate this democratization of data, ultimately producing a new type of data marketplace. Networks of companies could purchase data directly from users, with the data available on the blockchain. The availability of data peer to peer would allow



users—including companies traditionally excluded from the data ecosystem—to garner completely new insights. Here’s how it would work:

Buyers who offer money in exchange for training an AI model open up the possibility for AI to reach the hands of all types of developers and companies looking to gain insights from their existing data sets.

1. A buyer creates a smart contract containing a data set and commissioning an AI/ML model that must meet certain criteria. For example, an online retailer could write a contract that would release a data set of 20,000 orders and request an AI model that could identify fraudulent orders with greater than 90 percent accuracy, in exchange for some amount of ether.
2. The buyer publishes the smart contract to the Ethereum blockchain for anyone to access.
3. Machine learning engineers access the contract and download the data set. The race is on!
4. Each engineer uses the data to train an AI/ML model.
5. When an engineer has trained a model sufficiently, the engineer runs this model on the Ethereum blockchain.
6. If the model indeed meets contract requirements—that is, it identifies fake orders with greater than 90 percent accuracy—then the smart contract sends the model to buyer and issues payment to the engineer.²³

Buyers who offer money in exchange for training an AI model open up the possibility for AI to reach the hands of all types of developers and companies looking to gain insights from their existing data sets.

Table 2: Largest data breaches reported in 2018

Company	Parent	Sector	Number of accounts	Reported problem
Facebook		Web (social)	2,200,000,000	Hacked
Aadhaar	Ministry of Electronics and Information Technology, India	Government	1,100,000,000	Poor security
Twitter		App (social)	330,000,000	Poor security
MyFitnessPal	Under Armour	App (fitness)	150,000,000	Hacked
Neteests	Social Sweethearts	App	120,000,000	Poor security
Firebase	Google	App	100,000,000	Poor security

Source of data: David McCandless, "World's Biggest Data Breaches and Hacks," *Information Is Beautiful*, updated 15 Oct. 2018. www.informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks, accessed 16 Nov. 2018.



What if we could share the data among participants with a governance protocol in place?

What if we could share the data among participants with a governance protocol in place? For example, what if automotive manufacturers sold autonomous testing data to one another, or gave data sets away in exchange for the insights derived from them? Enabling this kind of sharing could help accelerate the development of new vehicles.

Autonomous vehicles are generating terabytes of data, with one car capable of producing hundreds of terabytes' worth of radar, lidar, and camera data for analysis (Table 3). Rather than settling for current data silos, we could attempt to create a data marketplace for automotive manufacturers.

The data would be on the blockchain, so that all players would have an understanding of data sources and timestamps. There would be no central party controlling, aggregating, or removing the data, decreasing the chances of manipulation. Participants uploading suboptimal data could be voted out of the blockchain by other participants.

Smart contracts

One of the biggest issues behind using data for AI is the privacy and sensitivity of the content. Smart contracts have the potential to help govern the actual use of these data in an AI environment. The ability for a smart contract to face "computational scrutiny" without relying on a third party could result in new governance platforms.²⁴

Dr. Henry Kim, associate professor and co-director of blockchain lab at York University, explained that we "may be able to use a governance smart contract." While the data are in use, entities must still "observe privacy" and confirm that they have access to the right data. Dr. Kim touched on the importance of a smart contract in dictating the governance policies of the data. The smart

Table 3: Data generated by car automation sensors

Sensor type	Quantity	Data generated per sensor
Radar	4-6	0.1-15 Mbit/s
Lidar	1-5	20-100 Mbit/s
Camera	6-12	500-3500 Mbit/s
Ultrasonic	8-16	<0.01 Mbit/s
Vehicle Motion, GNSS, IMU		<0.1 Mbit/s

Source of data: Stephan Heinrich, "Flash Memory in the Emerging Age of Autonomy," Flash Memory Summit, 7 Aug. 2017. www.flashmemorysummit.com/English/Collaterals/Proceedings/2017/20170808_FT12_Heinrich.pdf, accessed 6 March 2019.



contract itself could contain the set of rules for using those data and rewarding the corresponding party for providing the data. The question comes down to whether we may use that person's or party's data in exchange for payment?²⁵

In addition, AI has the potential to analyze smart contracts for errors.

Blockchain governance

For example, if a company chooses to use a particular data set, the governance details behind that data set could be stored on chain for all interested parties to verify. This is incredibly important as data privacy laws come under even more scrutiny. The Ponemon Institute found that 71 percent of companies would face global detriment to their businesses for failure to comply with EU *General Data Privacy Regulation* (GDPR).²⁶

Rather than having a grand mesh of data aggregated centrally by a single party, a company could extract the data sets relevant to its needs, verify the data, and then use the sets according to governance rules in the smart contract. An on-chain smart contract would be available for all participants to verify, compared to a simple two-way contract. Network participants could be voted out for malicious behavior.

In addition, AI has the potential to analyze smart contracts for errors. As millions of smart contracts are developed, engineers could input the contracts themselves into a training model to determine where errors might lie and whether they could construct a more efficient contract.

In a distributed AI system, a central system still manages the overall setup and has greater permissions over activities such as data aggregation.

Decentralized artificial intelligence

Now that we have discussed the applications of blockchain in existing centralized solutions, we can begin looking at the functional potential of blockchain in a decentralized AI system. *Decentralized artificial intelligence* (DAI) is a subfield of artificial intelligence that analyzes the interaction and activity of autonomous agents. DAI attempts to understand how agents distribute themselves, interact with one another, and adjust to changes in their environment, without the oversight of a central body.²⁷ An entirely new and autonomous data ecosystem could emerge.

Distributed AI versus decentralized AI

Distributed artificial intelligence establishes a different setup from a completely decentralized entity. In a distributed AI system, a central system still manages the overall setup and has greater permissions over activities such as data aggregation.

For example, imagine a system that tracks the number of photos taken by a group of smartphones in a given period.



A central entity oversees distributed artificial intelligence, which depends on the entity to use and process the data carefully. However, what if agents—sensors embedded in cars or factory equipment, or distributed throughout corporations—could begin learning and transacting with one another? That is where the combination of AI and blockchain gets really exciting.

The ability for one car to teach another car where kids tend to play on the street has the potential to improve safety dramatically for all.

DAI and blockchain's role

In DAI, a multiagent system could one day emerge. For example, in the future, an autonomous vehicle may be able to pick up riders, earn fares, and pay for insurance, maintenance, and gas all without human intervention. The car could then begin to participate in the sharing economy by offering rides to passengers. Through an AI model, and millions of rides later, the car could learn various trends that a human driver would never consider. The ability for one car to teach another car where kids tend to play on the street, for example—to transfer what it has learned about a neighborhood to a car that has never traveled in that neighborhood has the potential to improve safety dramatically for all. The veteran car could even sell its learnings to the new car on the block.

The car may begin positioning itself in areas with high densities of passengers. It may understand when it needs a maintenance check and could scan local mechanics for the best pricing. The vehicle may understand how to receive the most affordable insurance policy by adjusting its driving habits or “negotiating” with autonomous agents at other entities such as insurance companies and vehicle manufacturers.

DAI will need an operating platform to run on. Were this platform run by a centralized authority, the agent itself could be compromised through a technical back door, which allowed the centralized entity to take over. Blockchains are run not by some central organization but by the nodes themselves. These nodes decide how to govern the future of the network. As stakeholders, they must first approve any protocol upgrades or developments prior to implementation.

In a more decentralized world, autonomous agents need not rely on a handful of data sources but could move from one data source to another, leaving an auditable trail for the data used.

Autonomous agents need a platform on which they can vote and help direct the future of the protocol itself. Blockchains would not care nor discriminate whether a node were a human or an autonomous agent, but they could assign more weight to the votes of nodes that have learned more.²⁸

George Polzer, director of education at Blockchain 48 and task force chairman of the Enterprise Ethereum Alliance, sees blockchain “as an enabler ... to gather data from multiple nodes” helping to achieve “collective intelligence.”²⁹ In a more decentralized world, autonomous agents need not rely on a handful of data sources but could move from one data source to another, leaving an auditable trail for the data used. Imagine a massive number of nodes pooling their data, contributing to a data commons of sorts, rather than dealing with the very siloed and centralized pools of data that exist now.




In addition, blockchains would allow autonomous agents to interact without having to trust each other. Because data is immutable on the blockchain, agents can quickly verify that an agent is not falsifying its data for its own interest. The trust lies not in the agent itself but rather the cryptographically secure data produced by the agent's actions.

In addition, blockchains would allow autonomous agents to interact without having to trust each other.

Reviewing our autonomous vehicle example, we can see that a multiagent system needs to achieve consensus or agreement on terms. For example, how could the insurance company know that the vehicle had not altered its data before requesting a new policy?

Autonomous vehicles could leverage a blockchain by continuously uploading their data, through hashes, for other agents at other entities—insurance companies, vehicle manufacturers—to verify. This verification would provide an important level of trust and assurance that the data had not been tampered with for the agent's benefit.

Also, blockchains introduce an incentive layer on top of DAI. Governance protocols could require agents to stake or set aside their tokens in reserve while they interact with other agents. If a particular agent decided to act maliciously, other agents could vote to remove that agent from participating in the ecosystem and retain some portion of that agent's staked tokens as a penalty. The governance protocols could reward honest nodes for performing well (e.g., picking up passengers on time, driving safely) and could penalize malicious nodes (e.g., temporary denial of access, higher transaction fee, more tokens required in reserve).



Board Computer Chip Data Processing Solder Joint by Michael Schwarzenberger (blickpixel), 2014, used under Pixabay license, as of 10 March 2019.



Autonomous agents as AI representatives?

Imagine a system where one autonomous vehicle is elected to represent other autonomous vehicles for better bargaining power over insurance policies.

One possible economic structure could be the use of *delegated proof of stake* (DPoS). DPoS is a consensus mechanism used by other blockchain projects such as EOS and VeChain. Similar to the workings of the US legislative system, each token holder in these blockchains can vote on a representative to validate transactions on its behalf. A group of nodes, between 10 and 100, are the sole individuals who can validate transactions in the network.

Autonomous agents could begin representing other autonomous agents as delegates. As delegates, these agents would have to “stake” or put away some token in reserve, which would be slashed if the agent acted maliciously.

Imagine a system where one autonomous vehicle is elected to represent other autonomous vehicles for better bargaining power over insurance policies. Because each transaction would occur on the blockchain, other agents could verify that the delegate is performing in the interest of the network.

Toufi Saliba, CEO of Toda.Network and chair of the Association for Computing Machinery, described how we might see an environment where “machines can ensure each other’s governance accordingly; the machines, they can compete between each other.”³⁰ Saliba pointed to programs that follow a set of protocol rules and exchange information, yet still act individually. Each program plays a role in challenging malicious actors and removing them from a shared protocol. If a miner is malicious on the Bitcoin network, then it will be punished financially, as it is forced to spend vast amounts of money to achieve a 51 percent attack. In an autonomous world, autonomous miners can begin to punish malicious miners automatically, voting them off the network.

The crux to this governance is that no central authority would dictate which agents could join, leave, or take a percentage of the agents’ earnings. A group of autonomous agents would run the system: each would provide its resources for the greater good while still receiving rewards for its individual efforts.

The crux to this governance is that no central authority would dictate which agents could join, leave, or take a percentage of the agents’ earnings.

Saliba reminded us that people do not “want to acknowledge that there’s fear of attack from within.”³¹ Central authorities control the keys to vast amounts of permissioned information and data. This centralization puts these data at greater risk of hacking than data distributed across a cryptographically secure network, as we have seen happen at dozens of corporations (Table 2). Institutions that many people “trust today,” Saliba said, may be the same institutions that individuals “don’t know if they can trust tomorrow.” The simple yet groundbreaking topic of trust brings us to a new concept of organization and autonomous agents.



Looking to the future: DAO

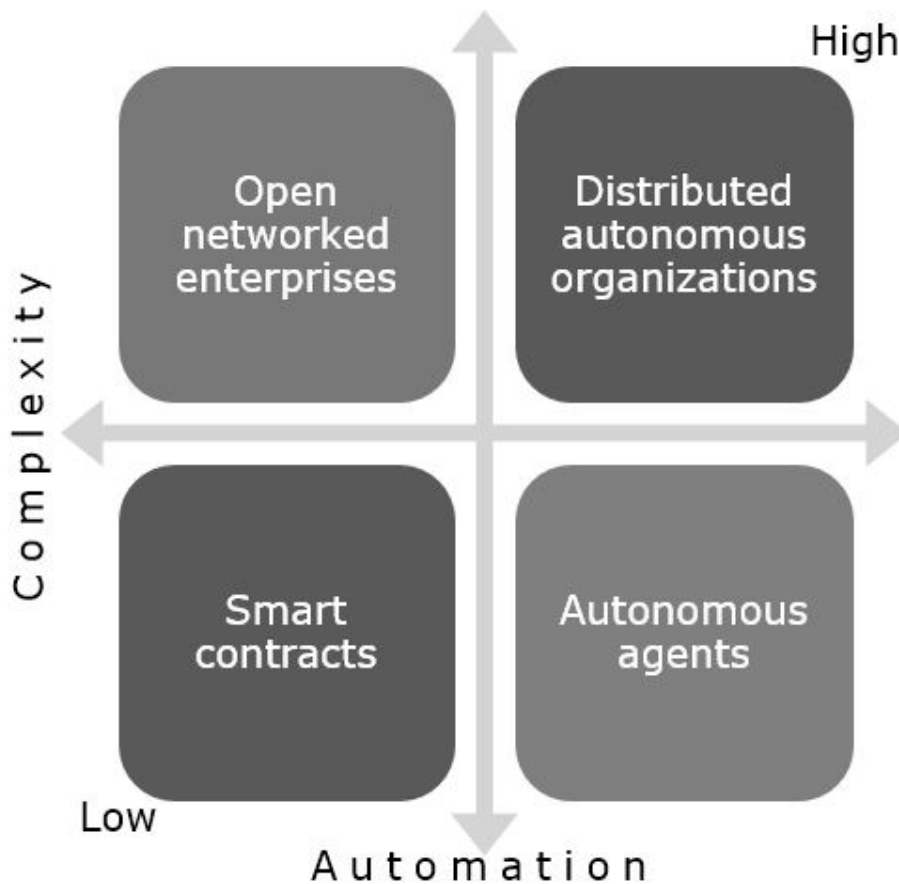
Autonomous agents can make their own decisions and learn over time through trial and error without human intervention or risk of human error. But can we create something more complex than an autonomous vehicle?

Autonomous agents can make their own decisions and learn over time through trial and error without human intervention or risk of human error.

Looking at the distributed economic entities visual (Figure 1), we can gain an understanding of where autonomous agents fit within the ecosystem. An autonomous vehicle or a vending machine that can understand customer demand and inventory management falls under relatively high automation and low complexity.

We can automate a few tasks such as replacing products in a vending machine or paying for insurance over time. However, what happens when we venture toward a complex corporation run on the

Figure 1: Distributed economic entities



Source of concept: Don Tapscott and Alex Tapscott, Blockchain Revolution: How the Technology Behind Bitcoin and Other Cryptocurrencies Is Changing the World (New York: Portfolio Penguin, 2018): 120.



An ambitious future would have organizations running entirely on their own, making business decisions, hiring employees, and responding to market changes.

blockchain? That is where the idea of a *decentralized autonomous organization* (DAO) comes into play. A DAO is an entity run entirely by its stakeholders rather than an executive team. Because of its decentralized nature, any stakeholder can propose a change to the organization's actions or structure with other stakeholders voting.

Over time, a DAO can begin to learn who its competitors are, what proposals stakeholders are supporting, and what market trends to follow. An ambitious future would have organizations running entirely on their own, making business decisions, hiring employees, and responding to market changes.

Imagine a US toy factory that knew exactly when to order raw materials, based on plastic production in China. The factory would use web scrapers to analyze toy searches and trends from millions of websites and videos to know which colors, styles, quantities, and trends to produce. The factory would then begin production, optimizing each piece of equipment for constructing a particular part of a toy. During production, the factory would contact the largest buyers in Europe, South America, and Asia and would negotiate with external agents to achieve the target profit margin. It would then ship the toys and receive payment without any human intervention. The factory would keep track of its earnings and analyze its competitors' earnings in its objective to become a market leader.

George Polzer argued that "anyone playing an intermediary role [such as] financial institutions or credit bureaus" would be disrupted by DAOs along with "any industry that is not currently transparent in transactions." These intermediaries could be governed by smart contracts that dictate the movement of information and value that these intermediaries have traditionally controlled. Smart contracts could "provide an independent, decentralized inner economy." The governance of DAOs is a critical topic that is currently being researched; however, "governance by a council (DPoS) ... [is] a novel idea worth experimenting and building."³²

Although the DAO model faced scrutiny by the Ethereum community in 2016, DAOs could be one of the most important advancements of artificial intelligence, with unprecedented use cases.

Blockchains provide a variety of use cases for AI development, but there are challenges with using a decentralized system.

Challenges with blockchains

Blockchains provide a variety of use cases for AI development, but there are challenges with using a decentralized system. For example, one of the most pressing issues with a blockchain is the issue of scaling.



Scalability


As people and things generate more data, storing those data in a decentralized manner can be challenging. Public blockchains are very expensive for storing vast amounts of data; and beyond the cost, quickly retrieving the data can be computationally intensive. Computing systems will need to retrieve, analyze, and compute algorithms in a very short period.

With terabytes of data passing between products and companies, existing solutions will not suffice. Fortunately, innovators are developing solutions for scaling.

Not only will data need to be collected and distributed to all relevant stakeholders, but that same data needs to be structured and trained within a short period of time. Current blockchain solutions do not offer the scalability required for such types of aggregation and computation (Figure 2). With terabytes of data passing between products and companies, existing solutions will not suffice.

Fortunately, innovators are developing solutions for scaling. One is the *InterPlanetary File System (IPFS)*. Its users download the IPFS protocols for storing and indexing files on their own and each other's devices, in a distributed network rather than on a central set of servers. Its storage capacity grows with the number of users and the available space on their collective machines. Users can host data and create a hash of the data set, so that others can locate and access the nearest copy of it, if not the original host's copy.

Figure 2: Comparison of transactions per second



Source of data: Raul, "Transactions Speeds: How Do Cryptocurrencies Stack Up to Visa or PayPal?" *howmuch.net*, 10 Jan. 2018. howmuch.net/articles/crypto-transaction-speeds-compared, accessed 6 March 2019.



For example, when an application needs to reference a large data set on a blockchain, it could refer to just the hash of the data rather than the whole data set. The hash should always match the data set, unless the set has been altered.

Governance

AI will require myriad data from various sources, but data owners will want to keep those data private and secure on the blockchain, so that only permissioned parties have some level of access.

Establishing the governance of these blockchains and the data they hold is incredibly important. AI will require myriad data from various sources, but data owners will want to keep those data private and secure on the blockchain, so that only permissioned parties have some level of access. We could see some integration of homomorphic encryption, which allows engineers to use data sets for training without compromising the confidentiality of those data.

In the midst of AI model development, user data stored on an immutable ledger may also become problematic. What if we wanted to transform and map the data from one raw format into another format so that we could work more easily with the data set downstream, say, for business analytics? Or what if governments placed limitations on the types of individual or household data that organizations may collect and use for commercial purposes?

How will we remove data sets that are supposed to be immutable and safe from deletion? Or is it a matter of changing the smart contracts governing access and usage of such data? Moreover, in these early days of development, as autonomous agents learn to interact, blockchains will likely face latency issues; the rate at which participants broadcast and validate transactions may be slow.

Conclusion and recommendations

The idea of an immutable, decentralized ledger providing an auditable data source for hungry AI models could vastly improve if not revolutionize existing models.

The convergence of blockchain and AI is an exciting frontier to explore. The idea of an immutable, decentralized ledger providing an auditable data source for hungry AI models could vastly improve if not revolutionize existing models. There is still plenty of development work, testing, and redevelopment to be done before this enterprise capability is ready for mass deployment. However, the future could not look more exciting.

AI is becoming more prevalent on the enterprise side. In numerous industries, this technology has shown immense utility, and its growth will continue with enterprise adoption. However, AI is extremely centralized in our current infrastructure. The question for enterprise leaders is, "Do we want to entrust autonomous agents to a handful of corporations, or do we want to advocate for a more decentralized architecture and open governance?"



A decentralized blockchain-based solution may allow a more democratized yet secure means of transmitting and using the data required for AI training models.

As scaling solutions mature, enterprise leaders should discuss where blockchain fits in the AI stack of their enterprise.

Blockchain can provide a decentralized infrastructure for the AI ecosystem. Currently, AI data sources and development are very centralized, preventing smaller players from working on novel AI solutions. A decentralized blockchain-based solution may allow a more democratized yet secure means of transmitting and using the data required for AI training models.

Blockchain adds an *incentive* layer on top of an Internet-based sharing economy. Such an infrastructure could open completely new channels for companies to share data sources and tap new resources and provide some incentive for doing so. Organizations could develop reputations for the quality and utility of the data they share.

Imagine the impact of millions of AI agents running on a few centralized servers. While the focus today may be on centralized AI, *decentralized* AI will lead to cars, machines, and even companies that are more autonomous. Pure autonomous systems will need decentralized platforms to run on. Blockchain may provide better insights over centralized, and possibly biased, data sets.


Current blockchain solutions are not ready for wide-scale deployment. Work is underway on scaling solutions ranging from layer-one improvements to layer-two scaling solutions such as the Lightning Network. Each solution comes with trade-offs in throughput as well as governance, and so no player dominates. As scaling solutions mature, enterprise leaders should discuss where blockchain fits in the AI stack of their enterprise.

Consider hybrid models of decentralized and centralized infrastructure. Which data sets could have greater value in such a hybrid model? If those data sets were ready, we could begin exploring blockchain platforms and determining which worked best. Private blockchain initiatives such as R3 and public blockchains such as Ethereum are competing for enterprise adoption. While private blockchains are often not as decentralized as public ones, initial tests on permissioned networks could inform more decentralized solutions down the road.

Distributed ledger technology cannot solve all of AI's problems. Blockchain may help address some, but not all, of the issues that AI developers face. For example, to learn and develop over time, these models need access to high-quality streams of data. The vetting of these data is crucial. While blockchain helps to track provenance, it does not yet verify accuracy or quality of inputs.

The great creation of value is yet to come. Remember, with the Internet, we did not immediately understand the potential for such use cases as e-commerce, social media, and cloud computing. These took years to reach mass adoption. Yet, as the infrastructure matured, start-ups and enterprise innovators alike created tremendous value by developing and running applications on top of the Internet. Distributed applications programmed to run on top of blockchains may generate as much or more value over time.





About the authors

Don Tapscott

Don Tapscott, executive chairman of the Blockchain Research Institute, is one of the world's leading authorities on the impact of technology on business and society. He has authored 16 books, including *Wikinomics: How Mass Collaboration Changes Everything*, which has been translated into over 25 languages. Don's most recent and ambitious book—*Blockchain Revolution: How the Technology Behind Bitcoin and Other Cryptocurrencies Is Changing the World*—was co-authored with his son, Alex Tapscott, a globally recognized investor, advisor, and speaker on blockchain technology and cryptocurrencies. According to Harvard Business School's Clayton Christensen, it is "the book, literally, on how to survive and thrive in this next wave of technology-driven disruption." The paperback version of the book, updated with new material covering recent developments in the blockchain industry, was published in June 2018.

In 2017, Don and Alex co-founded the Blockchain Research Institute. Its 80+ projects are the definitive investigation into blockchain strategy, use cases, implementation challenges, and organizational transformations. Don is a member of the Order of Canada and is ranked the second most influential management thinker and the top digital thinker in the world by Thinkers50. He is an adjunct professor at INSEAD and Chancellor of Trent University in Ontario. It is hard to imagine anyone who has been more prolific, profound, and influential in explaining today's technological revolutions and their impact on the world.

Anjan Vinod

Anjan Vinod is a member of Blockchain at Berkeley and a student at the Haas School of Business. He also works at OPEN Platform, a blockchain infrastructure for applications, where he leads product development and operations. OPEN raised a private round of financing from some of the top venture capital firms in the space including Draper Dragon and NEO Global Capital. Anjan's previous experiences include hedge funds, investment banking, and AT&T where he was involved with its IoT division.

Acknowledgments

The authors acknowledge the generous feedback on previous drafts from Lofred Madzou, project lead on artificial intelligence and machine learning for the World Economic Forum, and Christos Oikonomou, PhD student in decision sciences at INSEAD.






About the Brightline Initiative

Brightline is a Project Management Institute (PMI) initiative together with leading global organizations dedicated to helping executives bridge the expensive and unproductive gap between strategy design and delivery. Brightline delivers insights and solutions that empower leaders to successfully transform their organization's vision into reality through strategic initiative management. Learn more at www.brightline.org.

Brightline is also a founding member of the Blockchain Research Institute (BRI) and is republishing this light house white paper with the BRI's permission.

Please note that *Brightline* is a trademark of the PMI, serial number 87318956 in the US Patent and Trademark Office. All rights reserved.





About the Blockchain Research Institute

Co-founded in 2017 by Don and Alex Tapscott, the Blockchain Research Institute is an independent, global think tank established to help realize the new promise of the digital economy. For several years now, we have been investigating the transformative and disruptive potential of blockchain technology on business, government, and society.

Our syndicated research program, which is funded by major corporations and government agencies, aims to fill a large gap in the global understanding of blockchain protocols, applications, and ecosystems and their strategic implications for enterprise leaders, supply chains, and industries.

Our global team of blockchain experts is dedicated to exploring, understanding, documenting, and informing leaders of the market opportunities and implementation challenges of this nascent technology. Research areas include financial services, manufacturing, retail, energy and resources, technology, media, telecommunications, healthcare, and government as well as the management of organizations, the transformation of the corporation, and the regulation of innovation. We also explore blockchain's potential role in the Internet of Things, robotics and autonomous machines, artificial intelligence, and other emerging technologies.

Our findings are initially proprietary to our members and are ultimately released under a Creative Commons license to help achieve our mission. To find out more, please visit www.blockchainresearchinstitute.org.

Research management

Don Tapscott – Co-Founder and Executive Chairman
Kirsten Sandberg – Editor-in-Chief
Hilary Carter – Managing Director

Others in the BRI leadership team

Alisa Acosta – Director of Education
Luke Bradley – Director of Communications
Wayne Chen – Director of Business Development
Maryantonett Flumian – Director of Client Experience
Roya Hussaini – Director of Administration
Jody Stevens – Director of Finance
Alex Tapscott – Co-Founder



Notes

1. Deborah Bothun, Matthew Lieberman, and Anand S. Rao, "Bot.Me: A Revolutionary Partnership: How AI Is Pushing Man and Machine Closer Together," *Consumer Intelligence Series*, *PwCArtificialIntelligence.com*, PricewaterhouseCoopers International, 2017. pwcartificialintelligence.com, accessed 21 July 2018.
2. Deborah Bothun, Matthew Lieberman, and Anand S. Rao, "Bot.Me: A Revolutionary Partnership."
3. Pega, "What Consumers Really Think About AI," n.d. www1.pega.com/system/files/resources/2017-11/what-consumers-really-think-of-ai-infographic.pdf, accessed 8 Aug. 2018.
4. Louis Columbus, "10 Charts That Will Change Your Perspective on Artificial Intelligence's Growth," *Forbes.com*, Forbes Media, 12 Jan. 2018. www.forbes.com/sites/louiscolombus/2018/01/12/10-charts-that-will-change-your-perspective-on-artificial-intelligences-growth/#371a41124758, accessed 24 Aug. 2018.
5. Narrative Science, "Outlook on Artificial Intelligence in the Enterprise," Feb. 2019. narrativescience.com/wp-content/uploads/2019/02/Research-Report_Outlook-on-AI-for-the-Enterprise.pdf, accessed 4 March 2019.
6. John Markoff, "Computer Wins on 'Jeopardy!': Trivial, It's Not," *The New York Times*, The New York Times Company, 16 Feb. 2011. www.nytimes.com/2011/02/17/science/17jeopardy-watson.html, accessed 7 March 2019.
7. IBM, "Korean Air Is Using Watson to Search Vast Amounts of Data to Improve Operational Efficiency and On-Time Performance," IBM Corporation, n.d. www.ibm.com/watson/stories/airlines-with-watson, accessed 28 Aug. 2018.
8. "Machine Learning on AWS," Amazon Web Services, Amazon, n.d. aws.amazon.com/machine-learning, accessed 28 Aug. 2018.
9. Alphabet, "Google to Acquire Nest," Announcement, 13 Jan. 2014. abc.xyz/investor/news/releases/2014/0113.html, accessed 28 Aug. 2018.
10. Amazon, "Amazon.com to Acquire Kiva Systems, Inc.," Announcement, 19 March 2012. www.sec.gov/Archives/edgar/data/1018724/000119312512122135/d318297dex991.htm, accessed 28 Aug. 2018.
11. "The World's Most Valuable Resource Is No Longer Oil, but Data," *The Economist*, The Economist Group, 6 May 2017. www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data, accessed 5 Sept. 2018.
12. Sumeet Santani, "Why Data Is the New Oil," *Infospace*, Syracuse University iSchool, 26 Feb. 2018. ischool.syr.edu/infospace/2018/02/26/why-data-is-the-new-oil, accessed 9 Sept. 2018.
13. Dylan Curran, "Are You Ready? Here Is All the Data Facebook and Google Have on You," *The Guardian*, Guardian News & Media, 30 March 2018. www.theguardian.com/commentisfree/2018/mar/28/all-the-data-facebook-google-has-on-you-privacy, accessed 9 Sept. 2018.
14. Tom Simonite, "Some Startups Use Fake Data to Train AI," *Wired*, Condé Nast, 25 April 2018. www.wired.com/story/some-startups-use-fake-data-to-train-ai, accessed 15 Sept. 2018.
15. Steven Russolillo, "Initial Coin Offerings Surge Past \$4 Billion—and Regulators Are Worried," *The Wall Street Journal*, Dow Jones & Company, 14 Dec. 2017. www.wsj.com/articles/initial-coin-offerings-surge-past-4-billionand-regulators-are-worried-1513235196, accessed 25 Sept. 2018.
16. Satoshi Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," White Paper, www.bitcoin.org, 1 Nov. 2008. bitcoin.org/bitcoin.pdf, accessed 20 Sept. 2018.
17. Maria Korolov, "AI's Biggest Risk Factor: Data Gone Wrong," *CIO*, IDG Communications, 13 Feb. 2018. www.cio.com/article/3254693/artificial-intelligence/ais-biggest-risk-factor-data-gone-wrong.html, accessed 21 Sept. 2018.
18. Emel Aktas and Yuwei Meng, "An Exploration of Big Data Practices in Retail Sector," *Logistics*, vol. 1, no. 2 (12 Dec. 2017): 1-28. MDPI, www.mdpi.com/2305-6290/1/2/12, accessed 24 Sept. 2018.
19. Bob Tapscott, interviewed in person by Anjan Vinod, 14 May 2018.
20. Brigham Hyde, "Five Predictions for AI and Real-World Data in Oncology," *Forbes Technology Council*, Forbes Media, 14 Aug. 2018. www.forbes.com/sites/forbestechcouncil/2018/08/14/five-predictions-for-ai-and-real-world-data-in-oncology/#67a3d77441ec, accessed 2 Dec. 2018.



21. Bob Tapscott, interviewed in person by Anjan Vinod, 14 May 2018.
22. Bob Tapscott, interviewed in person by Anjan Vinod, 14 May 2018.
23. Besir Kurtulmus and Kenny Daniel, "Trustless Machine Learning Contracts; Evaluating and Exchanging Machine Learning Models on the Ethereum Blockchain," *Algorithmia*, Algorithmia Research, 26 Feb. 2018. algorithmia.com/research/ml-models-on-blockchain, accessed 15 Sept. 2018.
24. Wessel Reijers, Fiachra O'Brolcháin, and Paul Haynes, "Governance in Blockchain Technologies & Social Contract Theories," *Ledger Journal*, Vol. 1 (2016): 134-151. www.ledgerjournal.org/ojs/index.php/ledger/article/download/62/51, accessed 24 Sept. 2018.
25. Henry Kim, interviewed via telephone by Anjan Vinod, 29 May 2018.
26. Ponemon Institute and McDermott Will & Emery, "The Race to GDPR: A Study of Companies in the United States and Europe," *International Association of Privacy Professionals*, April 2018. iapp.org/media/pdf/resource_center/Ponemon_race-to-gdpr.pdf, accessed 26 Sept. 2018.
27. Vitalik Buterin, "DAOs, DACs, DAs and More: An Incomplete Terminology Guide," *Ethereum Blog*, 6 May 2014. blog.ethereum.org/2014/05/06/daos-dacs-das-and-more-an-incomplete-terminology-guide, accessed 27 Sept. 2018.
28. Marc Peter Deisenroth and Jun Wei Ng, "Distributed Gaussian Processes," *Journal of Machine Learning Research*, Workshop and Conference Proceedings, Vol. 37, 2015: 1-10. Proceedings of Machine Learning Research Press, proceedings.mlr.press/v37/deisenroth15.pdf, accessed 7 March 2019. Presented at the 32nd International Conference on Machine Learning in Lille, France, this paper builds on Gaussian processes (GP), which "allow for flexible modelling without specifying low-level assumptions in advance." GPs have helped to advance such areas as active learning, data visualization, geostatistics, optimization, robotics and reinforcement learning, and spatiotemporal modelling. Distributed Gaussian process models are "conceptually simple: they split the data set into small pieces, train GP experts jointly, and subsequently combine individual predictions to an overall computation."
29. George Polzer, interviewed online by Anjan Vinod, 3 Dec. 2018.
30. Toufi Saliba, interviewed via telephone by Anjan Vinod, 16 Nov. 2018.
31. Toufi Saliba, interviewed via telephone by Anjan Vinod, 16 Nov. 2018.
32. George Polzer, interviewed online by Anjan Vinod, 3 Dec. 2018.







blockchainresearchinstitute.org



brightline.org