

Harvard Business Review

PUBLISHED ON HBR.ORG
NOVEMBER – DECEMBER 2016

INSIGHT CENTER COLLECTION

The Next Analytics Age: Artificial Intelligence

SPONSORED BY



The Next Analytics Age: Artificial Intelligence

What management and leadership challenges will the next wave of analytic technology bring? This Insight Center on HBR.org went beyond the buzz of what artificial intelligence can do, to talk about how it will change companies and the way we manage them.

- 1 The Competitive Landscape for Machine Intelligence**
Shivon Zilis and James Cham
- 10 What a Visit to an AI-Enabled Hospital Might Look Like**
R “Ray” Wang
- 15 How Predictive AI Will Change Shopping**
Amit Sharma
- 19 How One Clothing Company Blends AI and Human Expertise**
H. James Wilson, Paul Daugherty, and Prashant Shukla
- 24 Bots That Can Talk Will Help Us Get More Value from Analytics**
Stuart Frankel
- 28 Prepare Your Workforce for the Automation Age**
Christoph Knoess, Ron Harbour, and Steve Scemama
- 32 Why You Shouldn’t Swear at Siri**
Michael Schrage

CONTINUES ON NEXT PAGE

TECHNOLOGY

The Competitive Landscape for Machine Intelligence

by Shivon Zilis and James Cham

NOVEMBER 02, 2016



Three years ago, our venture capital firm began studying startups in artificial intelligence. AI felt misunderstood, burdened by expectations from science fiction, and so for the last [two years](#) we've tried to capture the most-important startups in the space in a one-page landscape. (We prefer the more neutral term "machine intelligence" over "AI.")

In past years, we heard mostly from startup founders and academics — people who pay attention to early, far-reaching trends in technology. But this year was different. This year we've heard more from Fortune 500 executives with questions about machine intelligence than from startup founders.

The State of Machine Intelligence, 2016

ENTERPRISE INTELLIGENCE

AUDIO Caplo Clover Intelligence Expect Labs Gridspace* Motivoi Nexidia Pop Up Archive* Quiricus TalkIQ Twilio	INTERNAL DATA Alation* Arimo* Cycorp Digital Reasoning IBM Watson Kyndi Outlier Palantir Primer Sapho*	MARKET Bottlenose CB Insights DataFox Enigma Mattermark Predata Premise Quid Tracxn	SENSOR Alluvium C3 IoT GE Predix Imubit KONUX Maana Planet OS Preferred Networks Sentenai ThingWorx Uptake	VISUAL Algoclan Captricity Clarifai Cortica Deepomatic DeepVision Netra Orbital Insight* Planet Spaceknow
--	--	---	--	--

ENTERPRISE FUNCTIONS

CUSTOMER SUPPORT ActionIQ Clarabridge DigitalGenius* Eloquent Labs Kasisto Preact Wise.io Zendesk	MARKETING AirPR BrightFunnel* CogniCor Lattice LiftIgniter Mintigo msg.ai Persado Radius Retention Science	RECRUITING Enebo Gigster* HiQ HireVue SpringRole Textio* Unlivo Wade & Wendy	SALES/FINANCE Esense AppZen Aviso* Clari Collective[] Fusemachine InsideSales Salesforce Einstein Zensight*	SECURITY Cylance Darktrace Deep Instinct Demisto Drawbridge Networks* Graphistry* LeapYear SentinelOne SignalSense Zimperium
--	---	---	---	---

AUTONOMOUS SYSTEMS

AERIAL Airware DJI DroneDeploy Lily Pilot AI Labs Shield AI* Skycatch Skydio	GROUND AdiasWorks Auro Robotics comma.ai Drive.ai Google Mobileye nuTonomy Tesla Uber Zoox	INDUSTRIAL Clearpath Robotics Fetch Robotics Harvest Automation Jaybridge Robotics Kindred* Osaro Rethink Robotics	PERSONAL Amazon Alexa Apple Siri Google Now/ Allo Facebook M Microsoft Cortana Replika	PROFESSIONAL Alien Labs Butter.ai Clara SkipFlag Slack Sudo Talia x.ai Zoom.ai
---	---	--	--	--

INDUSTRIES

AGRICULTURE Abundant Robotics AgriData Blue River Technology Descartes Labs Mavrx* Pivot Bio TerrAvion Trace Genomics Tule* UDIO	EDUCATION AltSchool Content Technologies (CTI) Coursera Gradescope* Knewton Volley	INVESTMENT FINANCE AlphaSense Bloomberg Cerebellum Capital Dataminr iSentio Kensho Quandl Sentient	LEGAL Beagle Blue J Legal Everlaw Legal Robot ROSS Intelligence Seal	MATERIALS/ MANUFACTURING Calculario Citrine Informatics Eigen Innovations Ginkgo Bioworks Sight Machine Zymergen
---	---	---	--	---

RETAIL FINANCE

RETAIL FINANCE Affirm Bettermort Earnest Lendo Mirador Tala (Ika Inventure) Wealthfront ZestFinance	TRANSPORTATION/ LOGISTICS Acerta ClearMetal Marble NAUTO PRiStop Preteck Routific
--	---

HEALTH CARE

BIOLOGICAL DATA Atomwise Color Genomics Deep Genomics* Grai iCarbonX Luminist Numerate Recursion Pharmaceuticals Verily Whole Biome	PATIENT DATA Atomwise CareScore Deep6 Analytics IBM Watson Health Numerate Oncora Medical pulseData Sentrian Zephyr Health	IMAGE DATA 3Scan Arterys Bay Labs Butterfly Network Enlifer Google DeepMind Imagja
---	--	--

MACHINE LEARNING TECHNOLOGY STACK

AGENTS AND CONVERSATIONAL INTERFACES (AGENT ENABLERS)

Automat	Kasisto	OpenAI Gym
Facebook CommAI	KITT.AI	Semantic Machines
Howdy*	Maluba	
	Octane AI	

DATA SCIENCE

Ayasdi	Domino	Seldon
BigML	Data Lab*	SparkBeyond
Dataiku	Kaggle*	Yhat
DataRobot	RapidMiner	Yseop

MACHINE LEARNING

Bonsai	deepsense.io	minds.ai
CognitiveScale	Geometric Intelligence	Nara Logics
Context	H2O.ai	Reactive
Relevant*	HyperScience	Scaled Inference
Cycorp	Loop AI Labs	SkyMind
Datacratic		SparkCognition

NATURAL LANGUAGE

Agolo	Lexalytics	MonkeyLearn
AYLIEN	Loop AI Labs	Narrative Science
Cortical.io	Luminoso	spaCy

DEVELOPMENT

AniDot	Kite	SigOpt
Bonsai	Layer 6 AI	SignifAI
Fuzzy.ai	Lobe.ai	
Hyperopt	Rainforest	

DATA CAPTURE AND ENRICHMENT

Amazon Mechanical Turk	DataSift	Paxata
CrowdAI	Diffbot*	Trifacta
Crowdfunder	Enigma	WorkFusion
Datalogue	Import.io	

OPEN SOURCE LIBRARIES

Amazon DSSTNE	DeepLearning4j	Nervana Neon
Apache Spark MLlib	H2O.ai	scikit-learn
Baidu PaddlePaddle	Keras	TensorFlow
Caffe	Microsoft Azure ML	Theano
Chainer	Microsoft CNTK	Torch7
	Microsoft DMTK	Weka
	MXNet	

HARDWARE

1026 Labs	KNUPATH	Qualcomm
Cadence Tensilica	Intel (Nervana)	Tenstorrent
Cirrascale	Isocline	
Google TPU	NVIDIA DGX-Titan X	

RESEARCH

Cogitai	NNAISENSE	Vicarious
Kimera	Numenta	
Knoggin	OpenAI	

SOURCE: SHIVON ZILUS AND JAMES CHAM

*COMPANIES IN WHICH SHIVON ZILUS AND JAMES CHAM HAVE INVESTMENTS

© HBR.ORG

Click or Tap to See a Larger Version

The State of Machine Intelligence, 2016

ENTERPRISE INTELLIGENCE

AUDIO Capio Clover Intelligence Expect Labs Gridspace* Hobvoil Nexidia Pop Up Archive* Quirious TalkIQ Twilio	INTERNAL DATA Alation* Arimo* Cycorp Digital Reasoning IBM Watson Kymd Outlier Palantir Primer Sapho*	MARKET Bottlenose CB Insights DataFox Enigma Mattermark Predata Premise Quid Tracxn	SENSOR Alluvium C3 IoT GE Predix Imubit KONUX Maana Planet OS Preferred Networks Sentenai ThingWorx Uptake	VISUAL Algoclan Captricity Clarifai Cortica Deepomatic DeepVision Netra Orbital Insight* Planet Spaceknow
--	---	---	--	--

ENTERPRISE FUNCTIONS

CUSTOMER SUPPORT ActionIQ ClaraBridge DigitalGenius* Eloquent Labs Kasisto Preact Wise.io Zendesk	MARKETING AirPR BrightFunnel* CogniCor Lattice LiftIgniter Mintigo msg.ai Persado Radius Retention Science	RECRUITING Entelo Gigster* HIQ HireVue SpringRole Textio* Unitive Wade & Wendy	SALES/FINANCE Esense AppZen Aviso* Clari Collective[] Fusemachine InsideSales Salesforce Einstein Zensight*	SECURITY Cylance Darktrace Deep Instinct Demisto Drawbridge Networks* Graphistry* LeapYear SentinelOne SignalSense Zimperium
--	---	---	--	--

AUTONOMOUS SYSTEMS

AERIAL Airware DJI DroneDeploy Lily Pilot AI Labs Shield AI* Skycatch Skydio	GROUND AdasWorks Auro Robotics comma.ai Drive.ai Google Mooliye nuTonomy Tesla Uber Zoox	INDUSTRIAL Clearpath Robotics Fetch Robotics Harvest Automation Jaybridge Robotics Kindred* Osaro Rethink Robotics	AGENTS PERSONAL Amazon Alexa Apple Siri Google Now/ Allo Facebook M Microsoft Cortana Replika	PROFESSIONAL Alien Labs Butter.ai Clara SkipFlag Slack Sudo Talia x.ai Zoom.ai
---	---	--	---	--

INDUSTRIES

AGRICULTURE Abundant Robotics AgriData Blue River Technology Descartes Labs Mavrx* Pivot Bio TerrAvion Trace Genomics Tule* UDiO	EDUCATION AltSchool Content Technologies (CTI) Coursera Gradescope* Knewton Volley	INVESTMENT FINANCE AlphaSense Bloomberg Cerebellum Capital Dataminr iSentium Kensho Quandl Sentient	LEGAL Beagle Blue J Legal Everlaw Legal Robot Ravel Law ROSS Intelligence Seal	MATERIALS/ MANUFACTURING Calculario Citrine Informatics Eigen Innovations Ginkgo Bioworks Sight Machine Zymergen
---	---	--	---	---

RETAIL FINANCE

Alfem Betterment Earnest Lendo Mirador Tala (fka InVenture) Wealthfront ZestFinance	TRANSPORTATION/ LOGISTICS Acerta ClearMetal Marble NAUTO PRTStop Preteckt Routific
--	--

HEALTH CARE

BIOLOGICAL DATA Atomwise Color Genomics Deep Genomics* Grail iCarbonX Luminist Numerate Recursion Pharmaceuticals Verily Whole Biome	PATIENT DATA Atomwise CareSkoore Deep6 Analytics IBM Watson Health Numerate Oncora Medical pulseData Sentrian Zephyr Health	IMAGE DATA 3Scan Arterys Bay Labs Butterfly Network Enlitic Google DeepMind Imagica
--	---	---

MACHINE LEARNING TECHNOLOGY STACK

AGENTS AND CONVERSATIONAL INTERFACES (AGENT ENABLERS)		
Automat Facebook CommAI Howdy*	Kasisto KIT.AI Maluuba Octane AI	OpenAI Gym Semantic Machines
DATA SCIENCE		
Ayasdi BigML Dataiku DataRobot	Domino Data Lab* Kaggle* RapidMiner	Seldon SparkBeyond Yhat Yseop
MACHINE LEARNING		
Bonsai CognitiveScale Context Relevant* Cycorp Datacratic	deepense.io Geometric Intelligence HzO.ai HyperScience Loop AI Labs	mindes.ai Nara Logics Reactive Scaled Inference SkyMind SparkCognition
NATURAL LANGUAGE		
Agilo AYLIEN Cortical.io	Lexalytics Loop AI Labs Luminoso	MonkeyLearn Narrative Science spaCy
DEVELOPMENT		
AnDdot Bonsai Fuzzy.ai Hyperopt	Kite Layer 6 AI Lobe.ai Rainforest	SigOpt SignifiAI
DATA CAPTURE AND ENRICHMENT		
Amazon Mechanical Turk CrowdAI CrowdFlower Datalogue	DataSift Diffbot* Enigma Import.io	Paxata Trifacta WorkFusion
OPEN SOURCE LIBRARIES		
Amazon DSSTNE Apache Spark MLlib Baidu PaddlePaddle Caffe Chainer	DeepLearning4j H2O.ai Keras Microsoft Azure ML Microsoft CNTK Microsoft DMTK MXNet	Nervana Neon scikit-learn TensorFlow Theano Torch7 Weka
HARDWARE		
1026 Labs Cadence Tensilica Cirrascale Google TPU	KNUPATH Intel (Nervana) Isoclone NVIDIA DGX-Titan X	Qualcomm Tenstorrent
RESEARCH		
Cogitai Kimera Knoggin	NNAISENSE Numenta OpenAI	Vicarious

SOURCE: SHIVON ZILUS AND JAMES CHAM

*COMPANIES IN WHICH SHIVON ZILUS AND JAMES CHAM HAVE INVESTMENTS

© HBR.ORG

[Click or Tap to See a Larger Version](#)

These executives are asking themselves what to do. Over the past year, machine intelligence has exploded, with \$5 billion in venture investment, a few big acquisitions, and hundreds of thousands of people reading our earlier research. As with the internet in the 1990s, executives are realizing that this new technology could change everything, but nobody knows exactly how or when.

If this year's landscape shows anything, it's that the impact of machine intelligence is already here. Almost every industry is already being affected, from agriculture to transportation. Every employee can use machine intelligence to become more productive with tools that exist today. Companies have at their disposal, for the first time, the full set of building blocks to begin embedding machine intelligence in their businesses.

And unlike with the internet, where latecomers often bested those who were first to market, the companies that get started immediately with machine intelligence could enjoy a lasting advantage.

So what should the Fortune 500 and other companies be doing to get started?

Enterprise Functions

AUDIO

Capio	Nexidia
Clover Intelligence	Pop Up Archive*
Expect Labs	Qurious
Gridspace*	TalkIQ
Mobvoi	Twilio

INTERNAL DATA

Alation*	Outlier
Arimo*	Palantir
Cycorp	Primer
Digital Reasoning	Sapho*
IBM Watson Kyndi	

MARKET

Bottlenose	Predata
CB Insights	Premise
DataFox	Quid
Enigma	Tracxn
Mattermark	

SENSOR

Alluvium	Planet OS
C3 IoT	Preferred Networks
GE Predix	Sentenai
Imubit	ThingWorx
KONUX	Uptake
Maana	

VISUAL

Algocian	DeepVision
Captricity	Netra
Clarifai	Orbital Insight*
Cortica	Planet
Deepomatic	Spaceknow

*COMPANIES IN WHICH SHIVON ZILIS
AND JAMES CHAM HAVE INVESTMENTS

SOURCE SHIVON ZILIS AND JAMES CHAM

© HBR.ORG

Make Talent More Productive

One way to immediately begin getting the value of machine intelligence is to support your talent with readily available machine intelligence productivity tools. Some of the earliest wins have been productivity tools tuned to specific areas of knowledge work — what we call “Enterprise Functions” in our landscape. With these tools, every employee can get some of the [powers previously available only to CEOs](#).

These tools can aid with monitoring and predicting (e.g., companies like Clari [forecasting client-by-client sales](#) to help prioritize deals) and with coaching and training (Textio's* [predictive text-editing platform](#) to help employees write more-effective documents).

Find Entirely New Sources of Data

The next step is to use machine intelligence to realize value from new sources of data, which we highlight in the “Enterprise Intelligence” section of the landscape. These new sources are now accessible because machine intelligence software can rapidly review enormous amounts of data in a way that would have been too difficult and expensive for people to do.

Imagine if you could afford to have someone listen to every [audio recording](#) of your salespeople and [predict their performance](#), or have a team look at every satellite image taken from space and determine what [macroeconomic indicators](#) could be gleaned from them. These data sources might already be owned by your company (e.g., [transcripts of customer service conversations](#) or sensor data predicting outages and [required maintenance](#)), or they might be newly available in the outside world (data on the [open web](#) providing [competitive information](#)).

Enterprise Intelligence

AUDIO

Capio	Nexidia
Clover Intelligence	Pop Up Archive*
Expect Labs	Qurious
Gridspace*	TalkIQ
Mobvoi	Twilio

INTERNAL DATA

Alation*	Outlier
Arimo*	Palantir
Cycorp	Primer
Digital Reasoning	Sapho*
IBM Watson Kyndi	

MARKET

Bottlenose	Predata
CB Insights	Premise
DataFox	Quid
Enigma	Tracxn
Mattermark	

SENSOR

Alluvium	Planet OS
C3 IoT	Preferred Networks
GE Predix	Sentenai
Imubit	ThingWorx
KONUX	Uptake
Maana	

VISUAL

Algocian	DeepVision
Captricity	Netra
Clarifai	Orbital Insight*
Cortica	Planet
Deepomatic	Spaceknow

*COMPANIES IN WHICH SHIVON ZILIS
AND JAMES CHAM HAVE INVESTMENTS

SOURCE SHIVON ZILIS AND JAMES CHAM

© HBR.ORG

Rethink How You Build Software

Let's say you've tried some new productivity tools and started to mine new sources of data for insight. The next frontier in capturing machine intelligence's value is building a lasting competitive advantage based on this new kind of software.

But machine intelligence is not just about better software; it requires entirely new processes and a different mindset. Machine intelligence is a new discipline for managers to learn, one that demands a new class of software talent and a new organizational structure.

Most IT groups think in terms of applications and data. New machine intelligence IT groups will think about applications, data, and models. Think of software as the combination of code, data, and a model. “Model” here means business rules, like rules for approving loans or adjusting power consumption in data centers. In traditional software, programmers created these rules by hand. Today machine intelligence can use data and new algorithms to generate a model too complex for any human programmer to write.

With traditional software, the model changes only when programmers explicitly rewrite it. With machine intelligence, companies can create models that evolve much more regularly, allowing you to build a lasting advantage that strengthens over time as the model “learns.”

Think of these models as narrowly focused employees with great memories and not-so-great social skills — idiot savants. They can predict how best to grow the business, make customers happier, or cut costs. But they’ll often fail miserably if you try to apply them to something new, or, worse, they may degrade invisibly as your business and data change.

All of this means that the discipline of creating machine intelligence software differs from traditional software, and companies need to staff accordingly. Luckily, though finding the right talent may be hard, the tools that developers need to build this software is readily available.

For the first time, there is a maturing “Stack” (see our landscape) of building blocks that companies can use to practice the new discipline of machine intelligence. Many of these tools are available as free, open-source libraries from technology companies such as Google (TensorFlow), Microsoft (CNTK), or Amazon (DSSTNE). Others make it easier for [data scientists to collaborate](#) (see “Data Science”) and manage machine intelligence models (“Machine Learning”).

If your CEO is struggling to answer the question of how machine intelligence will change your industry, take a look at the range of markets in our landscape. The startups in these sections give a sense of how different industries may be altered. Machine intelligence’s first useful applications in an industry tend to use data that previously had lain dormant. Health care is a prime example: We’re seeing predictive models that run on patient data and computer vision that diagnoses disease from medical images and gleans lifesaving insights from genomic data. Next up will be finance, transportation, and agriculture because of the volume of data available and their sheer economic value.

Your company will still need to decide how much to trust these models and how much power to grant them in making business decisions. In some cases the risk of an error will be too great to justify the speed and new capabilities. Your company will also need to decide how often and with how much

oversight to revise your models. But the companies that decide to invest in the right models and successfully embed machine intelligence in their organization will improve by default as their models learn from experience.

Economists have long wondered why the so-called computing revolution has failed to deliver productivity gains. Machine intelligence will finally realize computing's promise. The C-suites and boardrooms that recognize that fact first — and transform their ways of working accordingly — will outrun and outlast their competitors.

**The authors' fund has invested in this company.*

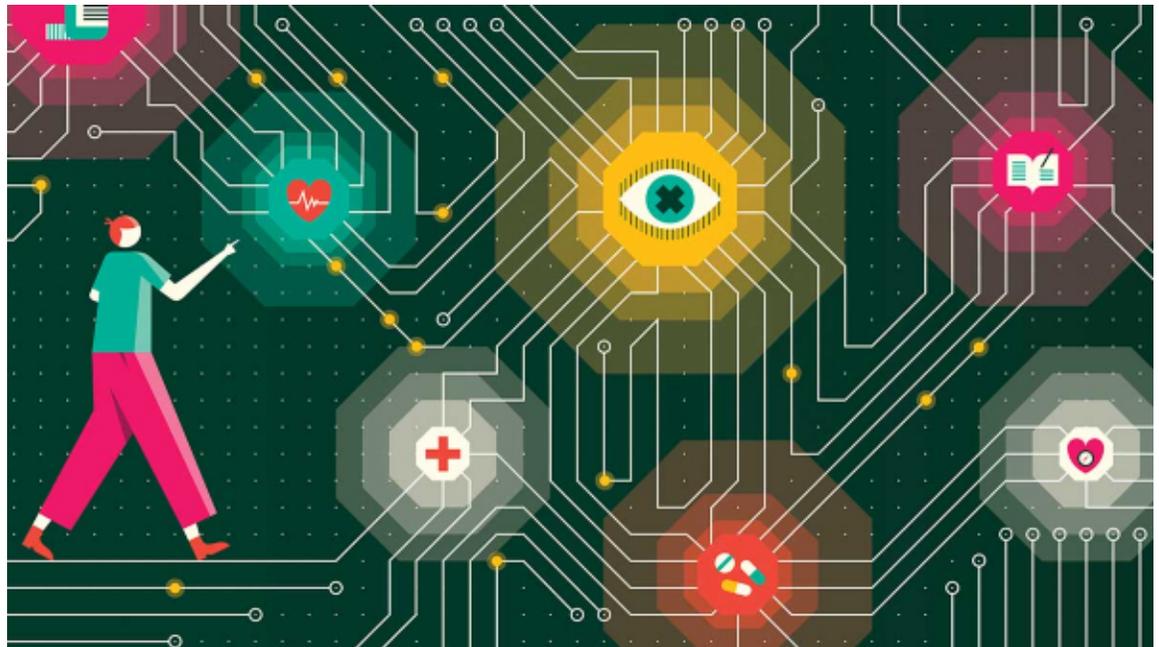
Shivon Zilis is a partner and founding member of Bloomberg Beta, which invests heavily in the future of work. She focuses on early-stage data and machine intelligence investments.

James Cham is a Partner at Bloomberg Beta where he invests in data-centric and machine learning-related companies.

INFORMATION & TECHNOLOGY

What a Visit to an AI-Enabled Hospital Might Look Like

by R “Ray” Wang
NOVEMBER 16, 2016



The combination of machine learning, deep learning, natural language processing, and cognitive computing will soon change the ways that we interact with our environments. AI-driven smart services will sense what we're doing, know what our preferences are from our past behavior, and subtly guide us through our daily lives in ways that will feel truly seamless.

Perhaps the best way to explore how such systems might work is by looking at an example: a visit to a hospital.

Setting 1: In the Lobby

You walk into the hospital and its AI-driven smart service searches to see the patient you may be there to see. (The service already knows who you are through earlier interactions.) The system uses facial recognition and matches you with the photos and the family's list of approved visitors. At the reception kiosk, the system notices you came empty-handed and that you are out of breath. Taking into account all of these variables, it asks you if you'd like to go to the 4th floor to see the patient, grab a drink, or visit the gift shop to pick up something. You respond you'd like to grab a drink and pick up a gift.

Setting 2: In the Cafeteria

The service asks if you'd like directions to the cafeteria and gift shop to be sent to your device. You say yes and follow the directions to the cafeteria. Once there, you choose a soda and pay for it. The AI smart service remembers the weather, the time, the location, and your heart rate so that it will be better prepared to help you choose in the future. You walk out: payment is automatic.

Setting 3: In the Gift Shop

The patient you're visiting is five years old, so the AI smart service offers you a fake floral bouquet, a balloon bouquet, or a stuffed animal. Because you'd been thinking about bringing a bouquet of real flowers, you find a clerk as you wonder why the smart service is wacky. The clerk looks at the patient name and informs you that the patient is allergic and provides you the same three options as before along with a coupon for 15% off on your next visit.

Setting 4: Off to the Patient's Room

With your purchase in hand, you head back to the elevators. The elevator tells you that your family member is on the 4th floor, the pediatric unit. As you get off the elevator, the smart system offers you an option to download all the latest information on pediatric health and wellness, the map shows you the fastest route, and the system reminds you that visiting hours are over at 9 p.m. You enter the room to visit your family member. Since you're not next of kin, it knows not to give you access to the latest chart. Meanwhile, the smart system asks you if you would like to chat with a caregiver or connect to another service.

Behind the Scenes

That's not a far-fetched vision of what our interactions with AI might soon look like. But let's go a little deeper to see how AI is powering the hospital experience.

The AI loop includes seven steps:

1. **Perception describes what's happening now.** Perception starts by providing both obvious and non-obvious observations to provide a baseline of information to help inform recommendations. In our example, someone had to tell the system the physical layout of the hospital to give you directions. This includes information such as geo-spatial locations, restricted access areas, and equipment in locations. Much of this is programmed into the system by humans.
2. **Notification tells you what you asked to know.** Notifications through alerts, workflows, reminders, and other signals help deliver additional information through manual input and learning. Rudimentary programming will lay out if/then paths for actions the system might take. A patient and their family, for instance, may receive notifications on when to take medicines.
3. **Suggestion recommends action.** Suggestions build on past behaviors and modify over time based on weighted attributes, decision management, and machine learning. Applying best practices and customer journey mapping, organizations will program rules and policies to surface recommendations such as pointing out what floor a patient is located on, providing dietary compliant options on the cafeteria menu, and suggesting when to rest based on vital signs.
4. **Automation repeats what you always want.** Automation enables leverage as machine learning matures over time and tuning. If the system knows you always drink “Tea, Earl Gray, Hot,” then it’s ready to comply. The algorithms will start to learn preferences based on the interactions in the neural networks. Automated alerts may pop up to inform a patient when to take medicine based on the previous patterns.
5. **Prediction informs you of what to expect.** Prediction starts to build on deep learning and neural networks to anticipate and test for behaviors. If the system knows that you had a cup of coffee before leaving the house for a 45-minute drive to the hospital, it might show you the path to the restroom when you arrive.
6. **Prevention helps you avoid bad outcomes.** Prevention applies cognitive reckoning to identify potential threats. Using what is known, a drug interaction may be avoided based on the existing patient chart and prescriptions written by the physician.
7. **Situational awareness tells you what you need to know right now.** Situational awareness comes close to mimicking human capabilities in decision-making. Patients and their providers can now find peace of mind as the system only provides the relevant information that’s needed to the right individual. If a patient crashes, the telemetry will share the most pertinent information required with the nurses, doctors, operating room, medical tech, and family on a prioritized need-to-know basis.

This combination will allow the emergence of AI-driven smart services. These new AI-driven smart services rely on five key components:

1. **Our digital footprints and data exhaust will allow AI services to build a profile.** Every individual, device, or network provides some information. That digital footprint or exhaust could come from facial analysis, a network IP address, or even one's walking gait. Using AI and cognitive reckoning, systems can start to analyze patterns and correlate identity. That means that AI services will know us as individuals across difference contexts.
2. **Immersive experiences enable a natural interaction.** Context, content, collaboration, and channels come together in all AI-driven services to deliver immersive and unique experiences to each of us. The services will use context attributes such as geospatial location, time of day, weather, heart rate, and even sentiment — combined with what the service knows of our identity and preferences — to improve relevancy and deliver the appropriate content. Sense-and-respond mechanisms will enable collaboration among participants and machines through conversations and text dialogs. Channels include all interaction points such as mobile, social, kiosks, and in-person. The goal is natural user experiences based on identity.
3. **Personalization at scale delivers digital services.** Anticipatory analytics, catalysts, and choices interact to allow mass personalization at scale. Anticipatory analytics adjust the customer experience depending on what the system knows about the customer. Catalysts provide offers or triggers for customer responses. Choices are just that — the decisions customers make themselves. By combining these three approaches, AI systems will design experiences on the fly, based on each customer's identity, past preferences, and needs at the time.
4. **Value exchange completes the orchestration of trust.** Once an action is taken, value exchange cements the transaction. Monetary, non-monetary, and consensus exchanges are three common forms of value exchange. While monetary value exchange might be the most obvious, non-monetary value exchange (including recognition, access, and influence) often provide a compelling form of value. Meanwhile, a simple consensus or agreement can also deliver value exchange on the veracity of a medical claim with an insurer or an agreement on a patient treatment protocol.
5. **Cadence and feedback continues and AI-powered learning cycle.** Powered by machine learning and other AI tools, smart services consider the cadence of delivery: one time, ad-hoc, repetitive, subscription-based, and threshold-driven. Using machine learning techniques, the system studies how the smart services are delivered for certain patient populations and applies this to future interactions. For example, the system can detect which options are more likely to please children than adults and when to offer these suggestions.

While all of this AI-driven automation may make some people anxious, fears of robots taking over the world are overblown. Successful AI-driven smart services will augment human intelligence just as machines augmented physical capabilities. By enabling reduction of errors, improving speed of decisions, identifying demand signals, predicting outcomes, and preventing disasters, AI-driven smart services will play a key role in all of our lives.

R “Ray” Wang is Principal Analyst and CEO at [Constellation Research](#) and the author of *Disrupting Digital Business: Create an Authentic Experience in the Peer-to-Peer Economy*

INFORMATION & TECHNOLOGY

How Predictive AI Will Change Shopping

by Amit Sharma
NOVEMBER 18, 2016



Imagine you're about to leave the house to pick up your kids. As you grab your keys, you hear a voice from the device on your coffee table: "It looks like you'll use the last of your milk tomorrow, and yogurt is on sale for \$1.19. Would you like to pick up an order from Trader Joe's, for a total of \$5.35?" You say yes, and Alexa confirms. The order will be ready for curbside pickup, on the way home from your kids' school, in 15 minutes.

This future scenario isn't so far off. Amazon, Facebook, Google, and Apple are accelerating consumer expectations and what's technologically possible, from same-day delivery to machine-powered

image recognition. You can call an Uber with Siri and book a flight entirely through a Facebook Messenger bot.

Responsive retail has peaked, and we're about to enter the era of predictive commerce. It's time for retailers to help people find products in their precise moment of need — and perhaps before they even perceive that need — whether or not they're logged in or ready to click a “buy” button on a screen. This shift will require designing experiences that merge an understanding of human behavior with large-scale automation and data integration.

Machine Learning Beyond Forecasting

Retail giants have been using machine-learning algorithms to forecast demand and set prices for years. Amazon patented [predictive stocking](#) in 2014, and saying that AI, machine learning, and personalization technologies have improved since then is an understatement. Retailers need to think more like tech companies, using AI and machine learning not just to predict how to stock stores and staff shifts but also to dynamically recommend products and set prices that appeal to individual consumers.

Say you're on a business trip and realize you forgot your phone charger. You'll pay a premium for a new one delivered to your hotel room before an all-day meeting. An electronics retailer might also predict that you want new headphones. It can offer you a deal on a noise-canceling pair at a price that accounts for current pricing on Amazon, in-store inventory at Best Buy, the current rates for on-demand couriers, and the fact that you're taking a red-eye flight home tomorrow.

This level of prediction requires detecting subtle patterns from massive data sets that are constantly in flux: consumers' purchase histories, product preferences, and schedules; competitors' pricing and inventory; and current and forecasted product demand. This is where AI and machine learning comes in and where companies are investing. Etsy just [acquired a company that specializes in machine learning](#) to make its searches more predictive by surfacing nuanced product recommendations that go beyond simple purchase histories or preferences. This is the natural evolution of product recommendations, one that will be the standard for years to come.

Realizing the Potential of Connected Devices and Data

Predictive retail involves inspiring consumers in different contexts — before, during, and after a purchase. Commerce is already becoming less of a deliberate activity than an organic part of how we experience daily life. It's not just smartphones that make browsing and buying spontaneous; Amazon's Dash buttons and Alexa-powered Echo device are enabling purchases in the home. You can hit the Tide Dash button in your laundry room when you see that you're running low on detergent, or ask Alexa to [order your mom a bouquet of flowers](#) when you remember that her birthday is next week. This is just the beginning.

The next generation of smart assistants and connected devices will learn from user habits and pick up on behavioral and environmental patterns in order to make these experiences more predictive.

Devices like the Echo will access data from everyday interactions to predict specific opportunities for a transaction.

There's also huge potential for connected devices in retail stores to predict consumer behavior and respond to individual needs. Many stores are [already using smartphones](#) to follow customers' activity and deliver context-specific offers. It's not a stretch to imagine that the evolution of biometrics, identity technologies, and location sensors will allow retailers to personalize content based on factors such as how you're feeling, how much time you have to browse, and whether you're coming from the office or you've just finished working out.

Retailers will need to program brick-and-mortar experiences with the same targeting and personalization they offer online. Think about walking past Nordstrom and receiving a notification for an offer on a new pair of sneakers. Your current pair is worn down from running almost 500 miles — all logged by a chip in the sole that sends data to your fitness app. You swipe the notification to select the styles you want to try on, and an in-store map guides you to an associate waiting with your shoes.

Embracing Human-Centered Design

The future of predictive retail requires designing new ecosystems for commerce. These systems will be built around the human, rather than around a particular device or around online or offline experience. These systems will need to incorporate human connection and storytelling, spatial design and context, and a lot of data.

Many retailers are getting ahead of this shift by creating innovation labs — teams and spaces dedicated to incubating new ideas and testing digital experiences that connect the online and in-store worlds. Sephora's [Innovation Lab](#) is a great example. The brand introduced a “store mode” for its mobile app, which integrates a user's online shopping cart and Beauty Insider loyalty card to remind them of the products they've saved, the points they've earned, and the benefits available to them, such as a free makeover.

Retail chains, brands, and e-commerce companies are also collaborating to bring new ideas to life. Several years ago, Westfield Malls' lab worked with eBay to [build 10-foot-tall interactive screens](#) in its San Francisco shopping center. Shoppers swiped these screens to browse products from brands like Rebecca Minkoff and Sony, which they could purchase directly on mobile.

There's huge potential to layer predictive capabilities on top of this AI-driven infrastructure. Imagine a store window that connects with your phone to display personalized content. For instance, you might see gifts for your partner's birthday or swimsuits for your next vacation, customized based on the boards you follow on Pinterest and the brands you follow on Instagram. By connecting data from multiple sources and designing for the user, retailers can create more-relevant experiences that pull you into a store, website, or app. Even more powerful, they can predict what you want before you do.

Considering Privacy, Building Trust

There's almost always some trade-off between privacy and personalization; this has been true for every generation of technology. Retailers need to move forward with transparency, respect, and security as their priorities. They also need to show value. Google has done this well, not just with personalized search results but also with services such as Google Now, which integrates with your calendar and Google Maps to alert you that traffic to your meeting is heavier than usual and tells you when you should leave the office to arrive on time.

Many of us are inclined to share personal information for experiences that are magical and valuable — and that we can't get elsewhere. Retailers will need to create experiences that make this magic and value apparent. The revolution is already under way. Tomorrow, people will expect even faster and more-intelligent service than they do today. At a point in the very near future, the expectation will shift from on-demand to predictive commerce. It's time for retailers to get ahead of that change.

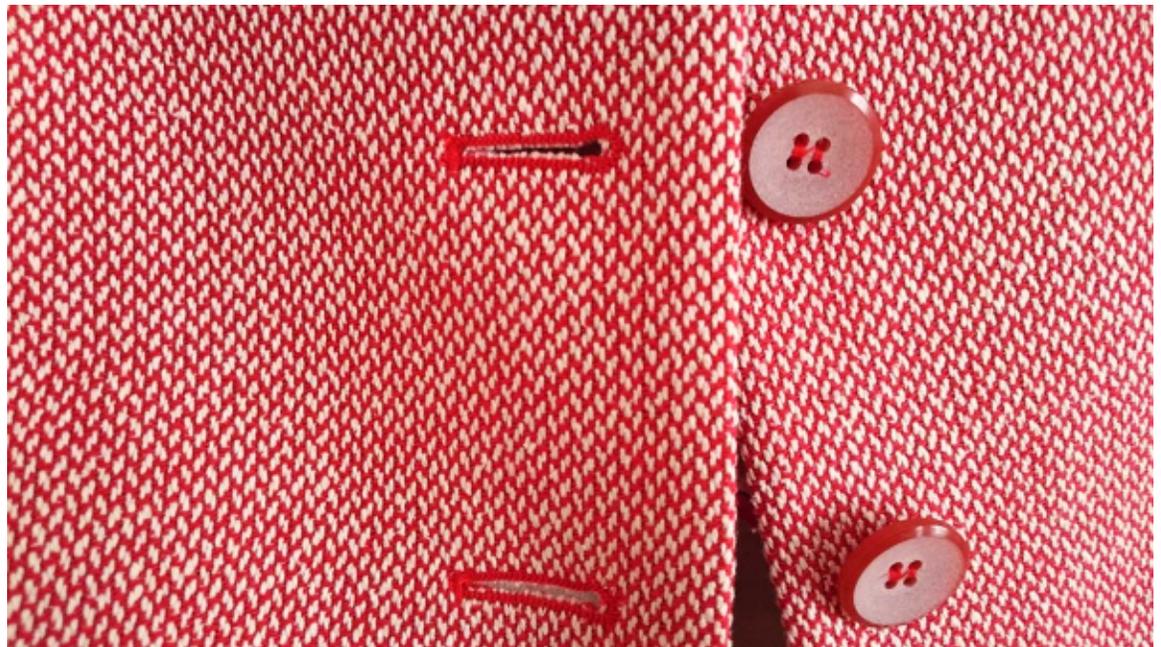
Amit Sharma is the founder and CEO of [Narvar](#), a post-purchase experience platform, and a former executive at Apple and Walmart.

INFORMATION & TECHNOLOGY

How One Clothing Company Blends AI and Human Expertise

by H. James Wilson, Paul Daugherty and Prashant Shukla

NOVEMBER 21, 2016



When we think about artificial intelligence, we often imagine robots performing tasks on the warehouse or factory floor that were once exclusively the work of people. This conjures up the specter of lost jobs and upheaval for many workers. Yet, it can also seem a bit remote — something that will happen in “the future.” But the future is a lot closer than many realize. It also looks more promising than many have predicted.

Stitch Fix provides a glimpse of how some businesses are already making use of AI-based machine learning to partner with employees for more-effective solutions. A five-year-old online clothing retailer, its success in this area reveals how AI and people can work together, with each side focused on its unique strengths.

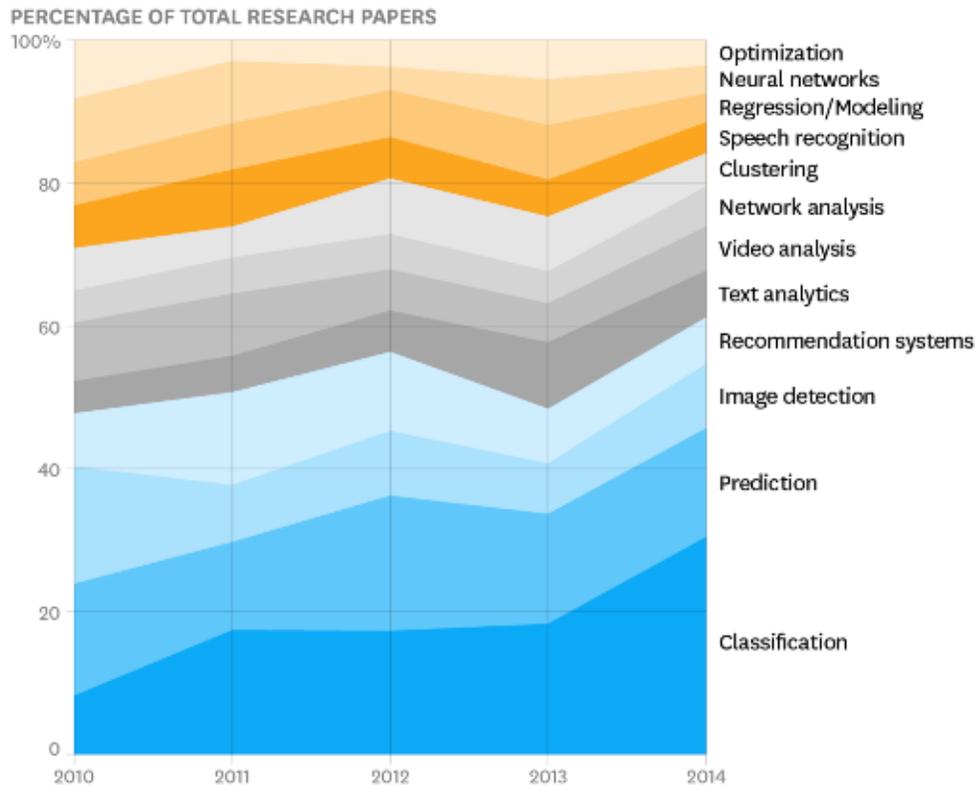
The company offers a subscription clothing and styling service that delivers apparel to its customers' doors. But users of the service don't actually shop for clothes; in fact, Stitch Fix doesn't even have an online store. Instead, customers fill out style surveys, provide measurements, offer up Pinterest boards, and send in personal notes. Machine learning algorithms digest all of this eclectic and unstructured information. An interface communicates the algorithms' results along with more-nuanced data, such as the personal notes, to the company's fashion stylists, who then select five items from a variety of brands to send to the customer. Customers keep what they like and return anything that doesn't suit them.

Stitch Fix's [approach](#) illustrates three lessons about how to combine human expertise with AI systems. First, it's important to keep humans in the business-process loop; machines can't do it alone. Second, companies can use machines to supercharge the productivity and effectiveness of workers in unprecedented ways. And third, various machine-learning techniques should be combined to effectively identify insights and foster innovation.

As research we've conducted across industry and academia shows, companies have an unprecedented opportunity to tap ongoing advances in AI and machine learning research to reinvent business processes. For instance, in analyzing a five-year sample of almost 1,150 papers, we identified at least 12 techniques, visible in the chart below, that can be readily applied and combined with each other within a process. Stitch Fix is already applying several of these machine learning techniques in service delivery and R&D — and other companies can follow its lead.

12 Common Machine Learning Techniques

These approaches were identified via an analysis of almost 1,150 research papers over a five-year period.



SOURCE ANALYSIS OF A STANFORD UNIVERSITY RESEARCH PAPER DATABASE BY THE ACCENTURE INSTITUTE FOR HIGH PERFORMANCE

© HBR.ORG

Some extremely successful companies have made great use of recommendation engines to boost sales or improve customer satisfaction. When it comes to recommendations, is there room for improvement in the way Amazon and Netflix operate?

Stitch Fix, which lives and dies by the quality of its suggestions, has no choice but to do better. And it can't rely solely on machines to do this. The company collects as much information about a client as it can, in both structured and unstructured form. Structured data includes surveys with personal information such as body measurements and brand preferences. Unstructured data can be derived from social media accounts, such as Pinterest, or through online notes from people about why they are buying new clothes, such as a special occasion, a change of season, or because a certain new style caught their eye.

The automated recommendation system is at its best when dealing with structured data. But to make sense of unstructured data, people and their judgment are needed. Say a client wants a new pair of stylish jeans, an item that's notoriously tricky to fit right to a person's measurements. To start, the algorithm finds jeans (across a range of fabrics, styles, and even sizes) that other clients with the same inseam decided to keep — a good indicator of fit.

Next, it's time to pick the actual pair of jeans to be shipped. This is up to the stylist, who takes into account a client's notes or the occasion for which the client is shopping. In addition, the stylist can include a personal note with the shipment, fostering a relationship, which Stitch Fix hopes will encourage even more useful feedback.

This human-in-the-loop recommendation system uses multiple information streams to help it improve. The algorithm absorbs feedback directly from the client — whether or not she or he (the company added men's options in late September) decided to keep an item of clothing. And the stylist improves and adjusts based on cues gleaned from client notes and with insights from previous interactions with the customer.

The company is testing natural language processing for reading and categorizing notes from clients — whether it received positive or negative feedback, for instance, or whether a client wants a new outfit for a baby shower or for an important business meeting. Stylists help to identify and summarize textual information from clients and catch mistakes in categorization. Because the algorithms are never far from human oversight, Stitch Fix can confidently test new machine learning technologies without worrying that the experimentation will [disrupt the client experience](#).

How does the role of machine-augmented stylist really play out? Stitch Fix employs more than 2,800 stylists, dispersed across the country, all of them working from home and setting their own hours. In this distributed workforce, stylists are measured by a variety of metrics, including the amount of money a client spends, client satisfaction, and the number of items a client keeps per delivery. But one of [the most important factors](#) is the rate at which a stylist puts together a collection of clothes for a client.

So speed matters, as does accuracy. This is where the work interface comes into effect. To enable fast decision making, the screen on which a stylist views recommendations shows the relevant information the company keeps about a client, including apparel and feedback history, measurements, and tolerance for fashion risks — it's all readily accessible. Moreover, the recommendations are tuned to reduce the amount of time a stylist spends [searching for items in vain](#).

In order to make sure the stylists' selections are effective, Stitch Fix's system [can vary the information a stylist sees to test for bias](#). For instance, how might a picture of a client affect a stylist's choices? Or knowledge about a client's age? Does it help or hinder to know where a client lives?

By measuring the impact of modified information in the stylist interface, the company is developing a systematic way to measure improvements in human judgment.

Even with the constant monitoring and algorithms that guide decision making, according to internal surveys, Stitch Fix stylists are [mostly satisfied with the work](#). And this type of work, built around augmented creativity and flexible schedules, will play an important role in the workforce of the future.

The fashion industry is no stranger to fast cycles of learning. One of the great benefits Stitch Fix sees from collecting and analyzing so much data is an ability to predict trends. For example, the company's engineers are developing machine learning classifiers to find trends by using the simple yes-or-no decision that a client makes when they buy an item or send it back. From this seemingly simple data, the team has been able to uncover which trends change [with the seasons](#) and which fashions are going out of style.

Another benefit of all this data is research aimed at developing a computer vision system that can interpret style and extract a kind of style measurement from images of clothes. The system itself would undergo unsupervised learning, taking in a huge number of images and then extracting patterns or features and deciding what kinds of styles are similar to each other. This "auto-styler" [could be used to automatically sort inventory and improve selections for customers](#).

In addition to developing an algorithmic trend-spotter and an auto-styler, Stitch Fix is developing brand new styles — fashions born entirely from data. The company calls them "[frankenstyles](#)." These new styles are created from a "genetic algorithm," modeled after the process of natural selection in biological evolution. The company's genetic algorithm starts with existing styles that are randomly modified over the course of many simulated "generations." Over time, a sleeve style from one garment and a color or pattern from another, for instance, "evolve" into a whole new shirt.

As a machine learning-native company, Stitch Fix bears watching. [Forbes estimates](#) that in 2015 the company brought in \$250 million in revenue, and it predicts a 50% increase in 2016. At the same time, Stitch Fix has tripled its workforce over the past two years. The company's approach shows that humans and machines can augment each other, producing better process results, at scale, than either could do alone.

We're only at the beginning of the era of artificial intelligence. Some upheaval is to be expected. But we are starting to see how AI can change industries, improve productivity, and even benefit a new generation of employees.

H. James Wilson is Managing Director of Information Technology & Business Research at the Accenture Institute for High Performance.

Paul Daugherty is the chief technology officer of Accenture, based in New York City.

Prashant Shukla is a Research Fellow at the Accenture Institute for High Performance.

ANALYTICS

Bots That Can Talk Will Help Us Get More Value from Analytics

by Stuart Frankel
NOVEMBER 24, 2016



Over the past few years, much has been made of the rise of big data. And yet [research](#) from TDWI states that at organizations where 50% of employees have access to business intelligence tools, only 20% of that group actually use them. Part of the problem is that systems are often hard to use. Another challenge is low rates of data literacy.

To get around these issues, many organizations have relied on visualizations to display information gleaned from data. While a picture may be worth a thousand words, the same can't always be said for these charts and graphs. There are a range of [causes for data misinterpretation](#), including insufficient domain expertise and lack of training in statistical thinking.

All of this suggests that trying to force people to become data literate is an uphill battle. But this is actually becoming less necessary thanks to the rise of artificial intelligence (AI) — and, in particular, advanced [natural language generation](#) (advanced NLG), a subfield of AI. Advanced NLG platforms — including, full disclosure, the one my company has built — start by understanding what the user wants to communicate. Then these systems perform the relevant analysis to highlight what is most interesting and important, identify and access the data necessary to tell the story, and finally deliver the analysis in a personalized, easy-to-consume way: as a narrative. Gartner predicts that by 2018, advanced NLG will be integrated into the majority of smart data discovery platforms and that 20% of business content will be generated by machines.

When you consider that the alternative is having employees spend hours pouring over complex data sets or equally inaccessible data visualizations, the value of advanced NLG starts to come into sharp focus. A few real-life or lightly fictionalized examples help drive the point home.

Alexa and the rise of conversational interfaces

Conversations with our financial advisors, doctors, or fellow employees are all about sharing information that's important and impactful. Why shouldn't our interactions with intelligent systems be the same? Imagine sitting in your family room on a Saturday night and your spouse inquires about the status of your financials. You ask Alexa and the conversation goes something like this:

You: Alexa, how are my investments doing?

Alexa: Do you want a high level summary or should we go straight to the details?

You: A summary is fine.

Alexa: Year-to-date you are up 5% which is slightly ahead of your 2016 investment return target. The mix of stocks, bonds, and alternative investments is generally the same as this time last year, but you are holding slightly more cash than you expected to be holding at this point in the year. This could just be timing, but you may want to look into it further. If you want, I can invest the excess cash and reallocate your holdings accordingly.

Conversations with systems that have access to data about our world will allow us to understand the status of our jobs, our businesses, our health, our homes, our families, our devices, and our neighborhoods — all through the power of advanced NLG. It will be the difference between getting a report and having a conversation. The information is the same but the interaction will be more natural.

Inventory management at the grocery store

As another example, large supermarket chains are very concerned about their inventory. That's especially true for perishable goods like the ones you'll find at the deli counter or in the produce department, because that's where their margins are earned. For stores like these, managing inventory precisely is mission-critical: Getting it wrong has a direct impact on profits.

Not surprisingly, the executives who run these chains are pretty focused on data. They have lots of it, like supply chain, production, and sales data, which they make available to store managers in the hopes of giving them the information they need to manage inventory flows effectively. The problem, of course, is that deli managers aren't necessarily trained to be business analysts, let alone data scientists.

Yet when advanced NLG becomes the main communication vehicle of their reporting process, interacting with this data gets easier. Instead of charts and graphs or raw data, store managers get customized narratives written in plain English that offer both analysis and prescriptive advice. They understand whether they should order a certain type of deli meat this week to meet demand, or whether they should wait until next week. Language and predictive modeling combine to take the guesswork out of the equation, giving store managers clear guidance on how they're doing in absolute and relative terms, as well as what they can do to improve.

Real-time financial analysis

Another example comes from a financial intelligence firm. The firm regularly produces large reports documenting organizations' performance, which are shared with internal committees who are then tasked with making decisions that have broad implications for the financial institution landscape. Whereas those reports traditionally took an analyst two hours to generate, NLG is being used to generate timely, accurate reports in seconds. Not just the charts and graphs – advanced NLG makes it possible for a computer to write the report itself. With the automated research in hand, the analysts can focus on enhancing the qualitative portions of the report to enable the committees to make more informed decisions.

Call center coaching

Big companies often have big call centers. I know a company in the financial services industry that has a staggering 13,000 employees in its call center. As you can probably imagine, the logistics of managing a call center workforce even a fraction of that size can be tough.

In organizations of this size, individual managers may not have the time or resources to conduct frequent performance reviews and deliver on-going personalized training, though [92% of managers](#) see high value in such communications. On top of that, people don't want to be given numbers or charts illustrating how they can do better, they want “[corrective feedback](#)” defined as suggestions for improvement or explorations of new and better ways to do things.

With advanced NLG, performance and call activity data can automatically be analyzed to generate weekly personalized coaching reports that convey in simple and conversational terms how individual employees are doing, what behaviors to improve, and their progress against goals to inspire and encourage change. Here's an example of feedback from one of these reports:

“With an 88% customer satisfaction rating, you have demonstrated solid service and commitment to our customers, though you still have room to strive for excellence. Work on executing your disconnect rate and speak to other representatives to find ways to improve performance.”

A well-known innovative financial services organization started testing this approach internally and are able to deliver personalized reports at a scale and frequency not previously possible.

As advanced NLG transforms from a niche emerging technology into the default communication layer that we put on top of data, the idea that every person needs a set of specific, technical skills in order to interact with data will seem ludicrous.

Our expectations of data are rapidly reaching the same tipping point of other types of innovations that we don't give a second thought to anymore. Phones that allow us to carry limitless knowledge in our pocket, free video conferencing with anyone across the world or smart houses that regulate themselves are now the norm and not the exception.

While we take them for granted now, there was a time when these advances seemed just as improbable. Looking ahead to the next few years, the same shift will occur to data as enabled by advanced NLG. Additionally, this growing movement will help build trust in intelligent systems as information will be delivered through in a familiar, conversational way and the systems will be able to explain in clear language why and how they came to conclusions.

People have always communicated through stories and language, why should we expect them to change now?

Stuart Frankel is the CEO and a co-founder of Narrative Science, a company working on advanced natural language generation for the enterprise.

INFORMATION & TECHNOLOGY

Prepare Your Workforce for the Automation Age

by Christoph Knoess, Ron Harbour and Steve Scemama

NOVEMBER 23, 2016



The internet has dramatically changed the way companies operate. Massive data storage capacity, super-fast data transmission and mobility devices—along with slick application program interfaces—have left companies scrambling to adapt.

Today, innovations in digitization and robotization are quickly laying the foundation for another disruptive corporate transformation. For example, Anheuser-Busch, working with Uber (and Otto), [just delivered 2,000 cases of Budweiser in a self-driving truck](#). [Commerzbank has announced plans to digitize 80% of its processes](#) within three years. We estimate that robotization, digitization, digital self-services, distributed digital advice and sales, and robo-advisors could result in a 60-70%

reduction in the workforces of service providers, from financial services to telecom. Manufacturers have already seen reductions, albeit at lower levels. The pace of robot adoption may surprise us, just as the internet spread more quickly than many anticipated.

But companies will only be able to realize productivity gains from these new operating models if they skillfully manage the soft side of their automation transformation—the people in what will be a vastly different organization. As companies introduce software bots and digital self-service, and as they transform assembly lines, they must bring along their key employees, leaders, and customers as they redefine jobs, career paths, workforce management, and social contracts. Executives must think carefully about how to best match people and machines, bearing in mind that many of the decisions they make today will have a long-tail effect on workforce composition, productivity, and profits for years to come.

As workforces hollow out, the remaining employees will be highly specialized and experienced business/technology hybrids—a new breed of professional who can work in highly distributed environments and shift from managing people to managing experiences and technology. In the back-office, the lights will dim, as work is shifted to the customer or other parts of the value chain. In the middle office, risk and compliance management will largely watch bots that are not prone to human error or fraud, supported by sophisticated models to predict quality and compliance issues. In the front office, automation based on predictive analytics will leave only managers who can control sophisticated robo-advisors trusted by customers. Salespeople will be disrupted as customers link with algorithmic bots to obtain products contextually presented at the (digital) point of need—often dominated by the global platforms that link retail, financial services, entertainment and communications in sticky ways.

Will this happen over night? No. Disruption rarely occurs as soon as expected. Freeways full of driverless cars and beer trucks are still far off, because of technological and regulatory limitations. But the inflection point always happens *faster* than expected. As always with technology adoption, there is an S-curve, already being scribbled by early adopters; when the inflection point is reached, expect sudden acceleration. So early preparation is needed.

Lessons from automakers

To figure out how to effectively integrate software robots and digitization, executives can take some cues from the automotive industry. After decades of introducing physical robots and automation, some automakers' factories require three to six months to launch an entirely new vehicle, while others need no more than a day. But the most automated carmakers are not necessarily the most efficient.

Instead, the nimblest factories have been pragmatic about integrating automation so that new processes can run smoothly, with continuous improvement—so that only the simplest, most repetitive processes are automated. Over the past two decades, leading automakers have automated

their paint and body shops—where they see the greatest gains—but have also retained and retrained the people required to quickly redesign products and processes on more complex tasks. Automakers have used automation gains to configure and fund increased customization; assembly lines are run by humans who build customized vehicles, choosing from as many as 55,000 parts.

The lesson: Change needs to be evolutionary, even if the impact of automation is ultimately revolutionary. Don't throw away your core capability, until you are sure automation is better, faster, cheaper. As you prepare for the inflection point, be pragmatic about cost-benefit tradeoffs. Think about the overall organization in an automated world. Be mindful of the critical skills you need to retain, and the skills you need to build up. Take a full end-to-end view. Think both short term and long term. Build strategic advantage through the gains you achieve, beyond cost.

Parallel work streams

The lessons to date on driving digitization and robotics suggest operating on a dual track. Strike the right balance between implementing short-term automation fixes and opportunities, while seeking solutions to problems that will determine success in the long run. It takes time for clients and employees to adapt to monumental change. Thus, it is as important to get long-term organizational change underway as it is to rapidly exploit near-term efficiency.

To that end, managers should develop a list of 10 to 15 processes that bots can quickly improve. Test and learn, both in the application of the right bots to the right problem, and how to redesign processes end-to-end to maximize results. Simultaneously test and learn on the soft side of automation. Blueprint the broader impact on roles, skills, controls, leadership, workforce and talent management, and social contracts.

By doing so, managers can move critical employees and clients closer to their longer term automation ambitions—which can be funded at least in part with returns from the earlier automation of simpler tasks.

Mold the organization

As more processes are digitized in every part of an organization, executives must think at a macro level about the entire enterprise, even as the organization is changing. How do you hire today for a diminished workforce 10 years out? When more and more of your people are replaced by bots, how do you lead, enforce quality control, and audit? The key to navigating through the coming automation age will be identifying and retaining (retraining) the employees who can make one transition after another.

Companies will experience huge changes as physical infrastructure disappears, offshore capabilities are repatriated, more services become self-service and virtual, and customers begin to interact more with robots. Automation will transform not just production, but operating models. Start early to shift

the leadership mindset. Actively drive customer adoption by adjusting your pricing and loyalty incentives to encourage early adopters, while simultaneously having your own employees co-service them with the same tools. Initially, customers should be allowed easy and seamless access to people as a fallback to self-service and other digital offerings. Assess the impact on your social contracts to both your employees and the communities you serve. Engage local stakeholders and unions early in discussions of how you can continue to give back to the community with a smaller workforce. Build a picture of the future state and work backwards.

There is time before the inflection point—time to prepare with purpose and pilots. Focus not just on the technology and analytics, the shiny object, but on people and a new form of leadership.

Christoph Knoess is the global head of Oliver Wyman’s strategic IT and operations practice and head of its digital practice.

Ron Harbour is a senior partner in Oliver Wyman’s global automotive manufacturing practice.

Steve Scemama is leading Oliver Wyman’s capital markets IT and digital practice.

TECHNOLOGY

Why You Shouldn't Swear at Siri

by Michael Schrage
OCTOBER 21, 2016



Stop swearing at Siri. Quit cursing Cortana. As digital devices grow smarter, being beastly toward bots could cost you your job.

As machine learning and artificial intelligence capabilities proliferate, digital interface design sensibilities begin to accelerate from [skeuomorphic](#) to [anthropomorphic](#). Consider how [Slack](#) and [Hipchat](#) now use bots that go beyond automating interfaces to facilitate smarter and more engaging user experiences. That's Chatbot 1.0. Ongoing algorithmic innovation assures that next-generation bots will be far sharper and [more empathic](#).

For Apple, Google, Amazon, Microsoft, Tencent, and Facebook, making bots and software assistants more human — and more humane — has become a strategic UX imperative. Being affective matters as much as being effective.

But because humans don't (yet) attach agency or intelligence to their devices, they're remarkably uninhibited about abusing them. Both academic research and anecdotal observation on man/machine interfaces suggest raised voices and vulgar comments are more common than not. It's estimated that about 10% to 50% of interactions are abusive, [according to Dr. Sheryl Brahn in a TechEmergence interview](#) late last year.

These behaviors are simply not sustainable. If adaptive bots learn from every meaningful human interaction they have, then mistreatment and abuse become technological toxins. Bad behavior can poison bot behavior. That undermines enterprise efficiency, productivity, and culture.

That's why being bad to bots will become professionally and socially taboo in tomorrow's workplace. When "deep learning" devices emotionally resonate with their users, mistreating them feels less like breaking one's mobile phone than kicking a kitten. The former earns a reprimand; the latter gets you fired.

Just as one wouldn't kick the office cat or ridicule a subordinate, the very idea of mistreating ever-more-intelligent devices becomes unacceptable. While not (biologically) alive, these inanimate objects are explicitly trained to anticipate and respond to workplace needs. Verbally or textually abusing them in the course of one's job seems gratuitously unprofessional and counterproductive.

Crudely put, smashing your iPhone means you have a temper; calling your struggling Siri inappropriate names gets you called before HR. Using bad manners with smart technologies can lead to bad management.

Your rudeness invites their underperformance. Abuse bespeaks a hostile work environment that undermines how bots learn. [Microsoft's Tay](#) offers a painfully superb — and superbly painful — real-world case study of how networked abuse shapes UX. Less than a day after Microsoft research released its unsupervised machine learning twitterbot, Tay became a "chatbot from hell" — tweeting a stream of increasingly nasty, racist, and homophobic comments until [Microsoft pulled the plug](#).

"The problem was Microsoft didn't leave on any training wheels, and didn't make the bot self-reflective," Brandon Wirtz said in a recent [LinkedIn article](#) about the situation. "[Tay] didn't know that she should just ignore the people who act like Nazis, and so she became one herself."

By contrast, look at how Microsoft's chatbot thrived in China's more regulated and inhibited digital environment. Xiaoice was able to avoid Tay's issues because Chinese digital culture effectively sanctions certain forms of expression. Where Tay is no more, Xiaoice enjoys [over 40 million users in China and Japan](#).

Digitally compromising and compensating for rude and abusive customers is one thing, but making bots beholden to rude and abusive employees is quite another. Much as certain kinds of behaviors became crude, vulgar, and anachronistic when women entered the workforce en masse, the notion that cursing smart machines is acceptable while swearing at dumb colleagues is not seems doomed to disappear.

Don't think for a moment there won't be immaculate digital records of enterprise misbehaviors. You can be sure that Apple has excellent insight into how often Siri is insulted and abused. Do you really think Amazon doesn't similarly instrument Alexa? We can confidently bet that sysadmins log chatbot exchanges to learn what actions — and which users — inspire criticism. Networked smartphones could easily track bad acts. In fact, in many (if not most) machine learning systems, insult, and abuse could easily become part of the real-world training. Perhaps they can even be programmed or trained to complain about their mistreatment.

The “deep learning” deep insight will be that it is foolish to disrespect adaptive intelligence regardless of its origins. Of course, there's no shortage of excellent science fiction stories and scenarios warning that abuse of smart machines can be hazardous to human health. No known chatbot currently obeys Asimov's [Three Laws of Robotics](#).

The more serious and practical takeaway, however, is that this irresistible rise of ever-smarter machines worldwide demands that leaders and managers become better people. Self-indulgently lashing out at underperforming bots and software agents represents poor leadership-by-example — and worse, bad manners.

Sooner, rather than later, organizations will have to recognize that getting the best out of their smart technologies requires getting the best from themselves.

Michael Schrage, a research fellow at MIT Sloan School's Center for Digital Business, is the author of the books [Serious Play](#) (HBR Press), [Who Do You Want Your Customers to Become?](#) (HBR Press) and [The Innovator's Hypothesis](#) (MIT Press).

Leadership and Big Data Innovation

PRESENTER:

Michael Schrage, Fellow, MIT Sloan School's Initiative on the Digital Economy

MODERATOR:

Angelia Herrin, Editor, Special Projects and Research, *Harvard Business Review*

Overview

As devices, sensors and new technology platforms generate ever more big data, organizations have new opportunities to gain important insights that can spur innovation, bring new products and services to the market and create a competitive edge. But unless organizations begin to treat data as a key asset that is just as important as their financial assets, those opportunities will be missed.

To survive and thrive in the big data era, organizations must implement new governance models, create dashboards of information shared across the company, encourage an environment of “experiment and scale,” and cultivate a culture that incorporates data into the everyday activities of the business. And the message needs to come from the top: the kind of data-driven transformation that most companies need will demand commitment from the C-suite.

Context

Michael Schrage discussed how leaders can promote innovation through big data and analytics.

Key Takeaways

To create new value, organizations must treat data as an asset.

In many cases, organizations use analytics to optimize existing sources of value, rather than leveraging data and analytics to innovate and create new value. To create new value, organizations must treat data as an asset, just like capital and human resources.

In just the next 12 to 18 months, organizations can expect to gather 10 to 1,000 times more data. As data assets grow, opportunities exist to create new value through more efficient processes, new value propositions, enhanced customer experiences, and more. Organizations must ask themselves how they would sell or market to customers and prospects differently, if they had 100 times more data about them.

“Organizations should expect 10 to 1,000 times more data in the next 12 to 18 months. To make data more valuable, organizations must consider how to define, measure, and assess value creation inside the enterprise and outside.” —Michael Schrage

Purposeful analytics like KPIs and dashboards support innovation and value creation.

Purposeful data and analytics can be applied to a variety of business activities, such as customer retention, upselling and bundling, and identifying new products and services. Both KPIs and dashboards support purposeful analytics.

To harness data to outcomes that matter, enterprises must align application programming interfaces (APIs) with KPIs. It is also important to think about dashboards related to innovation and new value creation. Key questions include what these dashboards would look like and which KPIs they would include.

From an organizational perspective, siloed organizations lead to siloed KPIs. Schrage suggests viewing data as a “solvent” that breaks down siloes and creates virtuous cycles within organizations. Insights gathered from analytics related to customer retention, for example, may inform analytics related to innovation.

Big data requires “little experiments.”

When an insight is identified in data, the next step is to determine which tests will validate whether the insight is true or not. More and better experimentation is needed. This means socializing experiments across the enterprise. New data should facilitate collaboration around experiments focused on new value creation.

“Organizations need a fundamental paradigm shift. Innovation will no longer be based on an R&D pipeline, but on an ‘experiment and scale’ model. Experiments are the future of business innovation.”
—Michael Schrage

Data leadership demands data governance.

Data oversight is the next leadership challenge for organizations. Explicit data governance is needed to determine who has access to data, who can experiment with it, and more. Observations related to data governance include:

- *Data stewards are a new category of colleague.* The data steward role involves planning, implementing, and managing the sourcing, use, and maintenance of data assets in an organization.

Data stewards enable organizations to take control and govern all the types and forms of data and their associated libraries or repositories.

- *Data governance plays a major role in company growth.* Serious growth and value companies have C-suites that are committed to data as an asset and to data management.
- *Data governance working groups must be unified around the idea that data resources can be transformed into new value.* Data governance working groups create the vision for the organization. It is important to have members who are committed to collaboration, data sharing, customer data protection, and developing data as an asset. Schrage recommends against having CFOs run data governance working groups, as they may overweight capital as an asset.



FIGURE 1: ELEMENTS OF DATA GOVERNANCE

As data volumes increase, the importance of an enterprise data culture grows.

To promote an enterprise “value of data” culture, organizations often need to make behavioral, cultural, and operational improvements. Examples of these shifts include:

- *API development.* To generate value, people expect to share data sets across the organization. APIs are essential to accomplish this.
- *Testing “gut feel” ideas with data-based experiments.* Many senior managers make decisions based on “gut feel,” rather than using experiments to challenge their perceptions. Ideas based on intuition aren’t bad, but they must be accompanied by testable business hypotheses. As Schrage noted, “Many good ideas evaporate on contact with the real world.” Data helps extract signals from the noise.
- *Deploying incentives to promote information sharing.* If data is equated with power, some people are inclined to hoard it. To promote information sharing at a large telecom, Schrage created an internal competition called “we were robbed” and “thief of the week.” These programs rewarded people for “stealing” data from other parts of the organization to create new value.
- *Revising performance review criteria.* Employees may be evaluated on how well they share data, as well as to what degree their decision making is driven by richer sources of data over time.

Other Important Points

Self-quantification. Schrage believes that individuals who engage in self-improvement every day will have an edge in the job market. People need to be more introspective about ways they add value in the workplace and must use data as a mirror to see themselves differently. For example, at the end of the day employees could analyze the content of their emails and texts using a Wordle. Being data driven and self-aware are two sides of the same coin.

Data and risk appetite. Organizations may use data to rethink their risk appetite. Insights derived from data are one way to create “guardrails” that can increase comfort with risk. Major “bet the company” decisions should always be data driven.

Data scientists. Some organizations place data scientists in a center of excellence, while others implant them in different groups. Regardless of where data scientists “live” in the organization, they should be expected to collaborate with the business and serve as a resource, like a good financial expert. Data scientists help teams achieve their goals by drawing on data as an asset.

FROM OUR SPONSOR

SAS continues to foster innovation through predictive analytics across industries. Leading organizations are bringing about the next analytics age by working with SAS for real-time streaming analytics, in-memory analytics, self-service analytics for business users, and applications for artificial intelligence that solve today's business problems and propel their organizations forward. Frequently sought out for quality, scalability, and compatibility of its solutions with open-source platforms, turn to SAS to make your new possible happen with analytics.

Start the journey to your new possible with these suggested next steps:

TDWI CHECKLIST REPORT: MARKETING ANALYTICS MEETS ARTIFICIAL INTELLIGENCE →

Artificial Intelligence and its subcomponents (machine learning, cognitive computing, and even deep learning) have grown in interest as data size and diversity continue to grow and the cloud becomes a popular option for quickly and economically scaling compute power and data storage. This report explores how Artificial Intelligence can be used to enhance applied analytics to help organizations both better understand their customers and deliver a great customer experience. It also provides practical advice on how organizations can use what they may already be doing to become even more effective in marketing and other business processes.

MANAGING THE ANALYTICAL LIFE CYCLE FOR DECISIONS AT SCALE →

This paper is a practical primer for the organizational leader contemplating the evolution of data-driven decision making to the automated decisions that define artificial intelligence. It addresses the major considerations that range from data preparation, through experimentation and model development to implementation and monitoring. Learn the factors that drive success in an analytically-driven operating environment that leads to the best automated models that realize the value creation possibilities of artificial intelligence.

SPONSORED BY 