Northwestern Kellogg School of Management

KE1046

January 17, 2018

ADAM ROBERT PAH, ALANNA LAZAROWICH, AND CHARLOTTE SNYDER

Evaluating the Cognitive Analytics Frontier

In the fall of 2014, Chad Kartchner '10, senior manager of marketing and product management at Honeywell Aerospace (HA), pondered how technology could transform the way aircraft were maintained. He wondered if he could develop an aircraft-wide system that could help Honeywell customers' crews manage their maintenance and repair processes more efficiently. Aircraft maintenance carried high stakes; malfunctioning parts and grounded planes could generate millions of dollars in replacement costs and delays.

Kartchner had been hearing a lot of buzz about cognitive computing, an artificial intelligence term referring to the use of computer models and algorithms to simulate human thought through self-learning systems, data mining, pattern recognition, and natural language processing.[1] Organizations across all industries seemed to be increasingly relying on sophisticated computer systems that could filter vast amounts of data and distill information that people could access quickly to make high-level business decisions, learn, and anticipate issues more effectively in complex situations. While attending a conference, Kartchner had listened to a keynote address demonstrating how cognitive analytics were revolutionizing cancer diagnostics. If oncologists could use emerging technologies to improve quality of care and patient experience in battles against fatal diseases, he mused, then perhaps HA could tap these resources to build a better decision support structure for aviation technicians that would ultimately increase air safety.

Kartchner's goal was to improve processes for capturing, analyzing, predicting, and transmitting data related to equipment maintenance, failure, and repair from myriad locations in a timely manner so that customers' technicians on the ground had the information they needed to perform their jobs effectively. He wondered if HA could build a connected platform on its own to enable customers to search a knowledge database or if it would be useful to explore how cognitive analytics providers could help.[2]

Honeywell Aerospace

Headquartered in Phoenix, Arizona, Honeywell Aerospace was a leading global manufacturer of aviation products. The company's product and service portfolio included aircraft engines, electric systems, thermal control systems, cockpit systems and displays, navigation, communications, cabin management and entertainment systems, flight safety, radar and surveillance systems, wheels and braking systems, as well as maintenance and service plans.[3] HA also manufactured turbochargers and automotive software solutions that improved the performance, efficiency, and security of passenger cars and commercial vehicles.

One of the largest divisions of Honeywell International, HA had more than 40,000 employees around the world and accounted for 39 percent of the company's overall sales. HA earned more than $15 billion in revenue in 2014, with sales to the U.S. government accounting for more than $3.6 billion.[4] Its customers included aircraft manufacturers, airlines, airports, operators, militaries, and defense and space contractors. HA's products were found on most commercial, defense, and space aircraft worldwide.

Aircraft Maintenance and Repairs

HA offered on-board monitoring and maintenance systems. As part of its maintenance, repair, and warranty programs, HA also offered a ground-based application interface that customers' technicians could use to perform activities such as searching manuals for troubleshooting advice, reading technical publications, ordering parts, and tracking order status. Customers could submit online technical requests and receive a call back from HA, or they could call a customer service line directly to make inquiries. Kartchner knew that the current ground-based process for aircraft troubleshooting had considerable room for improvement.

Typically, when an airplane was grounded, technicians were under intense time pressure, and for safety reasons, some operators required all fault codes (messages that correspond to specific failures in software or hardware) to be cleared before a plane could take off again. Fault codes were stored in an on-board maintenance computer and were generally transmitted automatically or called in by a pilot when they appeared. If a fault code had not been called in, a technician performing checks would often take a picture of the code with a cell phone, head back to his or her office, log in to a computer, and begin searching for fault codes in various maintenance databases. Problems with an airplane often triggered multiple fault codes that came in groupings, so the first challenge was to determine which ones really mattered to prioritize debugging.

Sheer volume of parts and the time-sensitive nature of repairs in the aviation industry made identifying problems and addressing them quickly exceptionally complicated. HA had tens of millions of hardware and software products on the market that it provided directly or that it sold to other manufacturers.[5] While most repair solutions could be found in various fault isolation manuals, maintenance manuals, repair manuals, and service bulletins, it was difficult for technicians to locate the appropriate sections of the appropriate manuals in a timely manner—let alone cross-reference these manuals and fault codes with the previous maintenance and repair history of an individual aircraft. For example, an Airbus 380 aircraft had approximately 4 million parts, with 2.5 million part numbers, produced by 1,500 companies in thirty countries.[6] Knowledge was fragmented in disconnected systems.

Once appropriate solutions were found, the technician printed troubleshooting guides and returned to the plane with a stack of paper to make the repairs. In addition, HA knew from the experience of its customer service department that as the complexity of a problem increased, technicians were less likely to seek answers from formal information tools and wanted the confidence of working directly with an expert.

Given the complexity of diagnostics and the time constraints of the technicians, there was a tendency in the industry to use an abbreviated process known as "shotgun maintenance." Technicians would often use the information from the plane's on-board maintenance computer to quickly formulate a plan that exchanged a number of parts—those that the technician thought were most likely to have caused a failure—in the hopes that the failures would be corrected. Failed parts were then sent back to manufacturers such as HA. When tested, many of these parts ultimately reported "no fault found"; they had been replaced unnecessarily because the lack of information, expertise, and time hampered solving complex repair problems.

In the near term, Kartchner hoped to help customers' crews quickly and accurately diagnose and fix problems when they occurred to reduce the incidence of "no fault found" exchanges and to decrease downtime for repairs. However, his larger goal was to give technicians an aircraft-wide maintenance solution that enabled them to be proactive rather than simply reactive, so they could anticipate problems and prevent them in the first place.

Internal Innovation

At the time, there were two competing schools of thought on how to approach diagnostics. A crowdsourced case-based approach, which entailed reviewing historical records to develop a knowledge base and using these records to diagnose current problems, was an evolution of the current service bulletin and repair manual standard. An emerging alternative approach was a model-based system that created a digital model of an aircraft, and diagnostics were developed, tested, and perfected in accordance with this digital model. Internally, HA leadership and engineers were solidly behind the model-based approach and its continued development.

As a leader in the field, HA took great pride in the development of its model-based diagnostic tool for on-board maintenance. This system integrated real-time sensor data to monitor the health of the aircraft during flight and offered both diagnostic and prognostic capabilities that enabled enhanced system reliability, faster repairs, and lower maintenance costs.[7] Using digital models of the aircraft, flight performance could be optimized; when fault codes appeared, the system could identify root causes and suggest appropriate repair processes quickly. However, this system was still under active development—thus far it was only able to model a contained system, such as the engine, and was not ready for public distribution to customers and partners.

Kartchner considered using internal resources to build a model-based ground system similar to HA's state-of-the-art on-board maintenance system. There was already significant support internally and with clients for this kind of approach, and buy-in seemed guaranteed.

In order to bootstrap an internally developed ground system, Kartchner considered supplementing the model-based approach currently under development at HA with a crowdsourced case-based approach, which would include crowdsourcing features that leveraged the abundance of knowledge among customers' technicians. However, Kartchner worried that this idea might run counter to the internal support behind model-based approaches. He also recognized that the plan involved overcoming a "chicken or the egg" strategy problem (with technicians being the chickens and information solutions being the eggs) because the platform would only become useful as technicians engaged with it. However, he was certain that a crowdsourcing feature would offer a virtuous circle and supply valuable information in real time to technicians. The challenge lay in incentivizing technicians, who were already pressed for time, to engage and provide the feedback that was so critical in generating network effects and ensuring the crowdsourced platform's success. Certainly it was possible to overcome these adoption problems through traditional methods—eBay, Uber, and Facebook were all prime examples of companies that had successfully launched two-sided marketplaces. However, Kartchner wondered if emerging cognitive technologies might help in reducing the ramp-up time for HA's new platform.

The New Frontier of Cognitive Analytics

The Internet of Things (IoT) had come to represent a new frontier for business across industries. IoT referred to devices that collected and transmitted data via the Internet. By 2008, there were more objects connected to the Internet than people.[8] Many of these devices were sensors in the so-called industrial Internet, which referred to increasingly connected machinery involved in infrastructure like oil and gas pipelines, wind turbines, and smart utility grids. These systems were often vast, expensive, and surprisingly vulnerable to breakdown. For example, a cracked ball bearing could ruin a $10 million turbine, or a multimillion-dollar centrifuge that ran just a few rotations too slowly or too quickly could burn out prematurely.[9] By 2014, there were an estimated 16 billion connected devices—everything from smart home devices and connected cars to wearable health trackers and bathroom soap dispenser sensors that provided alerts when supply was low—and those figures were projected to increase to 40 billion by 2020.[10]

IoT business opportunities were projected to add $10–15 trillion to global GDP in the next twenty years.[11] Organizations were racing to identify ways to harness the power of data, and emerging technology companies were helping them turn information into actionable insight.

Kartchner wanted HA's new solution to be mobile, digital, and cloud-based, so that ground crews anywhere in the world could access a knowledge database from their cellphones and get immediate results. He did not want the new system to be a push model that simply put out a giant body of knowledge in a variety of formats, manuals, and locations, leaving it to technicians to store, recall, and search for the information they needed. Instead, he wanted to develop a model-based system capable of incorporating crowdsourced solutions so that technicians anywhere in the world who identified useful ways of addressing technical issues could share their insights and best practices with anyone else who might need them. Kartchner began to evaluate two external companies that might provide additional support in this effort.

IBM Watson

IBM's supercomputer, Watson, stunned the world in 2011 when it beat the top two competitors in the history of the trivia game show *Jeopardy!*.[12] Using artificial intelligence, natural language processing, and machine learning, Watson was able to engage in dialogue with people and process information much like humans do, allowing the computer to quickly analyze and interpret large amounts of both structured and unstructured data.[13] Watson technology moved well beyond search engines such as Google that pointed searchers to documents that might contain an answer.[14] Instead, Watson was a question-answering machine that could understand questions posed in everyday language and respond with a precise answer.[15] Watson also learned over time, through feedback, to improve the relevance of its answers and its conversational skills.[16]

IBM believed that Watson's cognitive computing technology could enable people and computers to interact through natural language and augment human understanding of the world through big data insights.[17] John Kelly, head of IBM's research labs, hoped to bring Watson technology to every industry:

> *I want to create something that I can take into every other retail industry, in the transportation industry, you name it, the banking industry . . . Any place where time is critical and you need to get advanced state-of-the-art information to the front of decision-makers. Computers need to go from just being back-office calculating machines to improving the intelligence of people making decisions.*[18]

Following Watson's *Jeopardy!* win, IBM partnered with a number of healthcare organizations, including Memorial Sloan Kettering Cancer Center and the University of Texas MD Anderson Cancer Center, to develop systems that could build on Watson technology to generate personalized patient treatment plans.[19] In 2013, IBM introduced its Watson Engagement Advisor to enable businesses to improve connections with customers, and it later revealed its intention to create a cloud-based "Watson ecosystem" that enabled software application providers to build apps that incorporated Watson's cognitive computing intelligence.[20] In January 2014, IBM declared that it would invest $1 billion into its new Watson Group to encourage this development.[21] By July 2014, the Watson Group had more than 800 clients and partners in various stages of development,[22] and by October, around fourteen Watson ecosystem applications were in use.[23]

Companies that wished to become partners with the Watson Group could sign partner agreements and receive a toolkit, educational materials, and use of Watson's application programming interfaces for a specified period of time to design their own prototype.[24] IBM also provided some assistance along with recommendations for third-party vendors who could help interested partners train Watson for their specific business case and develop a market-ready product.[25] Once the prototype had been developed, IBM would either approve the prototype and begin to negotiate a commercial agreement, or it would opt not to move forward.[26] In a standard commercial agreement, IBM received a percentage of the commercialized product's revenue.[27] As part of the Watson ecosystem, partners benefited from IBM's marketing, as the technology giant touted use cases to drive further demand for Watson technology.[28] This resulted in incremental revenue for partners on their co-developed products.[29] However, before Kartchner could consider signing a partner agreement, he would need to work out non-disclosure and proprietary information agreements.[30]

Developing a customized user experience and user interface that incorporated Watson technology required significant amounts of time and resources on the part of partner companies (see **Figure 1**).

Figure 1: Sample Application Workflow When Using IBM Watson

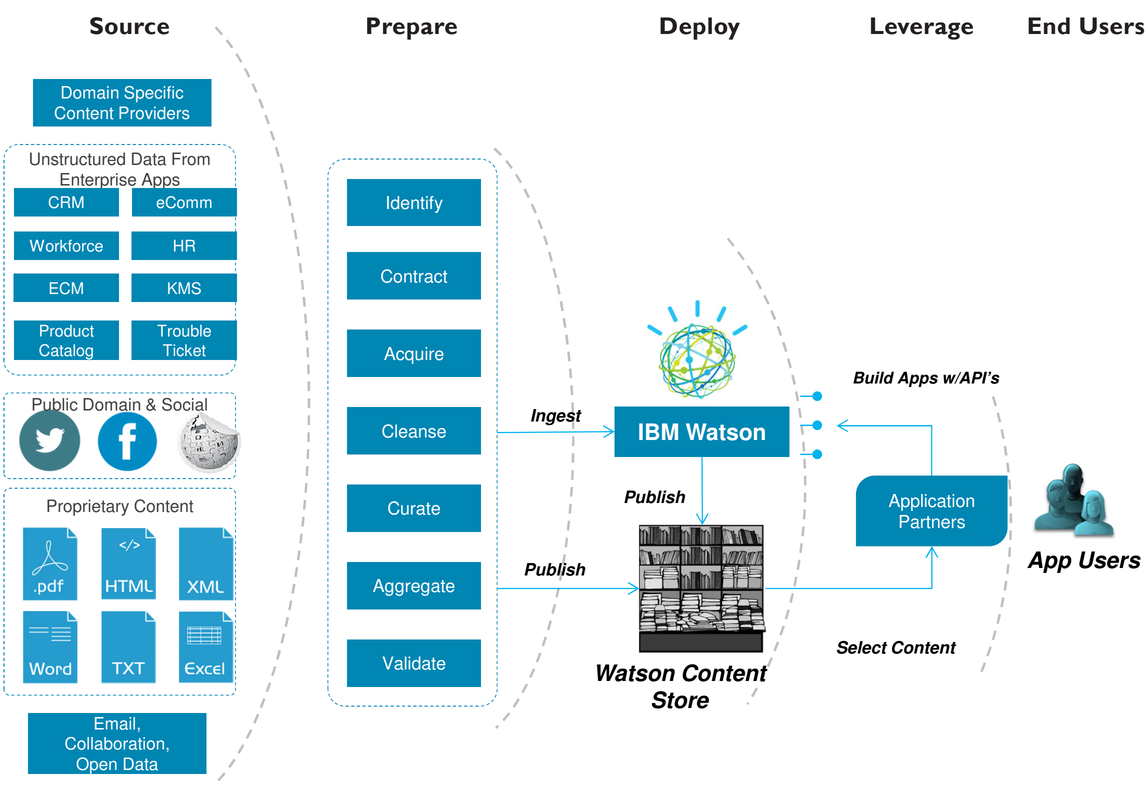

Source: Adapted from presentation by Sridhar Sudarsan (CTO, IBM Watson Ecosystem), "Getting Started with Watson: An Overview of the Watson Development Process," p. 51, IBM Watson Ecosystem Roadshow, Chicago, April 2, 2014, https://www.slideshare.net/SoniaBaratasAlves1/ibm-watson-56021913.

The first hurdle was training Watson to meet a partner's specific needs, which was referred to as extending Watson's capabilities.[31] Potential partners like HA had to pull together domain-specific content and data from their own proprietary information and documentation[32]—for example, more than 800 fault isolation manuals, more than 2,500 aviation maintenance manuals, and myriad other service bulletins, product catalogs, technical publications, and recall updates[33]—then upload and transform this information into a custom application.[34] In addition, partners needed to provide between 3,000 and 7,000 question-and-answer pairs[35] developed by subject matter experts to assist in the training, cleansing, curating, aggregating, and validating of information processes.[36] This process was known as corpus training.[37] Sourcing and preparing documentation and managing the actual ingestion process for Watson typically took two to three months.[38] Partners that had the engineering and development capabilities to manage the ingestion and corpus training processes could do it themselves; otherwise, partners paid IBM (or a third-party vendor) by the hour or the month to assist them.[39] Once Watson's capabilities had been extended, partner companies could begin building their user experience and user interface.[40]

Kartchner had less than six months to demonstrate a pilot version of any new solution in order to garner continued support for its development. To get a pilot version ready, he would need to first reach an initial agreement with IBM, schedule time for a team at HA to transform the repair manuals into a machine-readable format, schedule time for HA subject matter experts to develop thousands of question-and-answer pairs to support corpus training, and request time for the development team to build a user-facing application. Although HA had employees with the skills and knowledge necessary to do these tasks, Kartchner questioned his ability to requisition enough of their time for this new project given the employees' existing primary responsibilities. He also considered partnering with the third-party contractors IBM suggested to transform the repair manuals into a machine-readable format and build a user-facing application; however, the caveat with this approach was that it would introduce additional relationships that he would have to establish and manage in a short time period. Given Kartchner's past experience managing large technology-based products, he knew that increasing complexity, in both the number of steps and partners involved, made managing and delivering a project on time more difficult.

SparkCognition

Launched in April 2014 in Austin, Texas, SparkCognition was privately owned and backed by venture capital. It had strong ties with IBM: SparkCognition's president and CEO was appointed to the board of advisors for IBM Watson,[41] and its chairman of the board served as the general manager of IBM Watson between 2011 and 2014.[42]

SparkCognition had a handful of early clients but did not yet have a finished product or any client success stories. Its value proposition was to act as a full-service provider to assist client companies in extending Watson's capabilities. Specifically, SparkCognition handled the data ingestion and corpus training processes that integrated client documentation and proprietary information with data found in the public domain, and it could also create a user experience and design the user interface for its clients. By taking on corpus training and ingestion, SparkCognition claimed to be able to reduce the staff resources and the time that its clients would have to invest in these processes from two or three months to two or three weeks. Not only did SparkCognition claim to reduce training time, but its document preparation and data ingestion process also led to higher than average levels of search accuracy on an untrained corpus.

SparkCognition managed the IBM partnership for its clients, eliminating their need to sign a developer contract with IBM directly.[43] Lastly, by engaging SparkCognition for professional services, Honeywell would retain complete control over the intellectual property and revenue streams of the end product.

Although Kartchner was impressed with SparkCognition's technology, interest in the development of a new ground-based aircraft-wide maintenance system reached the highest levels of HA's senior management, and he was wary about the prospect of relying on a startup with a limited track record.

The Path Forward

Kartchner pondered his options. One thing he knew for certain was that these emerging technologies had the power to overhaul existing support systems for aviation technicians, and the winner in the new era of digital maintenance would have a strategic advantage in other data-driven business opportunities of the future, including predictive maintenance.

HA could explore building a new, connected, cloud-based platform on its own that employed crowdsourced features or pursue partnerships with IBM Watson or other companies like SparkCognition that could assist in increasing the company's cognitive capabilities. Kartchner recognized that his goal was larger than creating better search engines. He wanted ground crews to be able to ask questions and receive immediate, evidence-based recommendations to assist with their decision-making in real time. Ultimately, he hoped technicians on the ground would be able to proactively monitor assets to optimize systems, predict failures, and avoid downtime, just as their counterparts could in the air with HA's on-board maintenance systems. He just wasn't sure if the new ground-based solution would benefit significantly from leveraging emerging cognitive computing technologies or if HA was better off building the new system itself.

Endnotes

1 Margaret Rouse, "Cognitive Computing," WhatIs.com, accessed May 15, 2017, http://whatis.techtarget.com/definition/cognitive-computing.

2 Chad Kartchner (Honeywell Aerospace), in interview with the authors, April 28, 2017.

3 Honeywell Aerospace, https://aerospace.honeywell.com, accessed May 15, 2017.

4 Honeywell International Annual Report 2014.

5 Honeywell International, "Our Company," accessed May 15, 2017, https://www.honeywell.com/who-we-are/overview/our-company.

6 James Wallace, "A380 Ride & Factoids," SeattlePI.com, *Boeing and Aerospace News* (blog), March 29, 2007, http://blog.seattlepi.com/aerospace/2007/03/29/a380-ride-and-factoids.

7 Honeywell Aerospace, "Onboard Maintenance Systems," accessed May 15, 2017, https://aerospace.honeywell.com/en/products/cockpit-systems/onboard-maintenance-systems.

8 Dave Evans, "The Internet of Things [Infographic]," Cisco.com, *Inclusion and Diversity* (blog), July 15, 2011, https://blogs.cisco.com/diversity/the-internet-of-things-infographic.

9 Roger Kay, "SparkCognition Gains Traction," *Medium*, November 18, 2015, https://medium.com/@rogerkay/sparkcognition-gains-traction-fe66223368fd.

10 Gil Press, "Internet of Things by the Numbers: Market Estimates and Forecasts," *Forbes*, August 22, 2014, https://www.forbes.com/sites/gilpress/2014/08/22/internet-of-things-by-the-numbers-market-estimates-and-forecasts/#58e3b8ccb919.

11 Ibid.

12 John Markoff, "Computer Wins on 'Jeopardy!': Trivial, It's Not," *New York Times*, February 17, 2011, http://www.nytimes.com/2011/02/17/science/17jeopardy-watson.html.

13 Jo Best, "IBM Watson: The Inside Story of How the Jeopardy-Winning Supercomputer Was Born, and What It Wants To Do Next," *TechRepublic*, September 9, 2013, https://www.techrepublic.com/article/ibm-watson-the-inside-story-of-how-the-jeopardy-winning-supercomputer-was-born-and-what-it-wants-to-do-next.

14 Clive Thompson, "What Is I.B.M.'s Watson?" *New York Times*, June 16, 2010, http://www.nytimes.com/2010/06/20/magazine/20Computer-t.html.

15 Ibid.

16 Michael Goldberg, "Five Things to Know About IBM Watson, Where It Is and Where It's Going," *DataInformed*, January 14, 2014, http://data-informed.com/five-things-know-ibm-watson-going.

17 Thor Olavsrud, "In Pictures: 10 IBM Watson-Powered Apps That Are Changing Our World," *Computerworld Australia*, November 7, 2014, http://www.computerworld.com.au/slideshow/559079/pictures-10-ibm-watson-powered-apps-changing-our-world.

18 Thompson, "What Is I.B.M.'s Watson?"

19 "IBM Forms New Watson Group to Meet Growing Demand for Cognitive Innovations," press release, January 9, 2014, https://www-03.ibm.com/press/us/en/pressrelease/42867.wss.

20 Ibid.

21 "IBM to Invest $1 Billion to Create New Business Unit for Watson," Reuters, January 9, 2014, http://www.reuters.com/article/us-ibm-watson-idUSBREA0808U20140109.

22 Richard Byrne Reilly, "Inside IBM's Billion-Dollar Bet on Watson," *VentureBeat*, July 18, 2014, https://venturebeat.com/2014/07/18/inside-ibms-billion-dollar-bet-on-watson.

23 "IBM Unveils New Start Ups and Ecosystem Partners Delivering Watson-Powered Apps in the Cloud," press release, October 7, 2014, https://www-03.ibm.com/press/us/en/pressrelease/45023.wss.

24 "IBM Watson's Next Venture: Fueling New Era of Cognitive Apps Built in the Cloud by Developers," press release, November 14, 2013, https://www-03.ibm.com/press/us/en/pressrelease/42451.wss.

25 Presentation by Lauri Saft (director, Partner Recruitment, IBM Watson Ecosystem), "Welcome to Watson Innovation Day," IBM Watson Ecosystem Roadshow, Chicago, April 2, 2014, https://www.slideshare.net/SoniaBaratasAlves1/ibm-watson-56021913.

26 Fredrik Stenbeck, "Watson Ecosystem, How It Works," *Fredrik Stenbeck* (blog), February 9, 2016, http://fredrikstenbeck.com/watson-ecosystem-how-it-works, and Paul A. S. Ward (professor, software engineering, University of Waterloo), comment on April 25, 2014, https://www.quora.com/What-is-the-cost-of-licensing-hardware-to-get-an-IBM-Watson-starter-system-purchased-for-development.

27 Alpha Max (*pseud*.), "IBM's Next Big Thing—Part 2," *Seeking Alpha* (blog), May 11, 2016, https://seekingalpha.com/article/3974160-ibms-next-big-thing-part-2.

28 "IBM Watson Ecosystem Partners in Market Building Businesses," press release, September 24, 2015, http://www-03.ibm.com/press/us/en/pressrelease/47717.wss, and Saft, "Welcome to Watson Innovation Day."

29 Alpha Max, "IBM's Next Big Thing—Part 2."

30 University of Toronto, "Attachment to Agreement for IBM Watson Academic Initiative: IBM Watson Prototype Development Agreement," January 15, 2015, http://www.cs.utoronto.ca/~sengels/csc490/docs/WatsonProjectAgreement.pdf, and IBM developerWorks, "IBM Confidential Disclosure Agreement," accessed November 14, 2017, https://www.ibm.com/developerworks/community/terms/ibmconfidentiality/index.html.

31 Usman Shuja, "The Era of Exponential Impact," IBM, Client Success Field Notes, October 17, 2017, https://www.ibm.com/blogs/client-voices/era-exponential-impact; Ryan Boyles, "Bringing Cognitive IoT to the Edge," Storify, June 2, 2016, https://storify.com/therab/bringing-cognitive-to-the-edge; and "Hot Jobs at IBM," IBM Research, 2011: Volume 2, Hot jobs, accessed November 8, 2017, https://www-03.ibm.com/employment/news/201105_volume2_hotjobs_research.html.

32 Dale Lane, "Cognitive Computing: Putting IBM Watson to Work," https://www-03.ibm.com/systems/data/flash/uk/resources/infrastructure_matters/IBM%20Infrastructure%20Matters%20-%20Cognitive%20computing%20-%20putting%20Watson%20to%20work%20-%20Dale%20Lane.pdf, and Nick Heath, "Five Things You Need To Know Before Building an IBM Watson App," *TechRepublic*, October 27, 2014, https://www.techrepublic.com/article/five-things-you-need-to-know-before-building-an-ibm-watson-app.

33 Usman Shuja (SparkCognition), in interview with the authors, May 5, 2017.

34 Heath, "Five Things You Need To Know Before Building an IBM Watson App."

35 "IBM Watson Ecosystem: Getting Started Guide," Version 1.1, July 2014, p. 10, http://cognitive-science.info/wp-content/uploads/2014/10/IBM-Watson-University-Getting-Started-Guide-v1.7.pdf.

36 Heath, "Five Things You Need To Know Before Building an IBM Watson App."

37 "IBM Watson Ecosystem: Getting Started Guide," p. 5, and Heath, "Five Things You Need To Know Before Building an IBM Watson App."

38 Heath, "Five Things You Need To Know Before Building an IBM Watson App."

39 Saft, "Welcome to Watson Innovation Day," and Heath, "Five Things You Need To Know Before Building an IBM Watson App."

40 Presentation by Jodie Sasse (program director, Engagements, IBM Watson Ecosystem), "Getting Ready for Watson," IBM Watson Ecosystem Roadshow, Chicago, April 2, 2014, p. 32, https://www.slideshare.net/SoniaBaratasAlves1/ibm-watson-56021913.

41 Syed Talal, "IBM Appoints Amir Hussain as Board of Advisors for Watson," *ProPakistani*, October 31, 2014, https://propakistani.pk/2014/10/31/ibm-appoints-amir-hussain-board-advisors-watson, and SparkCognition, "Leadership Team," accessed November 14, 2017, https://sparkcognition.com/leadership-team.

42 CognitiveScale, "CognitiveScale at SXSW 2017," accessed November 14, 2017, https://www.cognitivescale.com/event-post/cognitivescale-sxsw-2017, and SparkCognition, "Leadership Team."

43 Usman Shuja (SparkCognition), in interview with the authors, May 5, 2017.