

## The Case for Using Robots in Intelligence Analysis

Puong Fei Yeh

---

### ***Is it possible that automation could determine the demand for intelligence analysts?***

From music to cars, the trend of automation—machines substituting human labor—is quickening. An article published in the *MIT Technology Review* in 2012 posed the question, “Can creativity be automated?” The author then told the story of Music Xray<sup>a</sup> which possesses technology to detect potential hit songs by using algorithms that compare a song’s constituent parts—tempo, rhythm, pitch—against those of historical top hits.<sup>1</sup> In March 2015, the e-commerce giant, Amazon, obtained US regulatory approval to begin testing its drone delivery service.<sup>2</sup> At the same time, Tesla Motors announced that it would roll-out autonomous driving technology.<sup>3</sup> Microsoft cofounder Bill Gates once commented that software substitution will ultimately affect workers ranging from drivers to waiters: “Twenty years from now, labor demand for lots of skill sets will be substantially lower, and I don’t think people have that in their mental model.”<sup>4</sup> Is it possible that automation could determine the demand for intelligence analysts?

---

#### ***Automated Analysis: It’s everywhere***

In a story heard all too often these days, my credit card company recent-

---

a. The company describes itself as a “song and talent filter,” linking artists and composers with potential employers.

ly stopped my card from being used for purchases in Kentucky that I had not made. For this action to occur, the company (or rather, its software) probably spotted purchases made several hundred miles apart and within minutes of each other. It drew the conclusion that I could not have been in two places at once, buying a TV at a WalMart in Kentucky and a coffee at Dunkin Donuts in Boston. Or maybe the software spotted a series of low value transactions that preceded the attempt to buy the TV.

Companies on Wall Street automate stock trades based on real-time analysis of news. In a span of two minutes during the early afternoon of 23 April 2013, the stock market fell 100 points before it recovered. The cause of the plunge was a fake tweet from a hacked Associated Press (AP) Twitter account claiming the White House had been attacked. The event underscored the extent to which companies analyze news in real-time to exploit market-moving news.<sup>5</sup> Dataminr Inc., a leading company in providing real-time social media analytics to financial institutions, analyzes hundreds of millions of tweets a day to uncover market-moving news. Last September when a journalist tweeted that Home Depot was the victim of a cyber attack, Dataminr sent out an alert to its subscribers nearly 15 minutes ahead of news wires and before Home Depot’s share price declined 2 percent.<sup>6</sup>

---

The views, opinions, and findings should not be construed as asserting or implying US government endorsement of its factual statements and interpretations or representing the official positions of any component of the United States government.

---

---

***The Open Source Indicators project aims to “develop methods for continuous, automated analysis of publicly available data in order to anticipate and/or detect significant societal events...”***

---

The Intelligence Advanced Research Projects Activity, the Intelligence Community’s (IC’s) blue-skies research arm, is sponsoring the Open Source Indicators (OSI) project.<sup>7</sup> OSI aims to “develop methods for continuous, automated analysis of publicly available data in order to anticipate and/or detect significant societal events such as political crises, humanitarian crises, mass violence, riots, mass migrations, disease outbreaks, economic instability, resource shortages, and responses to natural disasters.”<sup>8</sup> In early 2014, OSI claimed to have detected disease outbreaks in Latin America up to two weeks before they were reported by local media or official government authorities.<sup>9</sup>

---

***The Rise of Machine Reporters***

On 17 March 2014, the *Los Angeles Times* ran a nearly 150-word story that gained attention not for its content—a report of an earthquake aftershock—but for how it was written.

*“Earthquake aftershock: 2.7 quake strikes near Westwood*

March 17, 2014, By Ken Schwencke

*A shallow magnitude 2.7 earthquake aftershock was reported Monday morning four miles from Westwood, according to the U.S. Geological Survey. The temblor occurred at 7:23 a.m. Pacific time at a depth of 4.3 miles.*

*A magnitude 4.4 earthquake was reported at 6.25 a.m. and was felt over a large swath of Southern California.*

*According to the USGS, the epicenter of the aftershock was five miles from Beverly Hills, six miles from Santa Monica and six miles from West Hollywood.*

*In the last 10 days, there has been one earthquake of magnitude 3.0 or greater centered nearby.*

*This information comes from the USGS Earthquake Notification Service and this post was created by an algorithm written by the author. [Emphasis added]*<sup>10</sup>

Although the *Los Angeles Times* printed Schwencke’s name in the byline, the actual author was the program he designed, Quakebot.

From writing corporate earnings reports to sports stories, natural-language generation—the process of machines generating narratives based on data—is taking hold.<sup>11</sup> Participants in a 2014 study could not tell whether a recap of a football game was written by a human or machine.<sup>12</sup> Yahoo’s partnership with North Carolina-based Automated Insight enables millions of team owners in Yahoo Fantasy Football to receive automated status updates on their team’s performance, including automated trash talk.<sup>13,14</sup>

In 2014, AP, partnering with Automated Insights, automated the writing of corporate earnings reports. In March, the AP announced plans to expand the use of the automation technology to cover college sports.<sup>15</sup> Commenting on the technology’s impact, the AP noted automation allowed the production of 10 times as many stories per quarter as its reporters could previously accomplish.<sup>16,17</sup>

The IC’s venture capital firm, In-Q-Tel, has taken notice. In 2013, In-Q-Tel and Chicago-based Narrative Sciences announced a strategic partnership and technology development agreement to create a version of Narrative Science’s automated analysis software, *Quill*, for the IC.<sup>18</sup> Narrative Science began its work with automated sports reports and is now generating earnings previews for *Forbes*, mutual fund performance reports for financial institutions, and investment research reports.<sup>19</sup> The first paragraphs of a March 2015 report for *Forbes.com* illustrates the power of *Quill*.

*Wall Street is optimistic about Barnes & Noble, which is slated to report its third quarter results on Tuesday, March 10, 2015. Analysts project a profit of \$1.19 a share, a rise from 86 cents per share a year ago.*

*The consensus estimate remains unchanged over the past month, but it has increased from three months ago when it was \$1.10. Analysts are projecting earnings of 28 cents per share for the fiscal year. Analysts look for revenue to decrease 5% year-over-year to \$1.90 billion for the quarter, after being \$2 billion a year ago. For the year,*

revenue is projected to come in at \$6.05 billion.

Written narrative is important because it is the primary way information and insights are communicated to most people. As one writer noted, “just as words cannot really turn into pictures, pictures cannot replace words in terms of their ability to convey clear, (mostly) unambiguous information.”<sup>20</sup>

That the adoption of automated journalism is occurring during the Big Data era is not a coincidence. A 2012 study sponsored by information technology provider EMC and market research and analysis firm IDC concluded that 23 percent of the information in the “digital universe” was potentially useful if tagged and analyzed, but right now only 3 percent of that digital data is tagged and even less (under 1 percent) is analyzed.<sup>21</sup>

Today’s computing technology and advances in information storage mean volumes of past and real-time data being created at high velocity are now exploitable.<sup>22</sup> With increased amounts of data for analysis, there is increased demand for people and machines to make sense of *and* act on that data. The “global economy is beginning to operate truly in real-time, with constant streams of data showing where consumers are shopping, ships are traveling, energy and money are flowing,” according to the McKinsey Global Institute.<sup>23</sup>

Where companies like Automated Insights and Narrative Science fit into the Big Data landscape is that their technology, which is scalable, can analyze complex and high volumes of data and autonomously translate those insights into plain English.

---

***Today’s computing technology and advances in information storage mean volumes of past and real-time data being created at high velocity are now exploitable.***

---

According to Narrative Science CEO Stuart Frankel, “imagine the amount of time and money companies are spending taking data, trying to get a couple of interesting things out of that data, and then putting it into PowerPoint or Word. It takes an extraordinary amount of effort.”<sup>24</sup>

---

***Robots in Intelligence***

In spite of the numbers of agencies (17) that comprise the IC, analytical writing across the community follows a general formula and style. Some even view intelligence writing as somewhat mechanistic or borderline robotic. Many analysts use templates, which have specific guidance on how to structure an intelligence story according to the topic. Whether you are a puzzle solver—those who try to connect the proverbial dots—or a mystery framer—those who focus on “political and societal questions related to people, such as regional issues, national intent, or group intentions and plans”—the written narrative remains the principal means of delivering analysis to their audiences.<sup>25</sup>

Typically within the first paragraph, analysts describe a new development and give their bottom-line take on the development. The remaining paragraphs address reasons for the development and the last paragraph or paragraphs provide implications and outlook or suggest opportunities to capitalize or manage the risk of the development. Often below each paragraph are bulleted

points that provide “evidence” or supporting analytical reasoning.

The deliverable or the intelligence product can take many forms, ranging from PowerPoint slides to reports running several pages in length. Topics covered range from political and economic issues to terrorism and weapons development. The underlying narrative formula, however, for these stories is the same. Start with a development—nefarious guy in country X is preparing to do bad things in country Y—and explain to the intended audience the cause and significance of this development.

Story templates and data—both of which are in abundance in the IC—are essential for natural-language generation. To date, the technology’s rise to prominence is its capacity to autonomously analyze and generate stories based on quantitatively-rich data. Statistical data that can be compared across time and against other like data is particularly well-suited for natural-language generation. This is why the technology’s most well-known applications center on sports and financial-related topics.

To summarize a football game, for instance, one would want to know the key play or set of plays that led to the outcome. Was it time of possession or yards or total offensive yards gained that swung the game towards the final outcome? Using data from sports analytics platforms that provide continual updated odds of winning given a certain play, down and distance, and remaining time, natural-language generation can be used to automate this type of analysis.<sup>26</sup> The platform

---

**Stories that center on quantitative data (e.g., measurements) are probably most suitable for natural language generation.**

---

would then compare the game against a broader context. How unique was this outcome in the current season? How about compared to the past 10 seasons? The last step involves translating the analysis into the written narrative.<sup>27, 28</sup>

One also can tailor the narrative to suit different readers. Suppose you're a fan of the losing team, would you really want to read a recap of how the opposing team put up four touchdowns and 400 yards, breaking multiple records in the process? In fact, Big Ten conference officials approached NarrativeScience about writing stories that were less embarrassing for the losing team. When NarrativeScience covered little league games, its platform left out dropped fly balls to soften the blow to parents' egos.<sup>29</sup>

In the IC, the data that underpins analytical assessments are derived from across the "INTs" (i.e., human, signals, imagery, and measurement and signatures intelligence). The data can be qualitative or quantitative in nature. Qualitative data can come from the interpretation of person-to-person conversations or documentary information. Quantitative data can take the form of technical signatures or characteristics of a specific target. Imagery data can have both quantitative and qualitative aspects. Single or multiple sources of intelligence can be used to directly or indirectly support analytic reasoning. The evidence underlying analysis is messy and contradictory.

---

**Possibilities and Challenges in Automating Analysis**

What intelligence stories might natural-language generation write? Stories that center on quantitative data (e.g., measurements) are probably most suitable for natural-language generation. Measurement and signature intelligence gleaned from tests of an adversary's new ballistic missile, for example, would seem like a good fit. The data acquired during a test could include information on the composition of the exhaust or distinctive radar signatures and emissions or trajectories and distances from launch to impact. With such information from a current test together with data from previous tests or tests involving different missiles, comparisons can be made against current and historical data and patterns. One could structure a story on the significance of the most recent test or just focus on a discrete aspect of the test itself, for example, the composition of the missile exhaust.

The following is a hypothetical example of the possibilities of natural-language generation based on a declassified intelligence report about a 1980 Chinese intercontinental ballistic missile test.<sup>30</sup>

*Title: Missile Launch Preparations are underway involving Intercontinental Ballistic Missile*

*Bottom Line: Recent construction of tracking equipment facilities and activities at a missile production plant suggest launch preparations are underway.*

*Summary: In late November 1979, construction of two, new facilities to house launch alignment equipment occurred at the Shuangchengzi Surface-to-Surface Missile Complex. These facilities were larger than the old facilities, suggesting that the larger facilities house new types of alignment equipment. During the same month, activities observed at the Shanghai/Minhang Missile Production Plant were consistent with the transport of intercontinental ballistic missiles. The Shanghai/Minhang production plan is associated with CSL-2 space launch vehicles.*

The parts of the above story that are quantitative are the dimensions of the launch alignment facilities and signatures associated with the transport of missiles at a previously identified missile production plant. From this collection of data, a computer searches for insight asking analytical questions of the sort an analyst might ask.

- What missiles are associated with the missile production plant?
- With regard to physical dimensions, how is the new launch alignment complex different from the old complex?
- What is the significance of the missile transport activities?
- What is the significance of both occurrences—construction of new facilities and missile transport activities—occurring in the same time frame?

- Have these occurrences coincided before and if so what occurred afterwards?

Typical capabilities embedded in natural-language generation technology include regression analysis, time series modeling, and other statistical approaches. Translating these analytic insights into a narrative involves the computer understanding the style and structure of analytical products. For the hypothetical title, “Missile Launch Preparations are Underway Involving Intercontinental Ballistic Missile,” the computer needs to generate a headline that communicates the message of the report within a specified word limit. Similarly, for the bottom line, the computer needs to communicate the gist of the story while avoiding technical jargon that only specialists would understand. With automation, however, different narratives could be generated from the same data to suit different audiences. A policymaker’s analytical needs are different from those of a Chinese ballistic missile analyst.

A more challenging task for natural-language generation is creating stories based on qualitative information, although some qualitative data is more conducive to measurement than others. Measuring sentiment through analysis of social media probably is an easier task for natural-language generation than figuring out an adversary’s plans and intentions based on fragmentary reporting from multiple intelligence sources.

The United Nations-sponsored Global Pulse Initiative, in 2012, successfully predicted an increase in the price of foodstuffs in Indonesia several weeks ahead of government indices by measuring popular senti-

---

***A question separate from whether natural language generation can effectively process qualitative information is which can tell a better story, humans or computers?***

---

ment through Twitter.<sup>31</sup> Now imagine a natural-language generated story on the same topic that says, “based on the increasing intensity of popular discontent regarding foodstuff prices over the past three months, inflation probably has increased between X and Y percent. Compared with the past five years, the current estimated level of inflation represents a new high. The last time the inflation reached new levels, riots occurred in Jakarta and other major population centers. . . .”

Producing assessments on a government’s intent and plans are a greater challenge for natural-language generation since the underlying data is mostly unmeasurable. The intelligence reporting that analysts use to produce these types of assessments could span lengthy periods of time, consist of information from multiple sources across topics as diverse as domestic politics and defense industries. They also involve intuition.

For example, determining a country’s threshold for initiating a military offensive over a territorial dispute may require information on the capabilities of its military forces (current and projected), dynamics of senior political-military decision-making, and similar information on other internal and external actors. Assessing military force capability, alone, require judgments about forces’ readiness, structure, and sophistication of military equipment and personnel capabilities and will. Of these, sophistication of equipment may be most measurable. Now imag-

ine some of the information on the aforementioned factors is missing or fragmentary. In the financial services industry—a data rich environment for measurement—natural-language generation still has “trouble processing the qualitative information central to most analysts’ jobs.”<sup>32</sup>

A question separate from whether natural-language generation can effectively process qualitative information is which can tell a better story, humans or computers? In comparing two Disney corporate earnings reports—reports that natural-language generation currently process—*Slate* columnist Will Oremus noted, “[a] good story about Disney requires a journalist who already has a conception of what the company is about and why it’s important in the wider scheme of things. . . . [the] piece also draws on [Ryan Nakashima’s] understanding of the big abstract questions looming over 21<sup>st</sup>-century business management. Are big content brand acquisitions, in general, worth the money?”<sup>33</sup> For some kinds of storytelling, processing power, alone, may be insufficient, and therefore some automated intelligence stories do require a human touch.<sup>34</sup>

---

***Rebalancing Analytic Resources***

As is widely known, one of the big challenges confronting the IC is the ever increasing volume of data that requires timely analysis and the speed with which data is created and distributed. Adding another layer to these challenges is the IC’s requirement to maintain a watch over nearly

---

**Why not let robots help in analysis and writing and free up human analysts to engage in the exhaustive research and creative efforts necessary to confront the dangers we face now and will face in the future?**

---

the entire globe. US policymakers and the public expect the IC to be capable of providing warning in any part of the world where events may affect US national security and interests.<sup>35</sup> This often results in a situation in which a handful of countries (China, Iran, North Korea, and Russia) and topics (terrorism and proliferation of weapons of mass destruction) attract the most collection and analytical resources while remaining countries and topics attract much less.

Automation could help the community address the imbalance by permitting it to realign resources differently. For example, if analysis of developments like missile tests or changes of military force dispositions were automated and quickly communicated with little human involvement delivery of intelligence reports on these developments could be completed and disseminated much more quickly as time on analysis, drafting, editing, and reviewing would be reduced.

In the case of AP, automating news reports on corporate earnings

resulted in “far fewer errors” than appeared in manually written reports.<sup>36</sup> More importantly, automation allowed organizations to shift resources to their greatest need. For the AP, automation allowed it to “focus more reporters on higher-end enterprise stories *that break news that no one else has* [emphasis added].”

Ken Schwencke, the journalist and developer of *Los Angeles Times*’ Quakebot, said, “it [natural-language generation] saves people a lot of time, and for certain types of stories, it gets the information out there in usually about as good a way as anybody else would. The way I see it is it doesn’t eliminate anybody’s job as much as it makes everybody’s job more interesting.”<sup>37</sup>

Another potential benefit may be that IC producers of analysis may be able to more effectively tailor their products for different—and possibly a wider range of—consumers. Today, analysts tend to be focused on the needs—both in content and style—of specific audiences. Demands from multiple consumers compete for the time and attention of analysts and

their managers. Consequently, prioritization will lead to some customer’s needs being unmet, imperfectly met, or ignored. Natural-language article generation can help cover some of those gaps. During the football season, Yahoo Fantasy Football delivers personalized stories to millions of its users on their teams’ status. The goal of Automated Insights CEO Robbie Allen is to deliver tailored analysis on a massive scale— “[i]nstead of one story with a million page view[s], we’ll have a million stories with one page view each.”<sup>38</sup>

Earlier this year, Director of National Intelligence James Clapper testified before the Senate Armed Services Committee about the landscape of threats facing the United States.<sup>39</sup> Declaring that the new normal is “unpredictable instability,” he provided a grim picture of our world and the uncertainty in predicting how threats will unfold. As the same time, the IC faces increased demand for intelligence during a time of constrained budgets.

Why not let robots help in analysis and writing and free up human analysts to engage in the exhaustive research and creative efforts necessary to confront the dangers we face now and will face in the future?



---

**Endnotes**

1. Christopher Steiner, "Can Creativity be Automated?" *MIT Technology Review*, July 27, 2012, <http://www.technologyreview.com/news/428437/can-creativity-be-automated>.
2. Nick Wingfield, "Amazon Wins Approval To Test Delivery Drones Outdoors," *The New York Times*, March 19, 2015, <http://www.nytimes.com/2015/03/20/technology/amazon-wins-approval-to-test-delivery-drones-outdoors.html>.
3. Aaron M. Kessler, "Elon Musk Says Self-Driving Tesla Cars Will Be in the U.S. by Summer," *The New York Times*, March 19, 2015, [http://www.nytimes.com/2015/03/20/business/elon-musk-says-self-driving-tesla-cars-will-be-in-the-us-by-summer.html?\\_r=0](http://www.nytimes.com/2015/03/20/business/elon-musk-says-self-driving-tesla-cars-will-be-in-the-us-by-summer.html?_r=0).
4. Transcript of "From Poverty to Prosperity: A Conversation with Bill Gates," *American Enterprise Institute*, March 13, 2014, [http://www.aei.org/wp-content/uploads/2014/03/-bill-gates-event-transcript\\_082217994272.pdf](http://www.aei.org/wp-content/uploads/2014/03/-bill-gates-event-transcript_082217994272.pdf).
5. Dina ElBoghdady, "Market quavers after fake AP tweet says Obama was hurt in White House explosions," *The Washington Post*, April 23, 2013, ([http://www.washingtonpost.com/business/economy/market-quavers-after-fake-ap-tweet-says-obama-was-hurt-in-white-house-explosions/2013/04/23/d96d2dc6-ac4d-11e2-a8b9-2a63d75b5459\\_story.html](http://www.washingtonpost.com/business/economy/market-quavers-after-fake-ap-tweet-says-obama-was-hurt-in-white-house-explosions/2013/04/23/d96d2dc6-ac4d-11e2-a8b9-2a63d75b5459_story.html)).
6. Bradley Hope, "Startups Mine Market-Moving Data from Fields, Parking Lots—Even Shadows," *The Wall Street Journal*, November 20, 2014, <http://www.wsj.com/articles/startups-mine-market-moving-data-from-fields-parking-lotseven-shadows-1416502993>.
7. Sharon Weinberger, "Spies to use Twitter as crystal ball," *Nature*, October 17, 2011, <http://www.nature.com/news/2011/111017/full/478301a.html>.
8. <http://www.iarpa.gov/index.php/research-programs/osi>
9. Ray Locker "Data analysis allows researchers to predict disease outbreaks," *USA Today*, February 13, 2014, <http://www.usatoday.com/story/nation/2014/02/13/iarpa-anticipating-news-before-it-happens/5454511/Data%20analysis%20allows%20researchers%20to%20predict%20disease%20outbreaks>
10. Ken Schwencke, "Earthquake aftershock: 2.7 quake strikes near Westwood," *The Los Angeles Times*, March 17, 2014, <http://articles.latimes.com/2014/mar/17/news/earthquake-27-quake-strikes-near-westwood-california-rdivor>
11. Will Oremus, "Why Robot," *Slate*, February 5, 2015, [http://www.slate.com/articles/technology/future\\_tense/2015/02/automated\\_insights\\_ap\\_earnings\\_reports\\_robot\\_journalists\\_a\\_misnomer.single.html](http://www.slate.com/articles/technology/future_tense/2015/02/automated_insights_ap_earnings_reports_robot_journalists_a_misnomer.single.html)
12. Christer Clerwall (2014) Enter the Robot Journalist, *Journalism Practice*, 8:5, 519-531, DOI: 10.1080/17512786.2014.883116, <http://dx.doi.org/10.1080/17512786.2014.883116>.
13. Allen Weiner, "Fantasy football and the cold future of robot journalism," *The Kernel*, August 24, 2014, <http://kernelmag.dailydot.com/issue-sections/features-issue-sections/10097/fantasy-football-and-the-cold-future-of-robot-journalism/>.
14. Russell Brandom, "How to teach a robot to write," *The Verge*, July 11, 2014, <http://www.theverge.com/2014/7/11/5890693/how-to-teach-a-robot-to-write>.
15. Press Release, Automated Insights Blog, "Automation Helps AP Publish 10 Times More Earnings Stories", January 29, 2015, <http://blog.automatedinsights.com/post/109491692518/automation-helps-ap-publish-10-times-more-earnings>.
16. Automated Earnings Stories Multiply, AP, Erin Madigan White, January 29, 2015, <http://blog.ap.org/2015/01/29/automated-earnings-stories-multiply/>.
17. "AP, NCAA to grow college sports coverage stories," *The Associated Press*, March 4, 2015, <http://www.ap.org/content/press-release/2015/ap-ncaa-to-grow-college-sports-coverage-with-automated-game-stories>.
18. Press release announcing Narrative Science and In-Q-Tel strategic partnership and investment technology agreement, Narrative Science press release, June 5, 2013, <http://www.businesswire.com/news/home/20130605005246/en/Narrative-Science-Announces-Strategic-Investment-Technology-Development#.VPZ2iC60eFu>.
19. Tom Simonite, "Robot Journalist Finds New Work on Wall Street," *MIT Technology Review*, January 9, 2015, <http://www.technologyreview.com/news/533976/robot-journalist-finds-new-work-on-wall-street/>
20. Paolo Gaudiano, "As it turns out, a picture is not worth a thousand words," *Wired*, March 2014, <http://www.wired.com/2014/03/turns-picture-worth-thousand-words/>.
21. John Gantz and David Reinsel, "The Digital Universe in 2020: Big Data, Bigger Digital Shadow, and Biggest Growth in the Far East," IDC, December 2012, <http://www.emc.com/collateral/analyst-reports/idc-the-digital-universe-in-2020.pdf>.
22. Gartner IT glossary, <http://www.gartner.com/it-glossary/big-data/>.
23. McKinsey Global Institute, "Open Data: Unlocking innovation and performance with liquid information," McKinsey&Company, October 2013.
24. Tom Groenfeldt, "Lots of Data, One Analyst, Many Reports—Narrative Science," *Forbes*, September 5, 2013, <http://www.forbes.com/sites/tomgroenfeldt/2013/09/05/lots-of-data-one-analyst-many-reports-narrative-science/>.
25. Gregory F. Treverton, C. Bryan Gabbard, "Assessing the Tradecraft of Intelligence Analysis," The RAND Corporation, 2008; p.5, [http://www.rand.org/content/dam/rand/pubs/technical\\_reports/2008/RAND\\_TR293.pdf](http://www.rand.org/content/dam/rand/pubs/technical_reports/2008/RAND_TR293.pdf).
26. Ian Crouch, "The Sportswriting Machine," *The New Yorker*, March 26, 2015, <http://www.newyorker.com/news/sporting-scene/the-sportswriting-machine?mbid=rss>

27. Allen Weiner, "Fantasy football and the cold future of robot journalism," *The Kernel*, August 24, 2014, <http://kernelmag.dailydot.com/issue-sections/features-issue-sections/10097/fantasy-football-and-the-cold-future-of-robot-journalism/>.
28. NarrativeScience Whitepaper on "Narrative Analytics: From Data, to insight, to action," <https://www.narrativescience.com/narrative-analytics-white-paper>.
29. Steven Levy, "Can an algorithm write a better news story than a human reporter," *Wired*, April 24, 2012, <http://www.wired.com/2012/04/can-an-algorithm-write-a-better-news-story-than-a-human-reporter/>.
30. National Photographic Interpretation Center, First Broad-Ocean Area Missile Test from Shuangchengzi Missile Test Range Complex SSM, China, December 1980; CIA Records Search Tool (CREST), National Archives, College Park, sanitized copy approved for release 2010/08/18, <http://nsarchive.gwu.edu/nukevault/ebb489/docs/defsmac05.pdf>.
31. Ciara Byrne, "How the UN's New Data Lab in Indonesia Uses Lab in Indonesia Uses Twitter to Preempt Disaster," *Fast Company*, March 19, 2013, <http://www.fastcolabs.com/3007178/open-company/how-uns-new-data-lab-indonesia-uses-twitter-preempt-disaster>.
32. Stephanie Yang, "Can you tell the difference between a robot and a stock analyst?" *The Wall Street Journal*, July 9, 2015.
33. Will Oremus, "The Prose of the Machines," *Slate*, July 14, 2014, [http://www.slate.com/articles/technology/technology/2014/07/automated\\_insights\\_to\\_write\\_ap\\_earnings\\_reports\\_why\\_robots\\_can\\_t\\_take\\_journalists.html](http://www.slate.com/articles/technology/technology/2014/07/automated_insights_to_write_ap_earnings_reports_why_robots_can_t_take_journalists.html).
34. Joe Fassler, "Can the Computers at Narrative Science Replace Paid Writers," *The Atlantic*, April 12, 2012, <http://www.theatlantic.com/entertainment/archive/2012/04/can-the-computers-at-narrative-science-replace-paid-writers/255631/>.
35. John A. Kringen, "Keeping Watch on the World: Rethinking the Concept of Global Coverage in the U.S. Intelligence Community," *Studies in Intelligence* 59, No. 3 (September 2015), <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/csi-studies/studies/vol-59-no-3/keeping-watch-on-the-world.html>.
36. Erin Madigan White, "Automated Earnings Stories Multiply," Associated Press blog, January 29, 2015, <http://blog.ap.org/2015/01/29/automated-earnings-stories-multiply/>.
37. Will Oremus, "The First News Report on the L.A. Earthquake Was Written By a Report," *Slate.com*, March 17, 2014, [http://www.slate.com/blogs/future\\_tense/2014/03/17/quakebot\\_los\\_angeles\\_times\\_robot\\_journalist\\_writes\\_article\\_on\\_la\\_earthquake.html](http://www.slate.com/blogs/future_tense/2014/03/17/quakebot_los_angeles_times_robot_journalist_writes_article_on_la_earthquake.html).
38. Kint Finely, "In the Future, Robots Will Write News That's All About You," *Wired*, March 8, 2015, <http://www.wired.com/2015/03/future-news-robots-writing-audiences-one/>.
39. Remarks as delivered by The Honorable James R. Clapper, Director of National Intelligence, Opening Statement to Worldwide Threat Assessment Hearing, Senate Armed Services Committee, February 26, 2015. (<http://www.dni.gov/files/documents/2015%20WTA%20As%20Delivered%20DNI%20Oral%20Statement.pdf>)

