Artificial Intelligence and Online Spirituality

with Mohammad Yaqub Chaudhary, "Augmented Reality, Artificial Intelligence, and the Re-Enchantment of the World"; and William Young, "Reverend Robot: Automation and Clergy."

REVERENDROBOT: AUTOMATION AND CLERGY

by William Young

Abstract. Digital technology, including artificial intelligence, is having a dramatic impact on the professions of medicine, law, journalism, finance, and others. Some suggest that clergy will also be affected. We describe recent progress in designing artificially intelligent systems, suggesting that this is possible, perhaps even likely. We investigate ways in which technology currently is affecting ministry and outline some plausible scenarios in which digital systems could supplement or supplant clergy in some areas, specifically preaching and pastoral care. We also raise some theological issues raised by the use of digital systems in ministry.

Keywords: artificial intelligence; clergy; pastoral care; preaching; technology; theology

---

Interviewing father and son authors Richard and Daniel Susskind in April of 2016, NPR's On Point host Tom Ashbrook explored the thesis of their recent book, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Susskind and Susskind 2015). Society, the Susskinds aver, is on the brink of fundamental and irreversible changes in the way that professional expertise is delivered. Specifically, "increasingly capable machines will transform the work of professionals, giving rise to new ways of sharing practical expertise in society" (Susskind and Susskind 2015, 303). Current professions are an artifact, they say, "built to meet a particular set of needs [but] in their current form, will no longer be the best answer to those needs" (Susskind and Susskind 2015, 3). Citing professions the Susskinds believe at risk of serious disruption from automation, Ashbrook notes: "People have talked about lawyers and the law for a while. You've got doctors right there; we think of those immediately. But

William Young is Associate Professor of Instruction and Research Scientist in the Computer Science Department at the University of Texas at Austin, Austin, TX, USA. He is also ordained clergy with the Universal Fellowship of Metropolitan Community Churches; e-mail: byoung@cs.utexas.edu.

 ISSN 0591-2385

you also tick off accountants, management consultants, architects, teachers, journalists and, heaven forfend, even clergy" (Ashbrook 2016). Yes, the Susskinds believe that not even clergy are immune from the disruptive technologies that are altering the landscape of work in the twenty-first century.

Clergy, they argue, function largely as gatekeepers mediating access between believers and sacred texts or between worshipers and God, and traditionally operate within the context of localized religious communities. But current and emerging technologies provide alternative, nonlocalized means to deliver religious knowledge and analysis—Bible websites, streamed worship services, clergy blogs, digital-only congregations, video pilgrimages, online religious scholarship, and so on—which could have a profound impact on the religious status quo.

These platforms and systems can strengthen people's existing religious beliefs and can also support entirely new internet-based faiths. Sometimes, though, the services encourage people to question their inherited religions. Almost two-fifths of young practicing Christians in the United States use the internet to fact-check the statements that their religious leaders make (Susskind and Susskind 2015, 66).

Just as Gutenberg's press democratized access to the Bible in the fifteenth century and broke the clerical monopoly on scriptural knowledge, the internet is today democratizing access to the spiritual realm, with a potentially devastating impact on the current order, the Susskinds aver.

Not everyone agrees. Legal scholar Frank Pasquale believes that the Susskinds' claims are vastly overblown because "software has a difficult time mimicking the spontaneity, creativity, flexibility, and perceptiveness that are the hallmarks of good nurses, surgeons, psychiatrists, pediatricians" and, we assume, clergy. In a scathing rejoinder to the Susskinds, he mocks them as predicting "the emergence of bot preachers, ministering to avatars in Second Life's virtual cathedrals" (Pasquale 2016). (It's worth noting, however, that since 2007, a virtual "Anglican Cathedral" has existed in the Second Life virtual world, with daily worship, a weekly Bible-study class, and counseling services.) Nevertheless, many tasks currently performed by nurses, surgeons, psychiatrists, and pediatricians are susceptible to automation. Surely, the same is true of clergy.

An extensive 2013 Oxford University study (Frey and Osborne 2013) evaluated 702 common occupations. Using a sophisticated ranking algorithm, each was scored for likelihood of automation. Many current jobs are in the crosshairs of computerization. The study concludes, for example, that telemarketers, loan officers, and cashiers are likely to see their jobs replaced by automation with above 95 percent likelihood. Clergy, on the other hand, had only a 0.83 percent chance of losing employment to automation according to the study. Seminary students, take heart!

So, is emerging technology likely to disrupt the profession of clergy in the way that it promises to disrupt other white-collar professions? In a word, no. In the foreseeable future, churches would not be hiring/purchasing robot clergy to fill a vacant pulpit. But technology definitely will impact the way that clergy perform their functions, for good or ill, even in the short term. In the long run, who knows?

In this article, we examine the likely impact of digital technology, especially artificial intelligence (AI), on the profession of clergy. We identify two distinct technological trends that promise to disrupt the practice of working clergy in America. We will focus primarily on Christian clergy, but our conclusions differentially apply to non-Christian religious leaders as well.

The first trend we call online spirituality. This is the delivery of religious "products"—worship services, pastoral care and counseling, inspirational content, sermons, reflections, prayers, scriptural scholarship, and retreats—via digital media rather than within traditional congregational settings. A cursory Google search shows that these are already widely available. The modes of delivery may be nontraditional, but the content is still produced by human beings, often clergy working within the context of traditional churches. Tellingly, nearly all of the Susskinds' examples (Susskind and Susskind 2015, 61–66) of current incursions of technology into the sphere of religion are of this ilk.

A second, more futuristic possibility looms: that artificially intelligent digital systems actually will supplement or supplant human clergy in significant ways— writing sermons, delivering pastoral care, conducting scriptural/theological research, or performing sacramental functions. We will refer to this possibility, perhaps a bit whimsically, as digital clergy. To date, instances are few and mostly one-offs. For example, a robot has officiated at a Japanese wedding (Daily Mail 2010); another is available to perform funerals (Gibbs 2017). However, such technological incursions are rapidly disrupting other professions such as medicine, law, journalism, and finance.

Automated systems currently diagnose diseases, perform surgery, write wills, do legal research, compose news stories, and trade stocks. The fact that computer scientists and AI researchers have not as yet expended much effort automating many of the tasks traditionally performed by clergy does not mean that they could not do so or would not do so in the future, given sufficient incentives. How realistic then is the prospect that significant functions currently performed by human clergy can and will be performed by the "increasingly capable machines" heralded by the Susskinds?

The theological and pastoral issues raised by online spirituality have been discussed elsewhere (e.g., Campbell 2013). Some theologians also have discussed the issues entailed by AI (e.g., Herzfeld 2002). But the possibility of digital clergy raises some separate theological questions: To what extent can a digital system legitimately be said to minister? Does the

ontological status of a machine comport with our notion of a sacramental agent? Can a machine be ordained? Assuming that digital clergy can exist, how important is it that they function in the same way as human clergy? Must a digital minister have genuine religious conviction (faith)? What would that even mean?

Our goal in this article is twofold. First, we will describe some of the ways in which technology is currently changing the practice of ministry and some ways in which it promises—some would say threatens—to do so. Second, we will point to some theological and practical issues that these changes raise. The article is arranged as follows. In the following section, we discuss briefly the current state of the art in artificially intelligent systems, especially those ushering in the "postprofessional society" predicted by the Susskinds. The section on "Online Spirituality" describes how current trends in digital religion are affecting the practice of religion today and the likely effect on clergy in the future. Then, in the section on "Digital Clergy," based on trends in other professions, we describe some technologies that could automate key functions clergy perform today, specifically preaching and pastoral care. The section "Theological Issues" describes some theological issues raised by the use of digital systems to perform ministerial functions. In the final section, we draw some conclusions of this research.

The AI Revolution

AI promises to be one of the most disruptive technologies of the twenty-first century. A recent study panel of experts on AI from academia and industry provides the following definition of AI: "Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their nervous systems and bodies to sense, learn, reason, and take action" (Stanford University 2016). Progress in AI has been uneven since computer pioneer John McCarthy coined the term AI in 1956. Periods of progress and promise, often accompanied by excessive hype, have been punctuated by the so-called AI winters during which progress stalled and funding dried up. However, the past decade has seen accelerating progress. Consider the following:

- Video games powered by computer graphics, vision, and AI planning today constitute a larger entertainment industry than Hollywood.
- Automated systems routinely outcompete the best human experts at chess and even the popular TV show *Jeopardy*. For the first time, a computer program has become competitive at world-champion level Go, one of the most strategic of board games.
- Google's Translate utility now translates among 140 natural languages, rivaling skilled human translators (Lewis-Kraus 2016).

- A recent article in the Smithsonian's online magazine lists nine surprising tasks that robots are currently performing, including cooking dinners, filling prescriptions, checking guests into hotels, training athletes, and even riding camels (Matchar 2017).

Such advances in machine proficiency largely have been driven by two major factors: Moore's Law and advanced sophistication in machine learning (ML) algorithms.

Moore's Law describes exponential growth in the production of semiconductor electronics. Originally proposed by Intel founder Gordon Moore in 1965, Moore's Law predicted that component density (transistors per square unit) on integrated circuits would double approximately every 1.5–2 years, a trend that has now prevailed for over four decades, though futurist and inventor Ray Kurzweil has noted (Kurzweil 2001) that computers were actually doubling in power long before transistors were invented.

Accompanied by advances in storage capacity, materials science, fiber optic transmission, and multicore architectures, the result of this technological progress has been a remarkable explosion in the computational prowess of digital systems. In introducing Moore at a 2015 event, Intel CEO Brian Krzanich observed that, compared to Intel's first generation 1971 microchip, Intel's latest chip offered 3,500 times more performance, was 90,000 times more energy-efficient, and some 60,000 times cheaper (Friedman 2015). Krzanich added that, had a 1971 Volkswagen Beetle improved at the same rate as Intel microchips, "You would be able to go with that car 300,000 miles per hour. You would get two million miles per gallon of gas, and all that for the mere cost of 4 cents!" Moreover, that car would be the size of an ant. Your smartphone contains significantly more computing power than the machines that sent the Apollo astronauts to the Moon. Such hardware advances have provided the raw brainpower not only for the AI explosion but also for the digitalization of modern society.

The second major factor driving the AI revolution has been remarkable advancement in ML. The goal of ML is to train a computer to perform some cognitive task, for example, recognizing objects within images or interpreting spoken language. One common approach called supervised learning trains the digital system, often a neural network—a pliant digital lattice designed loosely to emulate the type of connectivity found in human brains—with a very large collection of samples of the items to be recognized, appropriately labeled. The system is then tested on unlabeled samples. A correct identification causes certain pathways through the network to be reinforced and others reduced. Over time recognition improves as the system literally rewires itself.

A digital translator, for example, might be trained on a large collection of English documents paired with good Chinese equivalents. It is then set loose on English text without corresponding translations. A human

or automated trainer scores each translation and this feedback prompts strengthening of successful pathways and degrading of unsuccessful ones. It is this approach that led to remarkable recent advances in Google Translate; after being converted to an ML approach, the system "demonstrated overnight improvements roughly equal to the total gains the old one had accrued over its entire [10 year] lifetime" (Lewis-Kraus 2016).

The holy grail of ML, however, is unsupervised learning, in which the network extracts meaningful patterns without being told explicitly what to search for. In principle, a digital system could learn English to Chinese translation without access to paired documents. According to a recent article in *Science*,

> That's possible because languages have strong similarities in the ways words cluster around one another. The words for table and chair, for example, are frequently used together in all languages. So if a computer maps out these co-occurrences like a giant road atlas with words for cities, the maps for different languages will resemble each other, just with different names. A computer can then figure out the best way to overlay one atlas on another. Voilà! You have a bilingual dictionary. (Hutson 2017)

The developing system is stressed by translating sentences from language A to language B and back again. If the result is identical to the original sentence, then the pathway is reinforced; otherwise, it is weakened. Recent advances using this approach promise great benefits for translations among languages, including obscure and endangered languages for which there are not vast corpora of documents with paired translations.

ML using neural nets is not a new technique; the earliest neural models date from the 1940s. But they fell out of favor for decades as other techniques exhibited faster progress. The main problem: simple neural nets can only do simple things. Real progress required the advances in hardware presaged by Moore's Law, the accumulation of vast corpora of training data, and development of "multilayered" architectures. Layering allows the system to search for patterns, then patterns of patterns, and so on. A system to understand text, for example, might comprise a lowest layer to recognize individual words, which are passed to a second layer that supplies a likely meaning gleaned from the context (adjoining words), which are passed to a third layer that assembles them into recognizable phrases, then into recognizable sentences, and so on.

The accelerating progress in intelligent computational systems along with the growing facility of robots on a wide variety of tasks raises an age-old question: How will these emerging technologies impact the workforce? Will they lead to widespread unemployment as many human skills become redundant? Or will they create new employment opportunities that offset their disruptive effects? Such questions have arisen at least since the first industrial revolution when English textile workers battled the introduction

of mechanized looms under the banner of the folkloric Ned Ludd. (Some scholars describe the rise of AI as a fourth industrial revolution [Schwab 2017] or second machine age [Brynjolfsson and McAfee 2014].) Studies, however, consistently have shown that technological innovation, while disruptive to certain professions, historically has produced net increases not only in productivity, but also in employment.

Economists steeped in the history of productivity improvement tend to point out that it is always bad in the messy middle of a transformation, but everything works for the better in the end, as productivity gains give rise to investments in new ventures that provide jobs that were never imagined before. At least historically that has always been true. But others insist that this time, history will not rhyme (Davenport and Kirby 2016, 227).

Historically, it has been blue-collar workers who were supplanted by automated systems. That trend continues. Tasks like ordering a hamburger, buying groceries, and shipping goods require drastically fewer human workers than previously. Self-driving cars are coming; but so are self-driving trucks (Freedman 2017) and even crewless cargo ships (Levander 2017). Robots now lay bricks, operate mining equipment, and make pizzas. Entire classes of jobs have been virtually eliminated: switchboard operators, filing clerks, travel agents, and assembly-line workers. *The MIT Technology Review* reports (Rotman 2017) that 83 percent of current jobs that pay less than $20 an hour are under threat from automation. According to a study commissioned by President Obama, the jobs "threatened by automation are highly concentrated among lower-paid, lower-skilled, and less educated workers" (White House 2016, 2).

Increasingly, however, white-collar and professional jobs are seeing incursions from automation. Physicians increasingly consult computer-based diagnostic assistants such as DXplain and VisualDx to confirm their diagnoses or suggest alternatives (Maron 2017). A legal assistance tool called ROSS, leveraging the brain power of IBM's artificially intelligent supercomputer Watson, can perform in seconds legal research that once took a skilled paralegal hours of work. Each quarter, Associated Press earnings reports for over 3,700 companies go "out to the wire" without any human intervention; the stories are written entirely by software, represent a twelve-fold increase over manual output, and contain fewer errors than the human-produced stories from years past (Miller 2015). Up to 70 percent of Wall Street trades are made by computerized high-frequency trading systems (Barrat 2013, 94). The professions of medicine, law, journalism, and finance likely will never be the same. But, surely, clergy are immune from the usurpations of automation. Or are they?

In the following two sections, we will outline ways in which technology may impact ministry both in the near term and in some plausible futures. First, we will outline ways in which current technologies are changing the religious landscape. Then, in the following section, we will speculate about

how future AI systems might usurp some of the traditional functions of clergy, specifically preaching and pastoral care.

Online Spirituality

Recall that by online spirituality we denote the delivery of religious products and services via digital media outside of traditional congregational settings. This is more than just using technology in church. Tech-savvy pastors have long supplemented their liturgy, congregational communication, evangelization, and church management with high-tech, particularly in mega-churches that can afford the hardware, software, and support staff needed to run a tech-heavy organization. Some mega-churches even use telepresence to allow the pastor to preach simultaneously at multiple campuses (Blake 2010).

Many churches, both large and small, stream worship services online or make them available to be watched anytime on YouTube. Increasingly, accessible web development tools, blogging sites, and new media make it easy even for small church pastors to have a significant digital presence. This facilitates access to their message, allowing congregants, including the elderly or infirm, remote access to congregational resources. It also allows potential congregants to sample the offerings of the local church prior to a visit. But offering online access to these products and services means they become available also to a broader audience, many of whom may have no interest in affiliating with or supporting the local church.

Church leaders recognize the value of the enhanced visibility that digital access affords. Pope Benedict XVI in his 2010 message for World Communication Day challenged ministers "to proclaim the Gospel by employing the latest generations of audiovisual resources (images, video, animated features, blogs, websites) [to] open up new vistas for dialogue, evangelization and catechesis" (Pope Benedict XVI 2010). Nearly every denomination now has a web presence, often providing extensive doctrinal, liturgical, and counseling resources. Mormon.org, for example, provides 24/7 access to Latter-Day Saints missionaries on call; it is reported that this multiplies the conversion rate of a missionary by sevenfold over door-to-door evangelization (Bosker 2014). Religious leaders increasingly communicate via social media; the Pope has more than 19 million Twitter followers, the Dalai Lama more than 10 million. Spirituality-related courses and web-based retreats are available from hundreds of providers.

A 2013 study (Buie and Blythe 2013) found over six thousand apps on the iTunes App Store related to spirituality and religion, including selections for education, prayer exchange, comparative religion, guidance for daily living, and teachings and sermons from specific religious leaders. A Pew Research study (Pew Research Center 2014) reported that around 20 percent of American adults used social media to share their religious

faith, "about the same percentage that tune in to religious talk radio, watch religious TV programs or listen to Christian rock music."

Electronic media provide access to the products of religion, but without any direct contact with clergy. It is possible to sample one hundred different sermons every day on YouTube without exhausting the ever-replenished pool. Many clergy also publish blogs to spread their message; a Google search for "clergy blog" returns almost 900,000 results. The facile online availability of religious goods and services almost certainly exacerbates the ongoing erosion of traditional religious observance, specifically attendance at brick and mortar houses of worship.

As religious products become ever more readily accessible, they are also increasingly becoming mere commodities. Why bother reading a blog or listening to a sermon from your local pastor when you have instantaneous access to some of the most charismatic religious leaders in the world? This trend is already evident as TV preachers such as Kenneth Copeland, T. D. Jakes, and Joel Osteen cultivate extensive "virtual" congregations, in addition to the mega-churches they pastor. Such factors could drive the religious landscape in the direction of what Duke University economist Philip Cook and others have called a "winner-take-all market," where a handful of superstar performers garner a larger and larger proportion of the available audience, and reap the lion's share of the rewards. This pattern, long common in entertainment and sports, is becoming increasingly prevalent in other fields—law, journalism, consulting, investment banking, management, and fashion (Frank 1994). It could become the new norm in the religious "marketplace" as well.

As most religious services become available remotely, the demand for local clergy shrinks. It is easy to imagine a plausible future in which a pool of clergy for hire fulfills residual demand for those occasions when a local clergyperson is required, say, to officiate at a wedding or funeral. Alternatively, a quick visit to the Universal Life Church (ULC) website can result in an online ordination. With more than twenty million ordained "ministers," ULC "has become one of Earth's largest and most active religious organizations" (Universal Life Church 2018). For the Susskinds and others, these trends suggest that religion and spirituality are trending in the do-it-yourself direction, potentially significantly reducing the need for local professional clergy.

Clearly, many individuals still will desire the fellowship of a local church community and an ongoing relationship with a spiritual mentor. Moreover, some studies have found that online religious activity generally serves to supplement rather than supplant offline church activity (Campbell and Garner 2016, 66–67). Nevertheless, the "nones" and the "spiritual but not religious" are among the fastest growing groups in the American religious landscape. Is this any surprise when the products of religion are so readily

available without the commitment, inconvenience, or expense of joining a local congregation on Sunday morning?

Digital Clergy

In this section, we venture beyond what is currently happening in the online realm to speculate about the following question: Could automated systems ever supplement or supplant clergy in such core functions as writing/delivering sermons or providing pastoral care? Current technological trends in established professions such as law, medicine, journalism, and finance suggest that the answer is yes. Whether or not it is a theologically or practically sound idea is a separate question. We will return to it in the section entitled "Theological Issues." For the remainder of this section, we consider what assumptions underlie our belief that digital clergy are likely to emerge in the coming decades, and then speculate about what that might look like, specifically in the areas of preaching and pastoral care.

Assumptions

Two key assumptions underlie our analysis. The first assumption is that developments in AI will continue on the current trajectory, leading to further advances in ML, neural networks, speech recognition, natural language processing, vision, and robotics. This assumption is not particularly controversial. The accelerating progress over the past decade described above has yielded a "constellation of mainstream technologies that are having a substantial impact on everyday lives" (Stanford University 2016). The resulting economic impact, both current and potential, on entertainment, finance, manufacturing, transportation, education, and healthcare means that vast amounts of capital are flowing into AI research. According to one estimate, tech companies poured between $20 and $30 billion into AI research in 2016 alone (Columbus 2017). Universities are hard pressed to hire or retain top AI talent because of the demand from industry. These factors ensure that "AI spring" has arrived and that advances likely will continue apace.

Slightly more controversial are questions of the eventual reach of computational competence. For example, will digital systems ever attain human-level machine intelligence (HLMI) (also called "general artificial intelligence" [GAI] or "strong AI") and reach a point where machines "can carry out most human professions at least as well as a typical human"? (Murphy 2012). Current AI systems are nowhere close to HLMI; existing applications are highly tailored to accomplish particular tasks, often sublimely competent within their narrow domain but useless outside it.

Consider, for example, IBM's artificially intelligent supercomputer Watson, which famously beat human champions Ken Jennings and Brad Rutter at *Jeopardy*! Though an AI tour de force, Watson was

purpose-built for the contest as, in essence, a specialized data mining program employing number crunching, pattern matching, and searching in a finely choreographed computational ballet. It brought to the contest an additional advantage: instantaneous access to more than 200 million pages of structured and unstructured content, including the full text of Wikipedia. Although Watson has since been applied by IBM to a variety of other problems, all are similarly analysis-intensive. Tellingly, none requires seeing or moving through the world. Cognitive scientist and Pulitzer Prize-winning author Douglas Hofstadter disparaged Watson as "just a text search algorithm connected to a database." Even though it is a fantastically competent text search algorithm, in this regard Watson exemplifies the paradox expressed by robotics pioneer Hans Moravec as follows: "It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility" (Moravec 1988). As yet, it is unclear when or if this will change.

Several recent surveys have asked AI researchers whether and when they believe HMLI might be achieved. AI Researcher Nick Bostrom summarizes the results as follows: "the combined sample gave the following (median) estimate: 10 percent probability by 2022, 50 percent probability by 2040, and 90 percent probability by 2090" (Bostrom 2014, 23). Note that few researchers surveyed doubted that HLMI would happen within this century. Moreover, Bostrom, like many others in the AI community, expect that HLMI is not the final stop on this route. "Just a short distance farther along the tracks, is superhuman-level machine intelligence. The train may not pause or even decelerate at Humanville Station. It is likely to swoosh right by" (Bostrom 2014, 5). Many commentators, including Bostrom, worry that if computers attain superhuman intelligence, it will have profound and potentially devastating effects on the human race (Barrat 2013; Bostrom 2014; Ford 2015; Tegmark 2017). The Stanford AI study panel, on the other hand, found "no cause for concern that AI is an imminent threat to humankind" (Stanford University 2016).

Our second assumption in this analysis is that the potential intellectual or economic payoffs will suffice to incentivize AI researchers to undertake automating functions currently performed by clergy. In other professional fields impacted by automation such as law and medicine, powerful economic drivers are at work. For instance, the CEO of ROSS Intelligence, which markets the Watson-powered legal assistance tool mentioned above, reports that the tool eliminates 20–30 hours of human time spent on discovery per case, for a dramatic cost savings to law firms (Mangan 2017). In healthcare, a 2017 report from Accenture predicts that clinical AI applications will save the industry $150 billion over the next 10 years (Collier et al. 2018). It is no surprise then that AI and robotics researchers are clamoring for a share of these lucrative markets.

Churches do not afford analogous opportunities for such economic windfalls. The potential religious market is large—a recent study (Randal 2017) estimates 384,000 congregations in the United States—but still relatively tiny compared to the number of small businesses in the country (around 28 million according to the U.S. Small Business Administration). A large proportion of churches' budgets is devoted to staff. A study by the Hartford Institute for Religion Research (Roozen 2008) found that Christian churches in the United States spend an average of 45 percent of their total operating budget on staff costs, with an average attendance to full-time staff ratio of 76:1. However, most U.S. churches have weekly attendance under one hundred (Earls 2016).

That implies that most churches only have one full-time employee, the pastor; and a large percentage of small churches cannot afford even a full-time minister. Many churches outsource noncore staff functions such as administrative support, bookkeeping, web development, and publishing. Except for the pastor's salary, there are few significant economic gains to be had from automation in a typical American church. In larger congregations with multiple staffers, economic benefits could be had from automation that eliminated staff positions, but most staff clergy, as opposed to administrative staff, provide direct ministry services that might be difficult to automate.

For these reasons, it is hard to see a strong economic incentive for AI researchers to develop applications specifically for the religious marketplace. However, that may not matter. Automation within churches might emerge instead as a niche market leveraging technologies developed for the much larger and more lucrative commercial sector. For example, software specialized for church management is a lucrative niche market but likely never would have emerged had not analogous software been developed for more mainstream businesses. Similarly, it is hard to imagine a well-funded research effort to build bespoke sermon writing software, but not an effort to customize existing essay writing applications to that purpose. Companies such as EssaySoft already offer essay writing software targeted to students. It is likely that automation that specifically impacts clergy, if such emerges, will be an adaptation of AI functionality developed for other domains.

In the remainder of this section, we will speculate about how that might look. Specifically, we will discuss how automation could impact two of the more time-consuming duties of the modern pastor: preaching and pastoral care. According to one study (McMillan 2002), these are the tasks to which the majority of a Christian pastor's time is devoted: preparing for preaching and worship (33 percent) and providing pastoral care (19 percent). Other significant pastoral tasks—administering the congregation's work and attending meetings (15 percent), teaching and training people for ministry (13 percent), and attending to denominational and community affairs (6

percent)—seem a bit less susceptible to automation, though AI is making significant inroads into teaching.

Preaching

Depending on a church's theology of the Word, preaching may be the pastor's most important weekly duty, or a relatively minor one. Automating the preaching task might refer to composing the sermon, delivering it, or both. Synthesizing spoken speech from a written text is relatively straightforward; speech synthesizers have been available since the early 1990s. Here, we will concentrate on the much more challenging task of composing sermons. Many preachers already avail themselves of online preaching resources, including sermon outlines and even full sermons. Is it possible that these could be produced by an automated system?

The AI subfield of natural language generation (NLG) has delivered some impressive results in narrow domains. Commercial software systems already produce a large volume of text for human consumption (Perera and Nand 2017; Gatt and Krahmer 2018). In most cases, however, these programs largely string together text in well-structured contexts such as reporting on an athletic competition, compiling a corporate earnings report, composing a simple will, or writing a business letter. The basic facts and structure of the narrative are known; the software merely puts flesh on the bones.

Composing a sermon arguably is a much more highly creative endeavor. A sermon must be theologically grounded, soundly argued, and culturally relevant. Ideally, it also will be moving and inspirational with rhetorical flourish—a far cry from a corporate earnings report. To complicate matters, a sermon's content should be informed by the personal theologies of the homilist and the intended audience. It is unlikely, for example, that a sermon written by a Church of Christ pastor would be suitable for a United Church of Christ congregation, or vice versa; the theological perspectives differ too radically. A commodity sermon-writing program presumably would be configurable to generate output aimed at different points on the theological spectrum. One approach might be to "train" a machine-learning sermon writing program on a corpus of sermons collected to reflect a specific tradition or theological standpoint.

Generating acceptable free-form compositions such as essays or sermons is currently considered beyond the state of the art in NLG technology. Examples where computer-generated "research" papers have been accepted for publication—more than 120 computer-generated nonsense papers have been unwittingly published by conferences and journals—speak more to the poor quality of academic peer review than to the high quality of NLG output. However, recent advances in NLG technology suggest that sermon-writing programs may not be too far in the future. In June 2018,

an IBM computer program called Project Debater (Metz and Lohr 2018) squared off against an Israeli college debating champion on the subject of government subsidies for space exploration. Debater stitched together a coherent argument on the topic, marshaling information from an extensive database, deciding which facts supported its own argument, and parsing and rebutting its opponent's arguments. One observer noted that Debater even displayed procatalepsis—the ability to anticipate and preemptively attack its opponent's arguments (Shankland 2018).

Currently, Debater's interactions are tightly constrained and include a significant proportion of canned content. Its debate topics are limited to around one hundred topics. Its prose is rather stilted and it sometimes cites irrelevant sources. Nevertheless, it illustrates some of the strengths of an automated system versus a human. For example, Debater has instant access to over 300 million news articles and scholarly papers indexed for quick search. The sheer volume of information accessible to such a system means that it draws from a much more extensive knowledge base than would be possible for any human. A machine-learning, sermon-writing AI with access to a similarly extensive corpus of existing sermons, theological resources, concordances, and scriptural exegeses could mine them for anecdotes, scriptural references, and theological arguments. It is not hard to imagine a system that could generate a sermon on any biblical text from any theological standpoint.

Search engines and online tools already greatly expand the homilist's available research reach. The pastor's physical library is increasingly being replaced by e-books, podcasts, websites, Bible software, and online college and seminary courses. As more capable AI-powered tools become available, they will be used. Although fully automated sermon generation is probably decades away, more reliance on technology in sermon preparation is just around the corner.

It is easy to imagine a progression of steps that lead from the current set of tools to fully automated sermon composition:

(1) Currently available spelling, grammar, and syntax checkers continue to evolve.

(2) Automated "research assistants" become available to assemble materials from the web on a pastor-supplied topic or sermon outline.

(3) Narrative planning systems become capable of producing a sermon plan or outline to be filled in by a human.

(4) Finally, AI systems produce fully developed sermons.

It is even possible that one or more of these steps could combine narrative generation with an "interest model" to gauge whether the sermon produced will engage the audience (McIntyre and Lapata 2009). Such

an enhancement to our sermon-writing tool suite might finally eliminate boring sermons entirely.

As NLG technology progresses, systems such as Debater will become increasingly adept, fluent, articulate, and consistent. Combined with advanced ML, these advances suggest that sermon-writing systems are not infeasible in the foreseeable future.

Pastoral Care and Spiritual Counseling

One of the most sensitive tasks of the pastor is spiritual counseling and care. It is hard to imagine that these could be automated. However, in a very real sense, many people already rely on digital systems to provide help and comfort. Virtual help agents or chatbots—Apple's Siri, Amazon's Alexa, Microsoft's Cortana, and other "digital assistants"—answer questions, make recommendations, and manipulate the physical environment through integration with web-based services. Users interact via a natural language interface; the systems are generally quite good at voice recognition, natural language understanding, and carrying out simple tasks such as finding a local restaurant or adjusting the thermostat. But could such chatbot technology evolve into a therapeutic or pastoral care role?

In some sense, it already has. The UK edition of *Wired* (Molteni 2017) reports that virtual help agents "have taken on surprisingly sensitive jobs in modern society: counseling Syrian refugees fleeing civil war, creating quiet spaces of contemplation for millions of Chinese living in populated cities, and helping Australians access national disability benefits." One notable entry in this space is Woebot, a "talk therapy chatbot" developed by a team of Stanford psychologists and AI experts. It dialogues with users in the style of a cognitive behavioral therapist, attempting to prompt users to recast negative thoughts into a more objective light. Patients are encouraged to open up about their emotional responses to life events. A study published in the *Journal of Medical Internet Research* (Fitzpatrick et al. 2017) found that some users reported significant reduction in symptoms of depression and anxiety after interacting with Woebot.

A different digital system developed at Northeastern University and Boston Medical Center adds an explicitly spiritual component to the dialogue, specifically talk about end-of-life issues (Utami et al. 2017). This "virtual conversational palliative care coach" is designed to help individuals "manage symptoms, reduce stress, identify and address unmet spiritual needs, and support advanced care planning." To test the acceptability to older individuals of a digital agent providing spiritual counseling, the development team provided various modes of interaction with differing degrees of spiritual engagement. In "prospiritual mode," the agent queries the user about his or her religious orientation and tailors the dialogue accordingly in a supportive fashion. Unsurprisingly, some respondents were

uncomfortable interacting with an automated system. Others, however, reported finding it "easier to talk to a computer agent [than a human]" about end-of-life issues.

This validates other studies finding that users are often significantly more candid and self-disclosing when conversing with a digital system compared to a human therapist. Apparently, the embarrassment factor is greatly diminished when conveying sensitive information to a nonhuman. "There's nothing like venting to an anonymous algorithm to lift [the] fear of judgment," notes Alison Darcy, one of the psychologists behind Woebot (quoted in Molteni 2017). This willingness to disclose to digital agents raises the concern that users might be seduced into oversharing sensitive personal data. This concern manifests itself with Woebot; the chatbot is built on Facebook Messenger that has access to participants' data.

Chatbots may provide solace or advice in verbal or written form. But what about the comforting presence provided by an embodied companion? Toy company Hasbro markets several cuddly animatronic dogs and cats designed to provide comfort and companionship primarily to elders no longer able to care for a live animal. Paro, an adorable mechanical seal developed at Japan's National Institute of Advanced Industrial Science and Technology, is used as an interactive therapeutic tool in hospitals and senior living facilities across the United States. One study (Robinson et al. 2013) found that sustained interaction with Paro decreased loneliness even more than structured activities with other humans. These cuddly "carebots" provide an embodied comforting presence similar to that provided by a pet, but with the additional benefit that some may provide additional services such as facilitating social media use, monitoring blood pressure, and ensuring that medications are taken on time.

If an animatronic seal, why not an animatronic humanoid? A number of humanoid robots have been introduced, including Pepper, a semihumanoid robot manufactured by SoftBank Robotics. Pepper is designed to function as a "genuine humanoid companion created to communicate with you in the most natural and intuitive way, through his body movements and his voice. Pepper gradually memorizes your personality traits, your preferences, and adapts himself to your tastes and habits" (Softbank Robotics 2018). Pepper has been deployed as a receptionist at several offices in the United Kingdom and restaurants in Japan, is able to identify visitors via facial recognition, and chat with prospective clients.

The designers of Pepper, as with many humanoid robots on the market or in development, did not attempt to replicate human facial features, perhaps with good reason. It was discovered by robotics pioneer Masahiro Mori in 1970 that an "almost human" level of realism in an artificial entity prompts a feeling of creepiness or revulsion in many observers, an effect Mori labeled "the uncanny valley" (Mori 2012). Jibo, a Japanese "family

robot" introduced in 2014, was deliberately designed to be less humanlike in appearance than its predecessor Kismet precisely for that reason. There are some notable exceptions to this tendency to avoid mimicking human features. For example, Bina48 (Hanson Robotics 2018) is a lifelike robot head created in the image of Bina Aspen Rothblatt, wife of author Martine Rothblatt. Not only is Bina48's face a replica of Bina Rothblatt's, the android is programmed with her "memories, attitudes, beliefs and mannerisms." It displays very lifelike facial expressions and movement. Still, the technology is nowhere near fooling a human interlocutor. As the technology improves and humans become increasingly comfortable interacting with nonhuman systems, carebots may increasingly simulate human features, gestures, and mannerisms.

An automated system potentially has significant advantages over a human therapist or pastoral care provider. AI systems would be available 24/7, be unfailingly attentive and patient, remember the details of every interaction, and pick up on subtle cues that a human might miss. They would have access to a vast corpus of clinical research literature and could be deployed almost anywhere nearly instantaneously. Moreover, the cost of "educating" a new provider would be miniscule and each successive generation would perform better than the preceding one, drawing upon an ever-expanding knowledge base of therapeutic techniques drawn from literature and its own and predecessors' interactions with clients.

Of course, pastoral counseling is quite a different modality than psychological counseling. As with preaching, pastoral care is deeply informed by the theology of the provider. A pastoral counselor's advice to a client concerned about his or her sexuality, for example, may differ radically depending on the religious tradition of the provider. Perhaps, an AI pastoral carebot would be configurable to provide advice and comfort at different points along the theological spectrum. As an example of what that might look like, the AI palliative care agent described above supports a variety of religious traditions including Christianity, Judaism, Islam, Hinduism, Buddhism, and Sikhism. It also functions within the context of atheism, spiritual humanism, and secular humanism. It tailors its dialogue to the religious orientation of the user.

Existing systems such as those described above fall far short of a fully automated pastoral care provider. However, the technology is advancing rapidly. Future digital pastoral care providers will be able to sense the emotional state of the client, not only through verbal cues but also through body language and other nonverbal cues. Current automated systems that attempt to display emotional intelligence are limited, often clunky in their interactions, and sometimes wildly inappropriate. But this is an area of intense research driven, as usual, by ML. There is little doubt that machines with enhanced emotional quotient are on the horizon.

Theological Issues

Some profound theological issues arise from the very prospect of advanced AI. To name just a few:

- What is the ontological status of a truly intelligent automated system, assuming such an entity is possible? Could such a system have a soul? Would it have rights?
- What would the existence of such a system say about the unique status of human beings created, according to Judeo-Christian theology, in the image of God?
- What constraints does morality impose on a scientific community racing to develop technologies that at the very least will have profound social and economic impacts and, according to some observers, could threaten the very existence of humanity?
- What are the ethical implications of automation reshaping the workplace with the concomitant effects on human dignity and autonomy?

No doubt, these issues will occupy the scientific and theological communities for decades to come. Some have been addressed elsewhere. See, for example, Herzfeld (2009) and Campbell and Garner (2016).

Here, we will focus narrowly on theological issues raised by the prospect of automated systems undertaking functions currently performed by ordained clergy. Our goal is not to hash out these issues that arise from digital ministry, but merely to raise them. If and when clergybots appear, theologians undoubtedly will engage these issues more fully.

Can a digital system minister? Assume that an automated system could provide precisely the same empathetic emotional and spiritual support as your current clergyperson. Is that even possible? Would it constitute ministry in the sense that we currently understand ministry, or in some new sense?

Can a digital system administer sacraments? Faith traditions that recognize sacraments may also specify who is allowed to confer the sacrament. For example, the Catechism of the Catholic Church specifies that "any person, even someone not baptized, can baptize, if he has the required intention." Such a baptism is always valid, but sometimes illicit. Suppose the baptism were conferred by a robot, pronouncing the appropriate trinitarian formula. Would this satisfy the requirement of "any person"? What about the "required intention"?

Where does faith enter in? Could one legitimately ascribe to a digital system faith in any meaningful sense? Would it matter? If a sermon is spiritually insightful and moving, would it matter that the author does not/could not believe a word of it? As with the AI palliative care agent described above, if a system functions equally competently within Buddhist,

Hindu, and Christian contexts, could it have any religious credibility or authority?

Could a digital system be ordained? Traditionally, religious authority has been established by a process of discernment of one's spiritual calling or divine appointment, or by undergoing an extensive period of training. What does it imply for religious authority if spiritual formation is replaced by uploading? What does ordination imply about the status of the ordinand? Is there a meaningful sense in which a digital system can be said to subscribe to the tenets of one or another denominational orthodoxy?

What about relationships? Many theologians within the Christian tradition (e.g., Karl Barth, Emil Brunner) aver that humans reflect the *imago Dei* precisely in their capacity for social relationships with God and other people. Can humans form genuine relationships with machines allowing for the type of intimate engagement required for ministry to occur?

These are just a few of the issues that arise from the prospect of digital clergy. There is a danger that even considering such questions might be seen as devaluing the profession of (human) clergy. However, as the digital systems become ever more capable, and human members of society become ever more comfortable interacting with automation, there may come a time when such considerations will become necessary.

Conclusions

Despite decades of effort, the field of AI research is still far from achieving artificial general intelligence (AGI). Current AI systems have been likened to "autistic savants," extraordinarily capable within a very limited range of tasks but effectively useless outside that range. But that is not from lack of trying. We saw above the consensus of AI experts that HLMI likely will be attained by the end of the century, if not well before. While AI experts have been notoriously overly optimistic in the past about the trajectory of AI progress, recent efforts have led to some breathtaking breakthroughs that lend credence to these predictions.

Existing systems show that intellectual and creative tasks we currently reserve for human beings can be automated. The fields of medicine, law, journalism, and finance are changing rapidly to adapt to the new technologies. Could the profession of clergy be next? Systems described above—Project Debater, the palliative care chatbot, and others—suggest that sermon writing, pastoral care, and certain other tasks currently performed by clergy could be automated were there sufficient incentives to do so. Focus in AI research is driven by intellectual curiosity on the part of researchers, but also increasingly by strong economic factors. Time will tell if clergy tasks will provide a tempting target for researchers, or if efforts in related fields may reach sufficient maturity to allow them to be repurposed to automate clerical tasks.

In the near term, the artifacts of online spirituality likely will become increasingly available, potentially exacerbating the ongoing decline of attendance at brick and mortar houses of worship. Although clergy are not currently in the crosshairs of automation, there is no guarantee that that situation will continue. As has happened in other professions, it seems likely that increasingly capable AI systems will encroach inexorably on the tasks and prerogatives that define us as human. Theologians will have to confront the issues raised by these developments. All that is certain is that progress in AI will continue apace. How the results will reshape the religious landscape is anyone's guess.

References

Ashbrook, Tom. 2016. *Digital Disruption in White Collar Jobs*. On Point Radio. Boston, MA: WBUR.

Barrat, James. 2013. *Our Final Invention: Artificial Intelligence and the End of the Human Era*. New York, NY: Thomas Dunne Books.

Blake, John. 2010. "'Virtual Preaching' Revolutionizes Sunday Sermons." Cable News Network. http://religion.blogs.cnn.com/2010/07/14/virtual-preaching-revolutionizes-sunday-sermons/ (accessed September 1, 2018).

Bosker, Bianca. 2014. "Hook of Mormon: Inside the Church's Online-Only Missionary Army." *Huffington Post*. http://www.huffingtonpost.com/2014/04/09/mormon-church-online_n_5024251.html (accessed June 18, 2018).

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. New York, NY: Oxford University Press.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W.W. Norton.

Buie, Elizabeth, and Mark Blythe. 2013. "Spirituality: There's an App for That (But Not a Lot of Research)." In *Proceedings of the 2013 Conference on Human Factors in Computing Systems*, 1–10. New York, NY: ACM Press.

Campbell, Heidi A., ed. 2013. *Digital Religion: Understanding Religious Practice in New Media Worlds*. New York, NY: Routledge.

Campbell, Heidi A., and Stephen Garner. 2016. *Networked Theology: Negotiating Faith in Digital Culture*. Grand Rapids, MI: Baker Academic.

Collier, Matthew, Richard Fu, and Lucy Yin. 2018. "Artificial Intelligence (AI): Healthcare's New Nervous System." Accenture. https://www.accenture.com/us-en/insight-artificial-intelligence-healthcare (accessed August 13, 2018).

Columbus, Louis. 2017. "McKinsey's State of Machine Learning and AI, 2017." *Forbes*. http://www.forbes.com/sites/louiscolumbus/2017/07/09/mckinseys-state-of-machine-learning-and-ai-2017 (accessed June 12, 2018).

Daily Mail. 2010. "Tie the 'Bot: Robot Registrar Performs First Ever Wedding Ceremony." https://www.dailymail.co.uk/sciencetech/article-1278973/Robot-registrar-performs-wedding-ceremony.html (accessed September 13, 2018).

Davenport, Thomas H., and Julia Kirby. 2016. *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines*. New York, NY: Harper Business.

Earls, Aaron. 2016. "Majority of American Churches Fall Below 100 in Worship Attendance." LifeWay. https://factsandtrends.net/2016/02/24/majority-of-american-churches-fall-below-100-in-worship-attendance (accessed June 17, 2018).

Fitzpatrick, Kara Kathleen, Alison Darcy, and Molly Vierhile. 2017. "Delivering Cognitive Behavior Therapy to Young Adults with Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial." *Journal of Medical Internet Research* 4:e19.

Ford, Martin. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York, NY: Basic Books.

Frank, Robert. 1994. "Talent and the Winner-Take-All Society." *The American Prospect*. http://prospect.org/article/talent-and-winner-take-all-society (accessed August 15, 2018).
Freedman, David H. 2017. "Self-Driving Trucks." *MIT Technology Review*. https://www.technologyreview.com/s/603493/10-breakthrough-technologies-2017-self-driving-trucks (accessed August 1, 2018).
Frey, Carl Benedikt, and Michael Osborne. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf (accessed July 12, 2018).
Friedman, Thomas L. 2015. "Moore's Law Turns 50." *New York Times*. https://www.nytimes.com/2015/05/13/opinion/thomas-friedman-moores-law-turns-50.html (accessed September 1, 2018).
Gatt, Albert, and Emiel Krahmer. 2018. "Survey of the State of the Art in Natural Language Generation: Core Tasks, Applications and Evaluation." *Journal of Artificial Intelligence Research* 61:65–170.
Gibbs, Samuel. 2017. "The Future of Funerals: Robot Priest Launched to Undercut Human-Led Rites." *The Guardian*. https://www.theguardian.com/technology/2017/aug/23/robot-funerals-priest-launched-softbank-humanoid-robot-pepper-live-streaming (accessed September 13, 2018).
Hanson Robotics. 2018. "Bina48." http://www.hansonrobotics.com/robot/bina48 (accessed September 10, 2018).
Herzfeld, Noreen L. 2002. *In Our Image: Artificial Intelligence and the Human Spirit*. Minneapolis, MN: Fortress Press.
———. 2009. *Technology and Religion: Remaining Human in a Co-Created World*. West Conshohocken, PA: Templeton Press.
Hutson, Matthew. 2017. "Artificial Intelligence Goes Bilingual—Without a Dictionary." *Science Magazine*. <http://www.sciencemag.org/news/2017/11/artificial-intelligence-goes-bilingual-without-dictionary>(accessed August 15, 2018).
Kurzweil, Ray. 2001. "The Law of Accelerating Returns." Kurzweil Accelerating Intelligence. http://www.kurzweilai.net/the-law-of-accelerating-returns (accessed September 14, 2018).
Levander, Oskar. 2017. "Autonomous Ships on the High Seas." *IEEE Spectrum* 54(2):26–31.
Lewis-Kraus, Gideon. 2016. "The Great A.I. Awakening." *The New York Times Magazine*. https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html
Mangan, Dan. 2017. "Lawyers Could be the Next Profession to be Replaced by Computers." CNBC. https://www.cnbc.com/2017/02/17/lawyers-could-be-replaced-by-artificial-intelligence.html (accessed July 14, 2018).
Maron, Dina Fina. 2017. "The Computer Will See You Now." *Scientific American* 317(2):24–25.
Matchar, Emily. 2017. "Nine Tasks Robots Can Do that May Surprise You." *Smithsonian Magazine*. https://www.smithsonianmag.com/innovation/nine-tasks-robots-can-do-that-may-surprise-you-180964729 (accessed June 20, 2018).
McIntyre, Neil, and Mirella Lapata. 2009. "Learning to Tell Tales: A Data-Driven Approach to Story Generation." In *Proceedings of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, 217–25. Stroudsburg, PA: Association for Computational Linguistics.
McMillan, Becky R. 2002. "What Do Clergy Do All Week?" *Pulpit & Pew*. http://pulpitandpew.org/node/823/ (accessed June 20, 2018).
Metz, Cade, and Steve Lohr. 2018. "IBM Unveils System that 'Debates' with Humans." *The New York Times*. https://www.nytimes.com/2018/06/18/technology/ibm-debater-artificial-intelligence.html (accessed July 13, 2018).
Miller, Ross. 2015. "AP's 'Robot Journalists' Are Writing Their Own Stories Now." *The Verge*. https://www.theverge.com/2015/1/29/7939067/ap-journalism-automation-robots-financial-reporting (accessed July 13, 2018).
Molteni. 2017. "The Chatbot Therapist Will See You Now." *Wired*. https://www.wired.com/2017/06/facebook-messenger-woebot-chatbot-therapist/ (accessed September 4, 2018).
Moravec, Hans. 1988. *Mind Children*. Cambridge, MA: Harvard University Press.
Mori, Masahiro. 2012. "The Uncanny Valley." *IEEE Robotics and Automation Magazine* 19(2):98–100.

Murphy, Kevin P. 2012. *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT Press.
Pasquale, Frank. 2016. "Automating the Professions: Utopian Pipe Dream or Dystopian Nightmare." *Los Angeles Review of Books*, March 15.
Perera, Rivindu, and Parma Nand. 2017. "Recent Advances in Natural Language Generation: A Survey and Classification of the Empirical Literature." *Computing and Informatics* 36(1):1–32.
Pew Research Center. 2014. *Religion and Electronic Media: One-in-Five Americans Share Their Faith Online*. Washington, DC: Pew Research Center.
Pope Benedict XVI. 2010. "The Priest and Pastoral Ministry in a Digital World: New Media at the Service of the Word." http://w2.vatican.va/content/benedict-xvi/en/messages/communications/documents/hf_ben-xvi_mes_20100124_44th-world-communications-day.htm (accessed June 30, 2018).
Randal, Rebecca. 2017. "How Many Churches Does America Have? More than Expected." *Christianity Today*. http://www.christianitytoday.com/news/2017/september/how-many-churches-in-america-us-nones-nondenominational.html (accessed September 2, 2018).
Robinson, Hayley, Bruce MacDonald, Ngaire M. Kerse, and Elizabeth Broadbent. 2013. "The Psychosocial Effects of a Companion Robot: A Randomized Controlled Trial." *Journal of the American Medical Directors Association* 14(9):661–67.
Roozen, David A. 2008. "Faith Communities Today: American Congregations 2008." Hartford Institute for Religion Research. Faithcommunitiestoday.org (accessed September 14, 2018).
Rotman, David. 2017. "The Relentless Pace of Automation." *MIT Technology Review*. https://www.technologyreview.com/s/603465/the-relentless-pace-of-automation/ (accessed July 31, 2018).
Schwab, Klaus. 2017. *The Fourth Industrial Revolution*. New York, NY: Crown.
Shankland, Stephen. 2018. "An IBM Computer Debates Humans, and Wins, in a New, Nuanced Competition." C-Net. https://www.cnet.com/news/an-ibm-computer-debates-humans-and-wins-in-a-new-nuanced-competition/.
Softbank Robotics. 2018. "Robots: Who Is Pepper?" https://www.softbankrobotics.com/emea/en/robots/pepper (accessed August 20, 2018).
Stanford University. 2016. "One Hundred Year Study on Artificial Intelligence." https://ai100.stanford.edu (accessed February 15, 2018).
Susskind, Richard, and Daniel Susskind. 2015. *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. New York, NY: Oxford University Press.
Tegmark, Max. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence*. New York, NY: Alfred A. Knopf.
Universal Life Church. 2018. "Become an Ordained Minister." https://www.themonastery.org/landing/get-ordained (accessed September 19, 2018).
Utami, Dina et al. 2017. "Talk about Death: End of Life Planning with a Virtual Agent." In *Proceedings of the International Conference on Intelligent Virtual Agents*, 441–50. Cham. Switzerland: Springer International.
White House. 2016. *Artificial Intelligence, Automation, and the Economy*. Washington, DC: Executive Office of the President.