THE FUTURE OF AUGMENTED JOURNALISM

A guide for newsrooms in the age of smart machines

The Future of Augmented Journalism:
A guide for newsrooms in the age of smart machines

By Francesco Marconi, Alex Siegman and Machine Journalist

AUTHORS' NOTE

"The Future of Augmented Journalism" report is the result of collaboration among two humans and multiple artificial intelligence systems. Sections augmented by smart machines will be annotated with the } ai symbol.

The insights in this report stem from interviews with dozens of experts in the fields of journalism, technology, academia and entrepreneurship. The resulting insights allow us to explore the ways that artificial intelligence is and may be used to augment journalism.

We hope to present you with the basic tools and expertise necessary to make your own informed decision regarding the future implementation of AI in your newsroom.

Prologue

In the summer of 2013, News department leaders at The Associated Press made a daring suggestion to their business-side colleagues. They recommended working out a deal with a fledgling startup in the artificial intelligence space to automate the creation of certain news content. A few months later, a deal was struck with the Durham, North Carolina-based Automated Insights to automate the production of narrative text stories directly from data, first in sports and, soon thereafter, for corporate earnings reports.

So, yes, you read that right. The journalists at AP were the ones who first suggested handing over some work to a computer program — to the robots, if you will. Ironic perhaps, on its face, the move proved pivotal for the news agency in combating two mega-trends in the business — the relentless increase in news to be covered and the human constraints associated with covering it.

Streamlining workflows, taking out grunt work, crunching more data, digging out insights and generating additional outputs are just a few of the mega-wins that have resulted from putting smart machines to work in the service of AP's journalism since those first experiments a few years ago.

And AP is not alone. Innovators throughout the news industry are collaborating with technology companies and academic researchers to push the envelope in a number of related areas, affecting all points on the news value chain from news gathering to production and distribution.

As you'll see in the following pages, artificial intelligence can do much more than churn out straightforward sports briefs and corporate earnings stories. It can enable journalists to analyze data; identify patterns, trends and actionable insights from multiple sources; see things that the naked eye can't see; turn data and spoken words into text; text into audio and video; understand sentiment; analyze scenes for objects, faces, text or colors — and more.

Broadly speaking, AI promises to reap many big rewards for journalism in the years to come. Greater speed, accuracy, scale and diversity of coverage are just some of the results media organizations are already seeing.

Of course, this new wave of technological innovation is no different than any other that has come before it. Success still relies on how human journalists implement these new tools. Artificial intelligence is man-made, meaning that all the ethical, editorial and economic influences considered when producing traditional news content still apply in this new age of augmented journalism.

To best leverage and responsibly use artificial intelligence in news, the first step is to understand the technology itself. This report was designed as a primer for that purpose and to inspire more discussion of where the movement goes from here.

Streamlining workflows, taking out grunt work, crunching more data, digging out insights and generating additional outputs are just a few of the mega-wins that have resulted from putting smart machines to work in the service of AP's journalism.

The potential of AI in journalism

- Attend to menial tasks and free journalists to engage in more complex, qualitative reporting.
- Enhance communication and collaboration among journalists.
- Enable journalists to sift through large corpuses of data, text, images and videos.
- Help journalists better communicate and engage with their audience.
- Empower the creation of entirely new types of journalism.

What is augmented journalism?

First of all, there are some terms to get straight. In one, very basic sense, this is all about "automation" — reducing human effort and squeezing time out of the many chores journalists must undertake to get the story and get the news out to the public.

But this new wave of automation in the business now incorporates many forms of what is commonly referred to as "artificial intelligence," or if you want to get even more geeky than that, "cognitive technologies."

The upshot is that the machine is in the news garden, and it's doing some thinking on its own. And that thinking, or intelligent processing, is aiding both the inputs and the outputs of journalism.

For a concrete example of how AI can impact the input side of what journalists do, consider the International Consortium of Investigative Journalists (ICIJ), which for two years has directed a team of nearly 400 journalists analyzing approximately 2.6 terabytes of leaked emails, documents and databases. The result was The Panama Papers.

The ICIJ didn't employ any artificial intelligence at the start of its research, but Matthew Caruana Galizia, the organization's web applications developer, wishes it had.

"We were dealing with a vast amount of documents, and ICIJ just didn't have the resources to investigate them all," Galizia said. "But by using artificial intelligence, we would have been able to make that process much faster for all the journalists involved and end up with the same result."

Julien Martin of the ICIJ's data and research unit is confident that once his team can leverage AI, journalists will be able to do more than ever before by amplifying their existing skills. "We don't expect artificial intelligence to manufacture evidence," Martin said. "We are augmenting journalists' work."

In other words, the ICIJ is practicing augmented journalism. It is using artificial intelligence to empower its journalists to report at greater scale and with greater accuracy.

On the output end, a project at AP to automate the production of corporate earnings stories provides a seminal case study. Historically, the financial news staff of AP was saddled every three months with the enormous human task of reporting earnings for as many public companies as possible. As we will cover in more detail later, an automation program introduced three years ago enabled the agency to increase its output of corporate earnings stories by an order of magnitude each quarter, essentially providing coverage of the entire U.S. stock market.

As is the case with any new, impactful technology, mass adoption can be preceded by mass confusion and, more often than not, dissent.

Greek philosopher Socrates once protested any technology that allowed men to easily express their thoughts in writing, claiming, "The written word is the enemy of memory."

Executive summary

This summary was generated using Agolo NLP (natural language processing) technology. The algorithm took the entire text of this report as input, and organized and summarized its content. We will expand on the inner workings of these algorithms later in the report.

- AI can enable journalists to analyze data; identify patterns, trends and actionable insights from multiple sources; see things that the naked eye can't see; turn data and spoken words into text; text into audio and video; understand sentiment; analyze scenes for objects, faces, text or colors — and more.
- Three principal concerns associated with augmented journalism are the risks inherent in unchecked algorithmic news generation, the potential for workflow disruptions and the growing gap in skill sets required to manage this new specialty area.
- Dan Keyserling, head of communications at Jigsaw, a technology incubator created by Google, explains the overarching concern — that algorithms are prone to bias, just like humans.
- Not only is it imperative to save time and money in an era of shifting economics, but at the same time, you need to find ways to keep pace with the growing scale and scope of the news itself.
- With social networks playing such a big role today in the expansion of the news ecosystem, news organizations need to keep constant track of what's trending among news consumers in real time.

} ai

Centuries later, Johannes Trithemius, a prolific German cryptographer, exclaimed that the printing press would prove detrimental to the integrity of monks responsible for transcribing religious texts. The typewriter and even the modern word processor also had their skeptics.

Now, there are similar disputes with respect to augmented journalism.

What would Socrates say?

Three principal concerns associated with augmented journalism are the risks inherent in unchecked algorithmic news generation, the potential for workflow disruptions and the growing gap in skill sets required to manage this new specialty area.

Risks associated with unchecked algorithmic news generation

Imagine a team of environmental journalists is looking to track the correlation between oil drilling and deforestation. The team is unsure where in the world to focus the story, so it uses artificial intelligence to find land areas that have experienced rapid deforestation following nearby drilling activity.

Our hypothetical team begins by feeding its AI system a series of satellite images that it knows represent deforestation via oil drilling, as well as a series of satellite images that it knows do not represent deforestation via oil drilling. Using this training data, the machine should be able to view a novel satellite image and determine whether the land depicted is ultimately of any interest to the journalists.

The system reviews the training data and outputs a list of four locations the machine says are definitely representative of rapid deforestation caused by nearby drilling activity. But later, when the team actually visits each location in pursuit of the story, it finds that the deforestation was not caused by drilling. In one case, there was a fire; in another, a timber company was responsible.

It appears that when reviewing the training data, the system taught itself to determine whether an area with rapid deforestation was near a mountainous area, because every image the journalists used as training data had mountains in the photos. Oil drilling wasn't taken into consideration. Had the team known how its system was learning, it could have avoided such a mistake.

Dan Keyserling, head of communications at Jigsaw, a technology incubator created by Google, explains the overarching concern — that algorithms are prone to bias, just like humans.

"We need to treat numbers with the same kind of care that we would treat facts in a story," Keyserling said. "They need to be checked, they need to be qualified and their context needs to be understood."

"We need to treat numbers with the same kind of care that we would treat facts in a story. They need to be checked, they need to be qualified and their context needs to be understood."

DAN KEYSERLING
HEAD OF COMMUNICATIONS
JIGSAW

Tom Kent, a former standards and ethics editor at AP, even went so far as to publish a checklist for using robot journalism. At the top of that list is the question to be considered ahead of all others: "How accurate is the underlying data?"

Jigsaw tackles geopolitical problems using artificial intelligence and is forced to consider the implications of using AI to create journalistic output every day.

"Something that Jigsaw thinks a lot about is how do we understand and mitigate potential biases in the creation of machine-learning models," Keyserling said.

Just as it is important to verify a source's reliability, it is equally important to consider the reliability of an AI system.

Potential for workflow disruption

Implementing any new system will disrupt the traditional existing workflow of a newsroom. The changes aren't drastic, but they are important to keep in mind when implementing AI in a newsroom.

Meanwhile, every job is composed of unique facets, some of which can be more easily automated than others.

Consider the journalistic work involved in the Panama Papers. Sifting through millions of documents is tedious and can be easily automated, but no machine can look at a link between two organizations, think "that's odd," make the appropriate calls, speak with the right people and do the investigative work necessary to produce any substantive output.

This means that a reporter's daily duties may change in an AI-assisted newsroom. Reporters may spend less time transcribing and manually poring over datasets and instead spend that time making calls and pursuing leads derived from an AI analysis.

AP has already witnessed firsthand impact of using artificial intelligence to redefine journalists' roles in the newsroom.

The Associated Press began using algorithms to produce automated earnings reports in 2014 and estimates that doing so has freed up 20 percent of journalists' time, allowing those reporters to engage in more complex and qualitative work.

"Through automation, AP is providing customers with 12 times the corporate earnings stories as before (to over 3,700), including for a lot of very small companies that never received much attention," said Lisa Gibbs, AP's global business editor.

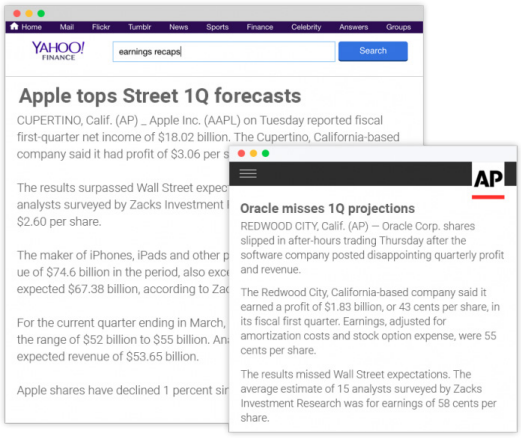

AP uses Automated Insight's "Wordsmith" to transform raw earnings data into 3,700 corporate earnings stories every quarter.
COURTESY AUTOMATED INSIGHTS

} ai tags generated through computer vision:
#text, #webpage, #font, #document

Where the chain of command may have once been:

Traditional workflow

- **REPORTER**
- **DESK EDITOR**
- **SECTION EDITOR**
- **COPY EDITOR**

It may now look more like this:

Adapted workflow

- **REPORTER**
- **REPORTER + AI**
- **DESK EDITOR + AI EDITOR**
- **SECTION EDITOR**
- **COPY EDITOR**

"With the freed-up time, AP journalists are able to engage with more user-generated content, develop multimedia reports, pursue investigative work and focus on more complex stories," Gibbs added.

Meanwhile, the automation routine itself requires maintenance and supervision. Story templates were built for the automated output by experienced AP editors. Special data feeds were designed by a third-party provider to feed the templates.

Continuing maintenance is required on these components as basic company information changes quarter to quarter, and though the stories are generated and published directly on AP's wires without human intervention, the journalists have to watch for any errors and correct them.

The bottom line in this case is that workflows changed for the better, but they changed significantly.

Growing gap in skill sets

So, then, what personal traits should an AI-assisted journalist cultivate?

First and foremost is a willingness to collaborate. Journalists who work well with data scientists and computational journalists are best positioned to thrive in an AI-assisted newsroom.

What is a data scientist? And what is a computational journalist?

Data scientists are individuals with the technical capabilities to implement the artificial intelligence systems necessary to augment journalism. They are principally scientists, but they have an understanding as to what makes a good story and what makes good journalism, and they know how to communicate well with journalists.

"It's important to bring science into newsrooms because the standards of good science — transparency and reproducibility — fit right at home in journalism," said Larry Fenn, a trained mathematician working as a journalist on AP's data team.

Computational journalists are principally journalists, but they have an understanding as to how artificial intelligence works and how AI can be used to augment their own journalism. Perhaps most importantly, they know how to communicate well with data scientists.

"We've put a lot of effort into putting more journalists who have programming skills in the newsrooms," said The New York Times' chief technical officer, Nick Rockwell.

Hearst, one of the nation's largest diversified media, information and services companies, recently announced the creation of the Native and Emerging Technologies unit. It will be responsible for launching data-driven products in categories that the operation believes will impact its various businesses lines in the future, including artificial intelligence, voice interfaces and augmented and virtual reality.

Using AI to analyze the writing in this report

With the help of IBM's Watson, we have determined that the earlier "What is augmented journalism" section was very analytical and open. The values in the "Language Style" and "Social Tendencies" sections work by assigning a score to the perceived writing style in the text. According to the algorithm, the section analyzed is very likely to be "perceived as intellectual, rational or systematic" and very likely to be "perceived as intellectually, curious, emotionally-aware, imaginative or open to change."

} ai

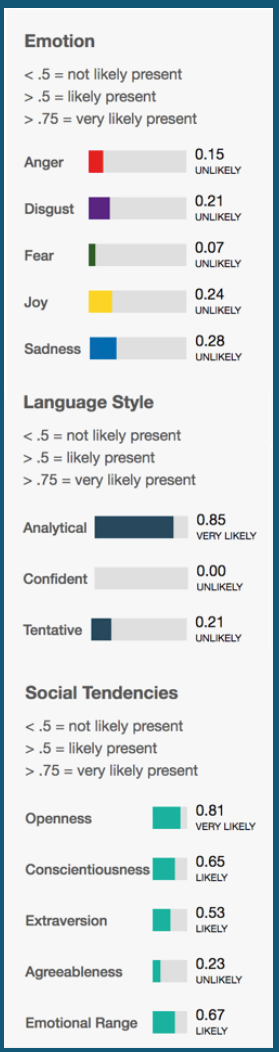

IMAGE COURTESY IBM WATSON

Chris Papaleo, executive director of emerging technology at Hearst, believes smart machines will come to drive change at a foundational level:

"Hearst sees AI as the common thread that will help us execute on this next wave of personal computing trends. Even beyond that, the skills and expertise we develop by leaning into these projects can be leveraged in the rest of the business, not just in an R&D lab. Can we develop models to help improve business operations like talent acquisition and retention, managing budgets, etc.? That's the transformative wave that we, as an industry, are embarking on.

"We are actively encouraging the heads of all the divisions to keep the principles of AI in mind," Papaleo added. "At this early stage of the technology, the best way to build expertise and future capabilities is to devote time and energy into training and experimentation."

How much science, math and computer coding does that traditional journalist need to understand so as to best interact with these data scientists and computational journalists? Not much, actually.

In fact, it's rather easy for everyone in an AI-assisted newsroom to collaborate. Journalists, data scientists and computational journalists are all storytellers. They all ultimately transmit information to readers in an engaging and digestible format. The only true difference between them is the medium they use to tell those stories.

"At this early stage of the technology, the best way to build expertise and future capabilities is to devote time and energy into training and experimentation."

CHRIS PAPALEO
EXECUTIVE DIRECTOR OF EMERGING TECHNOLOGY
HEARST

Automated summary, part one

- This new wave of automation in the business now incorporates many forms of what is commonly referred to as "artificial intelligence," or if you want to get even more geeky than that, "cognitive technologies."
- Julien Martin of the International Consortium of Investigative Journalists is confident that once his team can leverage AI, journalists will be able to do more than ever before by amplifying their existing skills.
- The Associated Press began using algorithms to produce automated earnings reports in 2014 and estimates that doing so has freed up 20 percent of journalists' time, allowing those reporters to engage in more complex and qualitative work.

} ai

What are the relevant technologies in this field?

Each of AI's unique subdomains offers an opportunity for journalists to augment their work. This report will cover those fields that are poised to be most relevant to journalism in the near future.

Machine learning

At the heart of several artificial intelligence systems is machine learning — a subdomain of AI that describes how a system can self-adjust without being told what adjustments to make.

Machine learning relies, in turn, on a process called deep learning, a method of breaking down a complex idea into a series of smaller, more approachable tasks that ultimately lead to a designated endpoint.

But for machines to learn, they need to be taught, as Columbia University data scientist Amir Imani explains:

"When children are growing up and learning how to perceive and understand their parents' facial expressions, they do so by constantly being exposed to them — seeing their faces and recognizing their tone.

"Just imagine that algorithms are doing the same thing, but instead of taking three years to learn how a person's facial expression is communicating a certain emotion, machine-learning algorithms go through thousands of images at once and try to categorize them."

Imani used this sort of machine learning alongside Quartz journalist Sarah Slobin in January 2017 to augment traditional reporting on the nature of U.S. President Donald Trump's inaugural speech, analyzing his facial expressions and determining the emotions expressed throughout his acceptance and inauguration speeches.

Imani's algorithm was trained to identify micro-expressions (e.g. raised eyebrows as an indicator of surprise) and subsequently equate those micro-expressions with pre-defined emotions.

Robert Farr, a technology director at AP, elaborated on the approach of deriving emotion-recognition through machine learning.

"To analyze for emotion, these machine-learning systems need to be able to first accurately detect a face and then landmark multiple points on each face. Based on these landmark points, it is possible to calculate the likelihood of certain expressions, like a smile or frown."

Farr added that while the AI is compelling, it remains flawed. For instance, the system is dependent on the orientation of a subject's face — head-on shots (like red-carpet or executive profiles) are preferable — and dependent on an unobstructed view of the subject.

Journalistically relevant subdomains of artificial intelligence

- **MACHINE LEARNING**
  - Supervised
  - Unsupervised
- **NATURAL LANGUAGE**
  - Generation
  - Processing
- **SPEECH**
  - Speech to text
  - Text to speech
- **VISION**
  - Image recognition
  - Computer vision
- **ROBOTICS**

How does AI feel about itself?

We used Microsoft Cognitive Services programming tools to detect the sentiment of all mentions of artificial intelligence in this report. The values below represent the level of certainty that the algorithm is detecting a given emotion. Values closer to "0" represent "maybe" while values closer to "1" represent "definitely."

sadness: 0.31443
joy: 0.204334
anger: 0.185282
disgust: 0.161522
fear: 0.109779

} ai

While no system is perfect, AI can nonetheless provide an additional perspective to a traditional news story, and as AI technologies continue to improve, the possible perspectives we can apply to stories will increase.

Possible perspectives include noting the dominant emotions of a CEO during an earnings call, tracking the tone of a politician over time when he or she speaks about a certain policy or topic, augmenting sports reporting with real-time biometrics and even automatically generating a soundtrack.

"Creative AI is something that enhances the lives and productivity of existing creators, whether they are journalists or musicians," said Drew Silverstein, co-founder of Amper Music, an AI-powered music composer.

In a recent project, AP's election development team built its own machine-learning algorithm that can assist in determining the probability of a political race outcome. The achievement is notable because vote counting requires extensive manual labor, and by augmenting the calling process with machine learning, the news agency can streamline a process where speed and accuracy are paramount. Expert human race callers now have better information at their fingertips faster.

"AP tested this approach extensively during the state primaries in 2016, and in the past two months, we have re-run elections in 10 states from last November. The system's recommended calls have held up in 100 percent of the races analyzed in those 10 states," said David Pace, a news editor at AP.

Elsewhere in the industry, The New York Times has used machine learning to enhance its campaign finance coverage.

"The New York Times has developed a lot of work with campaign finance data sets, which are typically hard to navigate and understand, often by design," said the Times' Rockwell. "That's the kind of problem that machine learning makes tractable. These algorithms can recognize patterns in data that may have even been deliberately obfuscated."

But what about more basic tasks or newsrooms that don't have their own CTOs or IT departments to fall back on?

Fortunately, you don't need to build your own unique machine-learning system to leverage the power of artificial intelligence.

For example, AP partnered with Graphiq, a company that uses more than 250 billion data points from sports, politics, weather and other sources to automatically generate clear and concise visualizations that allow the news agency to offer interactive graphics for dozens of stories each day.

"Most media outlets are eager to take advantage of AI," said Alex Rosenberg, vice president at Graphiq. "In some cases, those organizations may think, 'Oh, let's quickly spin up a team to build a tool or utilize custom software to execute an AI project.'

"And at least today, I believe that thinking is somewhat misguided. In fact, there are vendors such as Graphiq that are privatizing aspects of AI, making

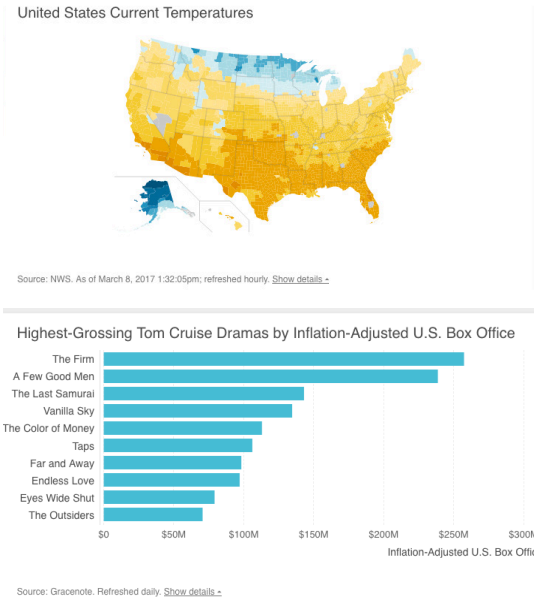

Graphiq uses AI to automate infographic generation and match its output to news articles in real time. IMAGES COURTESY GRAPHIQ

} ai tags generated through computer vision: #diagram, #illustration

Fortunately, you don't need to build your own unique machine-learning system to leverage the power of artificial intelligence.

it dramatically easier for anyone within a media organization to rapidly learn and utilize these tools."

Justin Pang, head of publishing partnerships at Google, elaborated on the possibility for collaboration between media outlets and AI system providers:

"It is estimated that media companies and user-generated content creates over 2 billion digital images and over 1 billion hours of video watch time every day. This explosion of unstructured data needs to be stored, analyzed and utilized and is an increasingly important area where media companies will closely partner with technology companies like Google that have supportive infrastructure and expertise in areas like machine learning."

Digging in a bit deeper into the detail, there are two general types of machine learning that companies like Google, Microsoft, IBM and Amazon can provide as a cloud service: supervised and unsupervised.

Supervised learning

Imagine a team of financial journalists suspect a local business is engaging in fraudulent activities, but they can't afford to allocate substantial manpower to investigate until they are certain. How might they use machine learning to aid their reporting?

The team could leverage supervised learning, a subset of machine learning, to easily flag any financial documents belonging to the suspected business that might suggest illegitimate operations.

This supervised learning system takes a set of labeled examples as its input and classifies new information as its output. In other words, the team feeds the machine financial documents belonging to businesses that either have or have not been convicted of illegal financial operations, and the system learns to differentiate between the two.

Then, when the team provides the system with novel information, which in this case are the documents belonging to the suspected business, it can classify documents as suspicious or not.

The same method could have been used by a sports journalism team looking to flag whether an athlete is on steroids (using the statistics of athletes convicted of using steroids as training data) or even by a team investigating a small town it suspects is home to illegal mining practices (using satellite images of areas that are known to be home to illegal mines as training data).

The accuracy of these supervised learning systems is, of course, important. The two most common errors in this sort of machine learning are terms that we borrow from statisticians —Type I (false negative) and Type II (false positive) errors.

A false negative would mean that the system labels a nefarious business as law-abiding, whereas a false positive means the system labels a law-abiding business as nefarious.

The more data a team has to teach its system, and the more accurate that data, the less likely the machine is to commit any errors.

The more data a team has to teach its system, and the more accurate that data, the less likely the machine is to commit any errors. No machine-learning algorithm is going to be 100 percent accurate.

Unsupervised learning

Unlike supervised learning, unsupervised learning takes an unstructured dataset as its input and is given no target output. The system has free rein to derive relationships between input and output.

Imagine a team of sports reporters is looking for possible storylines from the weekend's baseball games. It might instruct its system to take all of the weekend's box score data as the input and simply discover possible patterns.

That system may find that teams with managers whose last name begins with a vowel are more likely to lose morning games or that the further north a winter game is played, the lower scoring that game will be. Whatever it is, the computer isn't looking for anything in particular, but it may surface something worth investigating.

The data journalism team at AP, for example, used unsupervised machine learning to find hidden patterns in a data set of 140,000 human-entered incidents documented by the Gun Violence Archive (GVA). After discovering a host of errors in the initial data, AP used unsupervised machine learning to simplify the data and flag certain entries for further review without specific guidance.

AP was able to surface new story ideas that otherwise could have been missed. Journalists were able to find the "most typical" incidents related to gun violence, including whether the shooting was accidental or not, if it involved a child or even if a police officer was involved.

MACHINE LEARNING TIPS

- Machine learning is helpful for drawing conclusions from a large corpus of data.
- Use supervised learning to find the link between a known input and a known output.
- Use unsupervised learning if you're unsure of what you're looking for regarding a correlation between input and output.

Language

Generating language automatically is a key output from artificial-intelligence systems. Two of the most relevant fields of language processing already impacting journalism are Natural Language Generation (NLG) and Natural Language Processing (NLP).

"By pairing humans and software together, you can produce something that's much better than either one can do on its own," said Robbie Allen, founder of the Automated Insights, the company behind AP's data-to-text automation projects for sports and financial stories.

NLG

Natural language generation works similarly to the word game Mad Libs — it turns structured, templated data into a digestible written narrative.

Perhaps the most basic example of NLG is Quakebot, a tool developed by the Los Angeles Times to automatically generate a templated report of an earthquake within moments of the event itself.

"By pairing humans and software together, you can produce something that's much better than either one can do on its own."

ROBBIE ALLEN
CEO AND FOUNDER
AUTOMATED INSIGHTS

Consider this tweet from Quakebot at 1:33 a.m. on Feb. 22, 2017:

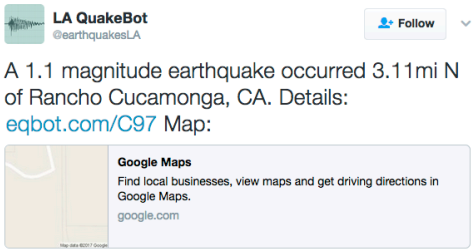

The template reveals itself immediately: "A ____ magnitude earthquake occurred ____ mi ___ of _____, ____."

Quakebot pulls the appropriate data from the U.S. Geological Survey to fill in those blanks, and it can tweet a simple story about an event moments after it happens.

The more complex stories being generated by AP follow the same basic formula. National language generation "enables somebody to tell a story around their data and do it at a scale that wasn't possible before," said Automated Insights' Allen.

Following the successful launch of automated earnings stories in financial news, AP partnered with Automated Insights again in 2016 to use natural language generation to produce automated coverage of Minor League Baseball games. In this case, consistent coverage of the lower levels of professional baseball was too sprawling to manage with human reporters. But based on the availability of reliable data from the games, the production of simple game recaps could be automated at an unprecedented scale, opening up brand new territory for AP.

NLP

Unlike natural language generation, natural language processing understands and contextualizes text. In other words, NLG writes and NLP reads.

Consider again the case of The Panama Papers. The ICIJ had to sift through millions of documents, each of which containing series of named entities (names of people, companies, institutions and places).

Natural language processing can help organizations such as the ICIJ determine the links that connect those named entities. For instance, if Mr. Smith sold a house to Mrs. Jones, NLP could recognize that Mr. Smith and Mrs. Jones are people, and that they are related by an action — the former sold a house to the latter. It may seem trivial, but it would take a human researcher years to decipher and plot the millions of relationships between individuals included in all the documents.

NLP is also common in the world of translation services, including a test program underway at AP.

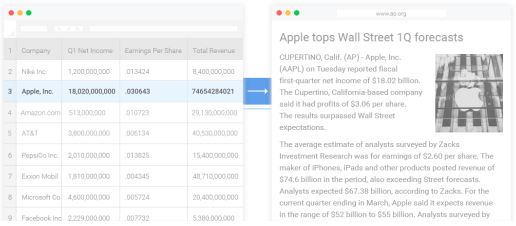

Natural language generation helps turn data sets into news content.
IMAGE COURTESY AUTOMATED INSIGHTS

} ai tags generated through computer vision: #text, #font, #webpage, #document

How does AI sum itself up?

Using Microsoft Cognitive Services once more, we leveraged natural language processing to analyze the text in this report. The algorithms can understand high-level concepts throughout the report and organize those concepts by a relevance score. You'll notice that some terms such as "Artificial neural network" and "Artificial intelligence systems integration," though not directly referenced in this document, were highlighted by the AI system as being relevant. This type of analysis is called concept expansion.

Machine learning: 0.970687
Artificial intelligence: 0.947929
Journalism: 0.670752
Artificial neural network: 0.636766
Natural language processing: 0.549341
AI systems integration: 0.423854
Unsupervised learning: 0.416939
Computer vision: 0.39715

} ai

"The goal is not to replace AP's current team of human translators but to increase their production so it can better serve news organizations in the Middle East, Latin America and the Caribbean," said AP's Farr. "Output will still need the manual editor without a doubt, but it may be useful as a baseline and therefore an efficiency that might enable a rise in volume."

"Researchers in NLP are teaching computers to understand the complexities of language," said Mark Hansen, executive director at The Brown Institute for Media Innovation. "That might start with questions of syntax or grammar — having a computer recognize the structure of a sentence, break it into words and label each word with its part of speech, for example.

"It also includes algorithms that get at the semantics of text, or what it means — having the computer highlight important 'entities,' like people and places, and identify the topics being discussed," Hansen added.

NLP is widely used in summarization technologies like those under development by Agolo, an AI company launched by Columbia University graduates and funded by Microsoft. (Note: To illustrate the effect of summarization, Agolo's technology was used to create the summaries that appear throughout this report, based on the analysis of the full text.)

"Agolo's NLP technology doesn't start by looking at structured data sets such as financial spreadsheets," Agolo's chief technology officer, Mohamed AlTantawy, said. "The algorithms analyze unstructured data, including news articles, reports and press releases, and turn it into an unstructured output. In our case, we generate summaries from all of those documents."

Agolo also leverages natural language processing to generate news summaries for chatbots and produce summaries from news articles that can be posted directly to social media.

LANGUAGE TIPS

- Natural language can help journalists work with data in an efficient, user-friendly manner.
- Just as you do with non-automated content, be sure you have a defined process for how to edit your automated stories.
- Test your system as many times as possible in a variety of situations and with a wide range of data.

Speech

A lot of discussion in journalism involves conversational interfaces and their impact for news consumption and delivery. Personal assistant bots such as Amazon Echo, Apple's Siri and Google Home depend on text to speech, while other tools such as automated transcription rely on speech-to-text functions.

Text to speech

Several news publishers including AP, The Wall Street Journal, NPR, BBC News, and The Economist have already started experimenting with audio interfaces.

"It's a very exciting time for voice assistants and on-demand audio news," said Tom Januszewski, a director of business development at AP. "The Amazon Echo, for example, has surprised everyone with its broad acceptance, and the large number of news organizations already participating on the platform. It shows these news organizations recognize the importance of the technology and that they want to get in early."

"There are some grammar guidelines to follow that just sound better when a smart speaker is reading something aloud."

KEN ROMANO
TEXT AND MULTIMEDIA PRODUCTS DIRECTOR
ASSOCIATED PRESS

This system works by analyzing text stories the newsroom provides and leveraging text-to-speech technology to broadcast the news content using a synthetic voice.

The underlying text also needs to be written in such a way that allows the machine to read it, and that's not always easy.

Ken Romano, a text and multimedia products director at AP, said certain practices such as avoiding direct quotations and including a person's title before his or her name, instead of after, result in a more accurate text-to-speech translation.

"What we've found is that content written like our stories for radio DJs result in better translations," he said. "There are some grammar guidelines to follow that just sound better when a smart speaker is reading something aloud."

Speech to text

How can a machine understand the spoken word? That's where speech-to-text technology comes into play.

Speech to text takes vocal patterns as input and matches them to an output. It derives meaning from sound, just like people do when they listen to one another. For all intents and purposes, the machine turns spoken word into text.

This technology also allows journalists to reduce the quantity of mundane tasks they must complete every day.

For instance, according to a recent survey of over 100 journalists in the U.S. conducted by the Reynolds Journalism Institute (RJI), reporters spend an average of three hours per week interviewing subjects and double that time transcribing notes from audio recordings of their interviews. With speech-to-text technology, an AI system could take over.

Recordly, a speech-to-text-powered transcription app developed at the RJI's Futures Lab, hopes to do exactly that.

"Recordly was born out of our own need to make the reporting and writing process more efficient," said Sintia Radu, a project lead. "We see transcribing as time-consuming, and we believe that beginning any storytelling process with this annoying activity is counterproductive. The writer is already fatigued when he or she needs to put the story together."

Meanwhile, AutoEdit, an open-source editor developed by Knight-Mozilla fellow Pietro Passarelli while at Vox Media, uses speech to text to automate routine post-production actions like captioning videos or selecting meaningful quotes from a video to attract readers.

Speech-to-text technology is being used as a reporting tool to transcribe interviews and conversations in real time.
IMAGE COURTESY RECORDLY

} ai tags generated through computer vision: #finger, #closeup, #hand, #arm, #watch

SPEECH TIPS

- Speech systems can make the most complex software user-friendly.
- For text-to-speech systems, such as voiceover technology, consider how to implement pronunciation hints and guidance for unusual names or phrases.
- Look to use speech-to-text systems to automate time-consuming, routine tasks such as transcription and captioning.

Vision

Beyond listening and transcribing, other smart systems being developed can record what the eye can see, or in some cases, even more than what the human eye can see in visual content. What's called "computer vision" can help editors quickly classify and organize vast corpuses of images and video, speed up the editing process and enable journalists to source evidence for investigative pieces.

In one high-profile case, AP used satellite imagery from a company called Digital Globe to secure high-resolution images of sea vessels in Southeast Asia to document critical evidence for an investigative project on abuses in the seafood industry that won the Pulitzer Prize for Public Service in 2016.

Digital Globe's computer-vision algorithms worked to reorient its satellite-based cameras to shoot the optimal and necessary images that ultimately provided a point of reference in the investigation beyond the reach of AP's reporting team.

This sort of image recognition works using a type of machine learning called "neural networks," according to Matt Zeiler, the founder of an image- and video-recognition platform called Clarifai. These neural-network algorithms work by mimicking the way humans are understood to perceive images.

"Nobody really knows exactly how the brain operates, but we do know the visual cortex processes the input to our eyes in multiple layers," Zeiler said. "Our eyes then group regions of images together and keep only the positive elements. Computer-vision algorithms use mathematical models to replicate this process through their own layering of images."

Aside from investigative work, these types of systems can be employed for much more mundane operational tasks. Clarifai, for example, uses vision technology to pull descriptive information out of photos and videos automatically, such as location and identifiable people, places and things. This can speed up the creation of valuable metadata that makes imagery easier to manage in the newsroom or to be discovered via search. Vidrovr is another young company in this space, and Google offers similar capabilities as part of Cloud Vision API, which automatically generated the image tags in this report.

In another production scenario, Wibbitz, a text-to-video platform, utilizes image recognition to create videos from compiled images and video footage that can match a given text automatically, producing rough-cut videos in seconds for further refinement by human editors.

"It's all about collaboration between humans and technology, allowing journalists to focus more on the content and less on the repetitive parts of creating a video," says Zohar Dayan, founder and CEO of Wibbitz.

"Our platform is enabling more individuals to work faster, better and in a more efficient way. It's also helping producers to create engaging videos at scale."

Ben Nadler, a product manager at AP, says image recognition and computer vision have already helped AP photographers and editors with both newsgathering and dissemination.

Computer vision can help editors quickly classify and organize vast corpus of images and video, speed up the editing process and enable journalists to source evidence for investigative pieces.

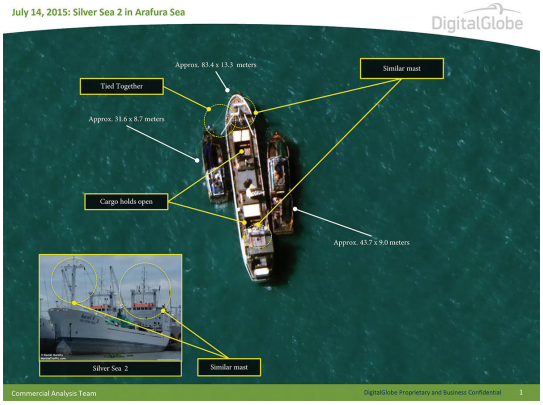

Satellite image provided by DigitalGlobe showing two fishing trawlers load slave-caught fish onto a refrigerated cargo ship.

} ai tags generated through computer vision:
#vehicle #aircraft

"Computer vision can process large numbers of photos very quickly by discarding outtakes, selecting the best images to go with a text story and adding tags and even captions," he said. "It fundamentally changes the way that photos and videos are used by making them much more searchable."

Another example of using computer vision to augment news is the possibility for analyzing real-life events. Imagine a hypothetical protest — there may be tear gas, mass confusion and excessive noise, all of which are likely to disorient a journalist. But if a journalist can capture video of the event from a safe distance, computer vision can analyze footage and help sense the physical movement of the crowd or even the sentiment.

Robotics

What about the hardware that enables journalists to gather the data that those AI systems rely on?

Robotic cameras and drones can provide journalists with unprecedented viewing access, while robotic sensors can measure human response to live events and measure natural occurrences — like the ones used to inform QuakeBot.

At the 2016 Summer Olympics, AP used a team of 11 robotic and 16 remote cameras that allowed photographers to place cameras in physically inaccessible areas as well as rapidly adjust the equipment to preset filming functions.

AP also used drones to report on the Middle East refugee crisis, capturing aerial photos of displaced Iraqis near Dibaga, southeast of Mosul.

And according to Sam McNeil, an AP video journalist based in the Middle East, drones can do more than just provide supporting imagery to journalism.

"Drones aren't just for simple aerial video or photographs, but can be used to gather scientific data about climatology and other fields to inform deeper journalism," McNeil said.

Robotics, when combined with other AI technologies such as machine learning, natural language, speech and vision, can enhance the utility of the hardware.

Earth TV, a German-based satellite company that airs live video broadcasts from around the world, operates a network of thousands of autonomous cameras to capture real-time events, from sunsets and skylines to traffic jams and festivals.

Earth TV also uses AI systems and machine learning to determine when to capture video and where to point its cameras to get the highest quality image — a task that, without AI and robotics, would be too labor-intensive to accomplish. For example, a camera may adjust its position if the air quality is particularly clouded by pollution or adjust its brightness settings if its subject is cast in shadows.

VISION TIPS

- AI vision can help speed up creative and editing processes in a newsroom.
- An image-recognition system is only as accurate as the data used to train it. Use a wide variety of images and, if the content is subjective, use several individuals of different backgrounds to provide the tagging data or annotations necessary to train the system.
- Like all other areas of AI, adjust your daily workflow as you integrate these systems.

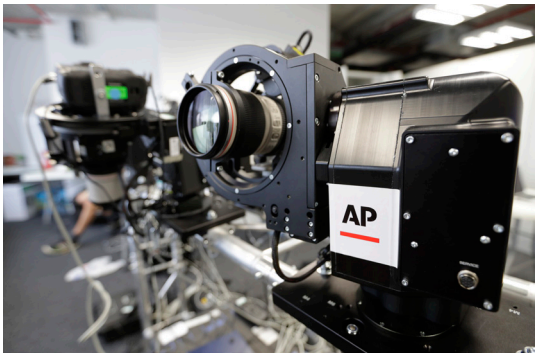

A robotic camera used by The Associated Press to capture unique images from angles not normally seen by the public.
AP PHOTO

} ai tags generated through computer vision: #cameraoperator, #machine, #singleLensReflexCamera, #filmmaking

"We invented an entirely automated smart production system with the goal to release the human operator from the stress of control and to channel the human impact to creatively teaching the robotic systems to control and improve themselves," said Earth TV CEO Nikolaus Lohmann.

Robots are also re-programmable — to change the robot's behavior, all that needs to be done is write a new program to its controlling computer.

AP used Raspberry Pis — a type of low-cost, small, easily programmable computer — to build an experimental prototype capable of measuring the ground vibration and noise level from entertainment and political venues as an indirect way of determining the most popular songs at a concert, the biggest plays of a game or even the most powerful quotes during campaign rallies.

"AP is investigating how smart hardware, commonly referred as the internet of things, can help journalists with their newsgathering needs," explained Vince Tripodi, AP's vice president of research and development.

ROBOTICS TIPS

- Drones can prove useful to journalists, but there are extensive restrictions against them. Make sure you know what, when and where you can fly.
- Robotics can work alongside machine learning systems to automate traditionally human-completed tasks.
- Just like AI, there are plenty of companies that rent or sell their services to smaller outlets and even freelancers.

Automated summary, part two

- Quartz used machine learning to augment traditional reporting on the nature of U.S. President Donald Trump's inaugural speech, analyzing his facial expressions and determining the emotions expressed throughout his acceptance and inauguration speeches.
- There are vendors such as Graphiq that are privatizing aspects of AI, making it dramatically easier for anyone within a media organization to rapidly learn and utilize these tools.
- Supervised learning systems take a set of labeled examples as its input and classifies new information as its output.

} ai

What will be the lasting impacts of AI on the news?

As this report has shown, these technologies are opening up new territory and changing journalism in ways no one might have predicted even a few years ago. And they arrive at a time when journalists and media companies are searching for new solutions to the challenges that the digital revolution has imposed on the news business.

Not only is it imperative to save time and money in an era of shifting economics, but at the same time, you need to find ways to keep pace with the growing scale and scope of the news itself.

With social networks playing such a big role today in the expansion of the news ecosystem, news organizations need to keep constant track of what's trending among news consumers in real time. Specialized monitoring systems, such as those provided by the startup NewsWhip, are providing news organizations with 21st-century radar systems that are helping them spot more trends and respond more rapidly in their coverage decisions.

"AP's newsroom uses tools like NewsWhip to get a better grasp for the stories that are capturing and holding people's attention," said Eric Carvin, AP's social media editor.

In addition to scale, scope and speed, the technologies are also moving the needle on accuracy. In the case of the automated earnings project at AP, the error rate in the copy decreased even as the volume of the output increased more than tenfold.

"The reason for the lower error rate is that algorithms don't make typos or arithmetic miscalculations," said AP's global business editor, Lisa Gibbs. "The errors are generally because of a problem with the data. If the data's bad, you get a bad story."

That's a conclusion that really distills the story of AI in journalism to its essence — "garbage in" still produces "garbage out." The new era of augmented journalism won't fulfill its apparent promise without the diligent management and care of the journalists learning how to put these new technologies into practice.

Justin Myers, appointed as the AP's first automation editor in response to the growing importance of machine-driven projects at the agency, summed it all up very plainly:

"Quality depends on two main things: explanation and testing. Editorial stakeholders must be able to explain what they need these systems to do, as well as flag problems they observe or expect. Technical implementers must be able to explain how their systems produced particular results, both to support editorial transparency and to resolve errors. And both groups need to be comfortable frequently testing these systems — and their assumptions — before and after launch."

NewsWhip developed a tool capable of tracking the most-discussed news topics by applying machine learning to social media data.
IMAGE COURTESY NEWSWHIP

} ai tags generated through computer vision: #screenshot, #machine

"The reason for the lower error rate is that algorithms don't make typos or arithmetic miscalculations. The errors are generally because of a problem with the data. If the data's bad, you get a bad story."

LISA GIBBS
GLOBAL BUSINESS EDITOR
ASSOCIATED PRESS

Justin Hendrix, executive director at NYC Media Lab, an organization that promotes collaboration between digital media and technology companies with universities, says AI ultimately represents a fundamental change in the way media operates:

"From text to images to video to audio, media is now completely machine-readable. Applications of data science, from machine and deep learning to artificial intelligence, are sweeping over the industry and changing the way media is produced, distributed, consumed and monetized. From automation to augmentation, media companies that embrace machines do so to find economic advantage. Efficiencies and optimization are possible in every area of the business, from content production to monetization. Increasingly, automation will change the balance sheets for media companies."

Hendrix also highlights the concerns associated with such changes:

"Where there is opportunity, there are also challenges. Chief among them are the organizational challenges of drawing in new talent that can advance and apply these technologies. There are also significant externalities to consider — from 'fake news' to filter bubbles, from privacy concerns to cybersecurity, there are profound challenges for individuals, companies and society in a machine-driven future."

"From automation to augmentation, media companies that embrace machines do so to find economic advantage."

JUSTIN HENDRIX
EXECUTIVE DIRECTOR
NYC MEDIA LAB

Conclusion

In the 1960s, historian Leo Marx captured the disruptive impacts of technology on pastoral America in a book entitled "The Machine in the Garden." In that masterwork, which has been required reading for generations of American history students over the past 50 years, Marx described the effects of technological disruption in the industrial age as more psychic than physical in many respects.

And so it is with AI, the machine in journalism's garden today. The impact of this technology will ultimately be a story about how the human journalists adapt, not how the machines work.

In that spirit, this report concludes with some parting thoughts about what to watch — and watch out — for as this story unfolds further.

The impact of this technology will ultimately be a story about how the human journalists adapt, not how the machines work.

Technology changes, journalism doesn't

Artificial intelligence can help augment journalism, but it will never replace journalism. AI might aid in the reporting process, but journalists will always need to put the pieces together and construct a digestible, creative narrative.

AI is susceptible to the same biases and errors as humans

Artificial intelligence is designed by humans, and humans make mistakes. Therefore, AI can make mistakes. Furthermore, an AI system is only as good as the data that goes into it.

AI is not a silver bullet

Artificial intelligence can't solve every problem. As the technology evolves, it will certainly allow for more precise analysis, but there will always be challenges the technology can't overcome.

Journalists can best leverage AI once they understand the technology

Artificial intelligence is complicated, and there are many ways it can be implemented in a newsroom, but just like any other technology, the more you know about a tool, the more effectively you can use it.

There are ethical considerations inherent in journalism's use of AI

Again, just because the tools of journalism change, that doesn't mean the rules of journalism change. As AI works its way into newsrooms, it is important to adhere to our existing standards and ethics.

Automated summary, part three

- AI technologies are opening up new territory and changing journalism in ways no one might have predicted even a few years ago.
- They arrive at a time when journalists and media companies are searching for new solutions to the challenges that the digital revolution has imposed on the news business.
- Not only is it imperative to save time and money in an era of shifting economics, but at the same time, you need to find ways to keep pace with the growing scale and scope of the news itself.
- With social networks playing such a big role today in the expansion of the news ecosystem, news organizations need to keep constant track of what's trending among news consumers in real time.

} ai

About the authors

Francesco Marconi is the manager of strategic planning and development at The Associated Press and an innovation fellow at The Tow Center for Digital Journalism at Columbia University. Follow him on Twitter at @fpmarconi.

Alex Siegman is a computational journalist and a master's candidate at Columbia University's Graduate School of Journalism. Follow him on Twitter at @alex_siegman.

"Machine Journalist" ({ ai }) is an amalgamation of various artificial intelligence systems.

THE FOLLOWING ASSOCIATED PRESS STAFF CONTRIBUTED TO THIS REPORT:

Mike Bowser, director of creative services
Eric Carvin, social media editor
Robert Farr, technology research director
Larry Fenn, data and web development journalist
Lisa Gibbs, global business editor
Tom Januszewski, director of business development
Jim Kennedy, senior vice president of strategy and enterprise development
Jake Kreinberg, marketing manager
Sam McNeil, video journalist
Justin Myers, automation editor
Ben Nadler, text and multimedia products manager
David Pace, news editor
Ken Romano, text and multimedia products director
Paul Shanley, director of international development and partnerships
Troy Thibodeaux, interactive newsroom data editor
Vincent Tripodi, vice president of research and development

THANK YOU TO THESE EXTERNAL EXPERTS WHO SHARED THEIR INSIGHTS:

Drew Silverstein (Amper Music)
Mohamed AlTantawy (Agolo)
Sage Wohns (Agolo)
Robbie Allen (Automated Insights)
Alex Moore (Boomerang)
Matt Zeiler (Clarifai)
Amir Imani (Columbia Data Institute)
Stephen Lepitak (The Drum)
Sebastian Huempfer (Echobox)
Chris Papaleo (Hearst)
Justin Pang (Google)
Julien Martin (ICIJ)
Matthew Caruana Galizia (ICIJ)
Nick Rockwell (New York Times)
Justin Hendrix (NYC Media Lab)
AJ Nelson (One Spot)
Jack Clark (OpenAI)
Sarah Sobin (Quartz)
John Keefe (Quartz)
Dan Keyserling (Jigsaw)
Zohar Dayan (Wibbitz)
Sintia Radu (Recordly)
Alex Rosenberg (Graphiq)
Joe Ellis (Vidrovr)
Mark Hansen (The Brown Institute for Media Innovation)
George King (The Tow Center for Digital Journalism)
Nikolaus Lohmann (Earth TV)
Russell Stevens (MIT Media Lab)
Paul Quigley (Newswhip)

AI TECHNOLOGY AND TOOLS MENTIONED IN THIS REPORT

Agolo is advanced summarization software that delivers real-time summaries of documents.

Auto Edit is an open-source editor that uses speech-to-text technology to automate video transcription.

Amazon Machine Learning helps developers leverage machine-learning technology without having to learn complex data science.

Amper Music enables users to create unique, professional-quality music through artificial intelligence, all with the click of a button.

Automated Insights provides real-time content automation services that transform data into digestible narratives, visualizations and applications.

Clarifai, an AI-based startup, uses sophisticated machine learning to classify images and videos, as well as to develop unique tools for a variety of organizations.

Earth TV is a satellite company that airs live video broadcasts from around the world and operates a network of thousands of autonomous cameras to capture real-time events.

Google Cloud Platform enables developers to build machine-learning systems, including computer vision, natural language, speech and translation.

Graphiq, a data aggregation and visualization company, transforms complicated data into contextually rich presentations and graphics.

IBM Watson provides some of the cognitive computing tools and API's necessary to build applications capable of understanding different forms of data, as well as those that can learn, reason and interact with humans.

Microsoft Cognitive Services is a digital toolkit that allows developers to build powerful applications that span the realms of vision, speech and language.

NewsWhip tracks and predicts the stories, events and individuals drawing attention from the public on social networks.

Raspberry Pi is a low-cost, small programmable computer.

Vidrovr is a computer vision company that automatically analyzes large video collections to and tags content so as to make it easily searchable.

Wibbitz is a text-to-video creation platform used by publishers such as Hearst and USA Today that automatically produces branded videos using content licensed from multiple sources.

RELEVANT RESEARCH AND ACADEMIC RESOURCES

Brown Institute for Media Innovation is a joint research program between the Stanford University School of Engineering and the Columbia University Graduate School of Journalism that funds researchers working to develop applications and technologies that can augment journalism via their annually awarded "Magic Grants."

MIT Media Lab is an interdisciplinary research laboratory at the Massachusetts Institute of Technology that devotes itself to projects at the convergence of technology, multimedia, sciences, art and design, as well as to groups such as the Laboratory of Social Machines and Civic Media that work to develop technologies and solutions for the journalism community.

Nieman Foundation at Harvard University works to promote and elevate the standards of journalism, while it's affiliate, Nieman Lab, publishes insights and industry updates on the latest innovations in the news media industry.

NYC Media Lab connects digital media and technology companies with universities and research centers to explore and develop new technologies, including artificial intelligence.

Open AI, a non-profit artificial intelligence research company associated with business magnate Elon Musk, aims to promote and develop "friendly" AI that will benefit humanity as a whole.

Reynolds Journalism Institute is a center affiliated with the Missouri School of Journalism that focuses on researching and testing new models of journalism in this era of technological advances.

Studio20 is a concentration at NYU Journalism School that explores innovating and adapting journalism for the web alongside digitally inclined media partners.

The Berkman Klein Center for Internet & Society is a research center at Harvard University that examines the ethical, social and cultural implications of new technology in society.

The CUNY Journalism School Research Center is dedicated to providing the latest research training, tools and resources for students and journalists.

The Knight Lab at Northwestern University is a collaboration between Medill School of Journalism and the McCormick School of Engineering that works to build tools that can help journalists tell better stories.

Tow Center for Digital Journalism is a research center at the Columbia Journalism School that explores the intersection between technology and journalism, focusing particularly on the ways technology can improve journalistic practices and consumption.

STAY INFORMED

AP Insights introduces readers to emerging technologies while presenting industry insights regarding their implications.

Columbia Journalism Review is a bimonthly publication of the Columbia University Graduate School of Journalism that explores news and media industry trends, often as they relate to new technologies.

Journalism.co.uk is a website that covers the intersection of technology and journalism.

MediaShift provides industry insights regarding journalism-based technologies.

Nieman Lab publishes insights and industry updates on the latest innovations in the news media industry.

NYC Media Lab newsletter is a weekly digital newsletter that details the most recent advances in AI, virtual reality, augmented reality and data science.

Storybench covers the world of digital storytelling and media innovation.

Tow Center blog explores the most recent information regarding the intersection of journalism and technology.

The definitive source of journalism innovation

Working on projects related to augmented journalism?
Get in touch with us: fmarconi@ap.org