

IBM USED NYPD SURVEILLANCE FOOTAGE TO DEVELOP TECHNOLOGY THAT LETS POLICE SEARCH BY SKIN COLOR

George Joseph, Kenneth Lipp

Illustration: Sally Thurer for The Intercept/Getty Images

In partnership with



IN THE DECADE after the 9/11 attacks, the New York City Police Department moved to put millions of New Yorkers under constant watch. Warning of [terrorism threats](#), the department created a plan to [carpet](#) Manhattan's downtown streets with thousands of cameras and had, [by 2008](#), centralized its video surveillance operations to a single command center. [Two years later](#), the NYPD announced that the command center, known as the Lower Manhattan Security Coordination

Center, had integrated cutting-edge video analytics software into select cameras across the city.

The video analytics software captured stills of individuals caught on closed-circuit TV footage and automatically labeled the images with physical tags, such as clothing color, allowing police to quickly search through hours of video for images of individuals matching a description of interest. At the time, the software [was also starting](#) to generate alerts for unattended packages, cars speeding up a street in the wrong direction, or people entering restricted areas.

Over the years, the NYPD has shared only occasional, small updates on the program's progress. In a 2011 interview with [Scientific American](#), for example, Inspector Salvatore DiPace, then commanding officer of the Lower Manhattan Security Initiative, said the police department was testing whether the software could box out images of people's faces as they passed by subway cameras and subsequently cull through the images for various unspecified "facial features."

While facial recognition technology, which measures individual faces at [over 16,000 points](#) for fine-grained comparisons with other facial images, has attracted significant [legal scrutiny](#) and [media attention](#), this object identification software has largely evaded attention. How exactly this technology came to be developed and which particular features the software was built to catalog have never been revealed publicly by the NYPD.

Now, thanks to [confidential corporate documents](#) and interviews with many of the technologists involved in developing the software, The Intercept and the Investigative Fund have learned that IBM began developing this object identification technology using secret access to NYPD camera footage. With access to images of thousands of unknowing New Yorkers offered up by NYPD officials, as early as 2012, IBM was creating new search features that allow other police departments to search camera footage for images of people by hair color, facial hair, and skin tone.


IBM declined to comment on its use of NYPD footage to develop the software. However, in an email response to questions, the NYPD did tell The Intercept that "Video, from time to time, was provided to IBM to ensure that the product they were developing would work in the crowded urban NYC environment and help us protect the City. There is nothing in the NYPD's agreement with IBM that prohibits sharing data with IBM for system development purposes. Further, all vendors who enter into contractual agreements with the NYPD have the absolute requirement to keep all data furnished by the NYPD confidential during the term of the agreement, after the completion of the agreement, and in the event that the agreement is terminated."

In an email to The Intercept, the NYPD confirmed that select counterterrorism officials had access to a pre-released version of IBM's program, which included skin tone search capabilities, as early as the summer of 2012. NYPD spokesperson Peter Donald said the search characteristics were only used for evaluation purposes and that officers were instructed not to include the skin tone search feature in their assessment. The department eventually decided not to integrate the analytics program into its larger surveillance architecture, and phased out the IBM program in 2016.

After testing out these bodily search features with the NYPD, IBM released some of these capabilities in a 2013 product release. [Later versions](#) of IBM's software retained and expanded these bodily search capabilities. (IBM did not respond to a question about the current availability of its video analytics programs.)

Asked about the secrecy of this collaboration, the NYPD said that "various elected leaders and stakeholders" were briefed on the department's efforts "to keep this city safe," adding that sharing camera access with IBM was necessary for the system to work. IBM did not respond to a question about why the company didn't make this collaboration public. Donald said IBM gave the department licenses to apply the system to 512 cameras, but said the analytics were tested on "fewer than fifty." He added that IBM personnel had access to certain cameras for the sole purpose of configuring NYPD's system, and that the department put safeguards in place to protect the data, including "non-disclosure agreements for each individual accessing the system; non-disclosure agreements for the companies the vendors worked for; and background checks."

Civil liberties advocates contend that New Yorkers should have been made aware of the potential use of their physical data for a private company's development of surveillance technology. The revelations come as a city council bill that would require NYPD transparency about surveillance acquisitions continues to languish, due, in part, to [outspoken](#) opposition from [New York City Mayor Bill de Blasio](#) and the [NYPD](#).

A rare look inside the New York Police Department's lower Manhattan security center, where cops monitor surveillance cameras, environmental sensors and license plate readers around the clock. Mayor Michael Bloomberg and Police Commissioner Ray Kelly announced that subway cameras are also being monitored in the center -- officially called The Lower Manhattan Security Coordination Center. Modeled after London's "Ring of Steel," the NYPD opened its coordination center in 2008. Today cops monitor feeds from over 1159 CCTV cameras with the number increasing to 3,000 as the program expands. (Photo by Timothy Fadek/Corbis via Getty Images)

Inside the New York City Police Department's lower Manhattan security center on Sept. 20, 2010, where cops monitor surveillance cameras, environmental sensors, and license plate readers around the clock. Photo: Timothy Fadek/Corbis via Getty Images

Skin Tone Search Technology, Refined on New Yorkers

IBM's initial breakthroughs in object recognition technology were envisioned for technologies like self-driving cars or image recognition on the internet, said Rick Kjeldsen, a former IBM researcher. But after 9/11, Kjeldsen and several of his colleagues realized their program was well suited for counterterrorism surveillance.

"After 9/11, the funding sources and the customer interest really got driven toward security," said Kjeldsen, who said he worked on the NYPD program from roughly 2009 through 2013. "Even though that hadn't been our focus up to that point, that's where demand was."

IBM's first major urban video surveillance project was with the Chicago Police Department and began around 2005, according to Kjeldsen. The department let IBM experiment with the technology in downtown Chicago until 2013, but the collaboration wasn't seen as a real business partnership. "Chicago was always known as, it's not a real — these guys aren't a real customer. This is kind of a development, a collaboration with Chicago," Kjeldsen said. "Whereas New York, these guys were a customer. And they had expectations accordingly."

The NYPD acquired IBM's [video analytics software](#) as one part of the Domain Awareness System, a shared project of the police department and Microsoft that [centralized](#) a vast web of surveillance sensors in lower and midtown Manhattan — including cameras, license plate readers, and radiation detectors — into a unified dashboard. IBM entered the picture as a subcontractor to Microsoft subsidiary Vexcel in 2007, as part of a project worth \$60.7 million over six years, according to the [internal IBM documents](#).

In New York, the terrorist threat "was an easy selling point," recalled Jonathan Connell, an IBM researcher who worked on the initial NYPD video analytics installation. "You say, 'Look what the terrorists did before, they could come back, so you give us some money and we'll put a camera there.'"

A former NYPD technologist who helped design the Lower Manhattan Security Initiative, asking to speak on background citing fears of professional reprisal, confirmed IBM's role as a "strategic vendor." "In our review of video analytics vendors at that time, they were well ahead of everyone else in my personal estimation," the technologist said.

According to internal IBM planning documents, the NYPD began integrating IBM's surveillance product in [March 2010](#) for the [Lower Manhattan Security Coordination Center](#), a counterterrorism command center [launched by Police Commissioner Ray Kelly in 2008](#). In a ["60 Minutes" tour](#) of the command center in

2011, [Jessica Tisch](#), then the NYPD's director of policy and planning for counterterrorism, showed off the software on gleaming widescreen monitors, demonstrating how it could pull up images and video clips of people in red shirts. Tisch did not mention the partnership with IBM.

During Kelly's tenure as police commissioner, the NYPD quietly worked with IBM as the company tested out its object recognition technology on a select number of NYPD and subway cameras, according to IBM documents. "We really needed to be able to test out the algorithm," said Kjeldsen, who explained that the software would need to process massive quantities of diverse images in order to learn how to adjust to the differing lighting, shadows, and other environmental factors in its view. "We were almost using the video for both things at that time, taking it to the lab to resolve issues we were having or to experiment with new technology," Kjeldsen said.

At the time, the department hoped that video analytics would improve analysts' ability to identify suspicious objects and persons in real time in sensitive areas, according to Conor McCourt, a [retired NYPD counterterrorism sergeant](#) who said he used IBM's program in its initial stages.

"Say you have a suspicious bag left in downtown Manhattan, as a person working in the command center," McCourt said. "It could be that the analytics saw the object sitting there for five minutes, and says, 'Look, there's an object sitting there.'" Operators could then rewind the video or look at other cameras nearby, he explained, to get a few possibilities as to who had left the object behind.

Over the years, IBM employees said, they started to become more concerned as they worked with the NYPD to allow the program to identify demographic characteristics. By 2012, according to the internal IBM documents, researchers were testing out the video analytics software on the bodies and faces of New Yorkers, capturing and archiving their physical data as they walked in public or passed through subway turnstiles. With these close-up images, IBM refined its ability to search for people on camera according to a variety of previously undisclosed features, such as age, gender, hair color (called "head color"), the presence of facial hair — and skin tone. The documents reference meetings between NYPD personnel and IBM researchers to review the development of body identification searches conducted at subway turnstile cameras.

"We were certainly worried about where the heck this was going," recalled Kjeldsen. "There were a couple of us that were always talking about this, you know, 'If this gets better, this could be an issue.'"

According to the NYPD, counterterrorism personnel accessed IBM's bodily search feature capabilities only for evaluation purposes, and they were accessible only to a handful of counterterrorism personnel. "While tools that featured either racial or

skin tone search capabilities were offered to the NYPD, they were explicitly declined by the NYPD,” Donald, the NYPD spokesperson, said. “Where such tools came with a test version of the product, the testers were instructed only to test other features (clothing, eyeglasses, etc.), but not to test or use the skin tone feature. That is not because there would have been anything illegal or even improper about testing or using these tools to search in the area of a crime for an image of a suspect that matched a description given by a victim or a witness. It was specifically to avoid even the suggestion or appearance of any kind of technological racial profiling.” The NYPD ended its use of IBM’s video analytics program in 2016, Donald said.


Donald acknowledged that, at some point in 2016 or early 2017, IBM approached the NYPD with an upgraded version of the video analytics program that could search for people by ethnicity. “The Department explicitly rejected that product,” he said, “based on the inclusion of that new search parameter.” In 2017, IBM released Intelligent Video Analytics 2.0, a product with a body camera surveillance capability that allows users to detect people captured on camera by “ethnicity” tags, such as “Asian,” “Black,” and “White.”

Kjeldsen, the former IBM researcher who helped develop the company’s skin tone analytics with NYPD camera access, said the department’s claim that the NYPD simply tested and rejected the bodily search features was misleading. “We would have not explored it had the NYPD told us, ‘We don’t want to do that,’” he said. “No company is going to spend money where there’s not customer interest.”

Kjeldsen also added that the NYPD’s decision to allow IBM access to their cameras was crucial for the development of the skin tone search features, noting that during that period, New York City served as the company’s “primary testing area,” providing the company with considerable environmental diversity for software refinement.

“The more different situations you can use to develop your software, the better it’s going to be,” Kjeldsen said. “That obviously pertains to people, skin tones, whatever it is you might be able to classify individuals as, and it also goes for clothing.”

The NYPD’s cooperation with IBM has since served as a selling point for the product at California State University, Northridge. There, campus police chief Anne Glavin said the technology firm IXP helped sell her on IBM’s object identification product by citing the NYPD’s work with the company. “They talked about what it’s done for New York City. IBM was very much behind that, so this was obviously of great interest to us,” Glavin said.

A rare look inside the New York Police Department's lower Manhattan security center, where cops monitor surveillance cameras, environmental sensors and license plate readers around the clock. Mayor Michael Bloomberg and Police Commissioner Ray Kelly announced that subway cameras are also being monitored in the center -- officially called The Lower Manhattan Security Coordination Center. Modeled after London's "Ring of Steel," the NYPD opened its coordination center in 2008. Today cops monitor feeds from over 1159 CCTV cameras with the number increasing to 3,000 as the program expands. (Photo by Timothy Fadek/Corbis via Getty Images)

A monitor showing surveillance footage of a New York street on Sept. 20, 2010, viewed inside the New York City Police Department's Lower Manhattan security center. Photo: Timothy Fadek/Corbis via Getty Images

Day-to-Day Policing, Civil Liberties Concerns

The NYPD-IBM video analytics program was initially envisioned as a counterterrorism tool for use in midtown and lower Manhattan, according to Kjeldsen. However, the program was integrated during its testing phase into dozens of cameras across the city. According to the former NYPD technologist, it could have been integrated into everyday criminal investigations.

“All bureaus of the department could make use of it,” said the former technologist, potentially helping detectives investigate everything from sex crimes to fraud cases. Kjeldsen spoke of cameras being placed at building entrances and near parking entrances to monitor for suspicious loiterers and abandoned bags.

Donald, the NYPD spokesperson, said the program’s access was limited to a small number of counterterrorism officials, adding, “We are not aware of any case where video analytics was a factor in an arrest or prosecution.”

Campus police at California State University, Northridge, who adopted IBM’s software, said the bodily search features have been helpful in criminal investigations. Asked about whether officers have deployed the software’s ability to filter through footage for suspects’ clothing color, hair color, and skin tone, Captain Scott VanScoy at California State University, Northridge, responded affirmatively, relaying a story about how university detectives were able to use such features to quickly filter through their cameras and find two suspects in a sexual assault case.

Join Our Newsletter
Original reporting. Fearless journalism.
Delivered to you.
I'm in

“We were able to pick up where they were at different locations from earlier that evening and put a story together, so it saves us a ton of time,” Vanscoy said. “By the

time we did the interviews, we already knew the story and they didn't know we had known.”

Glavin, the chief of the campus police, added that surveillance cameras using IBM's software had been placed strategically across the campus to capture potential security threats, such as car robberies or student protests. “So we mapped out some CCTV in that area and a path of travel to our main administration building, which is sometimes where people will walk to make their concerns known and they like to stand outside that building,” Glavin said. “Not that we're a big protest campus, we're certainly not a Berkeley, but it made sense to start to build the exterior camera system there.”

Civil liberties advocates say they are alarmed by the NYPD's secrecy in helping to develop a program with the potential capacity for mass racial profiling.

The identification technology IBM built could be easily misused after a major terrorist attack, argued Rachel Levinson-Waldman, senior counsel in the Brennan Center's Liberty and National Security Program. “Whether or not the perpetrator is Muslim, the presumption is often that he or she is,” she said. “It's easy to imagine law enforcement jumping to a conclusion about the ethnic and religious identity of a suspect, hastily going to the database of stored videos and combing through it for anyone who meets that physical description, and then calling people in for questioning on that basis.” IBM did not comment on questions about the potential use of its software for racial profiling. However, the company did send a comment to *The Intercept* pointing out that it was “one of the first companies anywhere to adopt a set of principles for trust and transparency for new technologies, including AI systems.” The statement continued on to explain that IBM is “making publicly available to other companies a dataset of annotations for more than a million images to help solve one of the biggest issues in facial analysis — the lack of diverse data to train AI systems.”

Few laws clearly govern object recognition or the other forms of artificial intelligence incorporated into video surveillance, according to Clare Garvie, a law fellow at Georgetown Law's Center on Privacy and Technology. “Any form of real-time location tracking may raise a Fourth Amendment inquiry,” Garvie said, citing a 2012 Supreme Court case, [United States v. Jones](#), that involved police monitoring a car's path without a warrant and resulted in five justices suggesting that individuals could have a reasonable expectation of privacy in their public movements. In addition, she said, any form of “identity-based surveillance” may compromise people's right to anonymous public speech and association.

Garvie noted that while facial recognition technology has been heavily criticized for the risk of false matches, that risk is even higher for an analytics system “tracking a

person by other characteristics, like the color of their clothing and their height,” that are not unique characteristics.

The former NYPD technologist acknowledged that video analytics systems can make mistakes, and noted a study where the software had trouble characterizing people of color: “It’s never 100 percent.” But the program’s identification of potential suspects was, he noted, only the first step in a chain of events that heavily relies on human expertise. “The technology operators hand the data off to the detective,” said the technologist. “You use all your databases to look for potential suspects and you give it to a witness to look at. ... This is all about finding a way to shorten the time to catch the bad people.”

Object identification programs could also unfairly drag people into police suspicion just because of generic physical characteristics, according to Jerome Greco, a digital forensics staff attorney at the Legal Aid Society, New York’s [largest public defenders organization](#). “I imagine a scenario where a vague description, like young black male in a hoodie, is fed into the system, and the software’s undisclosed algorithm identifies a person in a video walking a few blocks away from the scene of an incident,” Greco said. “The police find an excuse to stop him, and, after the stop, an officer says the individual matches a description from the earlier incident.” All of a sudden, Greco continued, “a man who was just walking in his own neighborhood” could be charged with a serious crime without him or his attorney ever knowing “that it all stemmed from a secret program which he cannot challenge.”

While the technology could be used for appropriate law enforcement work, Kjeldsen said that what bothered him most about his project was the secrecy he and his colleagues had to maintain. “We certainly couldn’t talk about what cameras we were using, what capabilities we were putting on cameras,” Kjeldsen said. “They wanted to control public perception and awareness of LMSI” — the Lower Manhattan Security Initiative — “so we always had to be cautious about even that part of it, that we’re involved, and who we were involved with, and what we were doing.” (IBM did not respond to a question about instructing its employees not to speak publicly about its work with the NYPD.)

The way the NYPD helped IBM develop this technology without the public’s consent sets a dangerous precedent, Kjeldsen argued.

“Are there certain activities that are nobody’s business no matter what?” he asked. “Are there certain places on the boundaries of public spaces that have an expectation of privacy? And then, how do we build tools to enforce that? That’s where we need the conversation. That’s exactly why knowledge of this should become more widely available — so that we can figure that out.”

This article was reported in partnership with the Investigative Fund at the Nation Institute.

