Security tool guarantees privacy in surveillance footage

March 28 2022, by Rachel Gordon

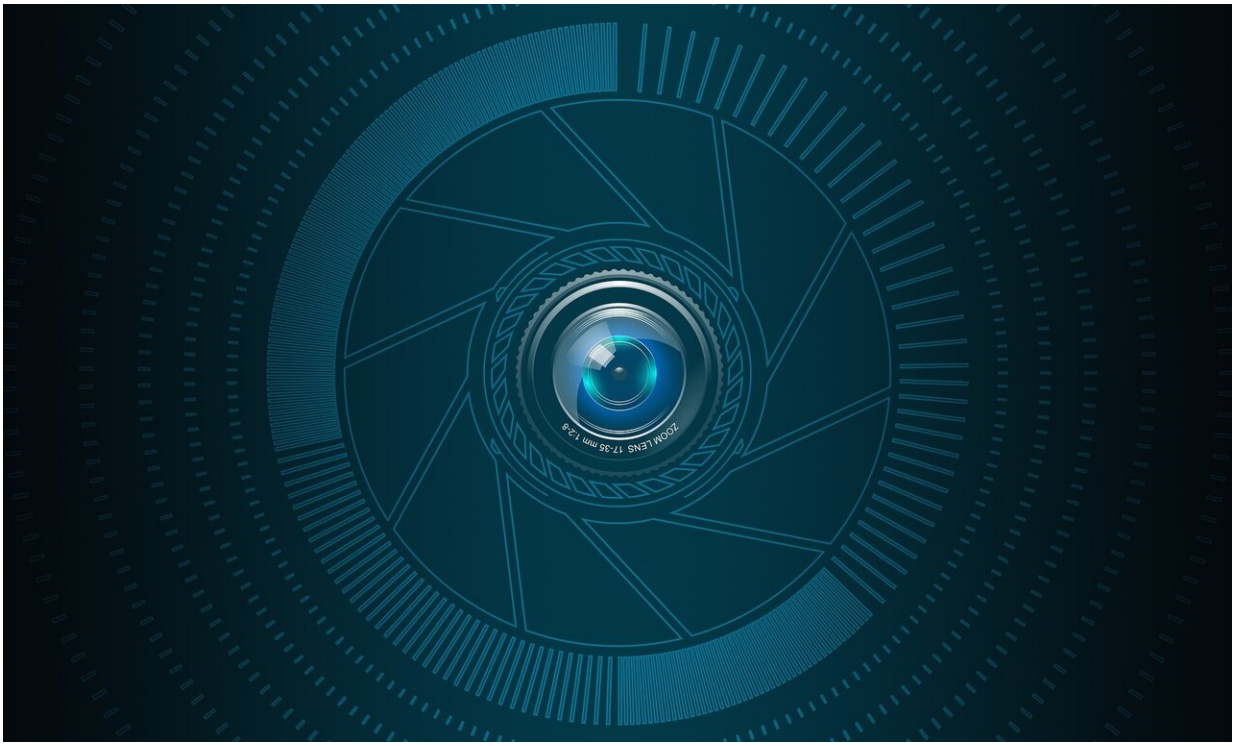

Credit: Pixabay/CC0 Public Domain

Surveillance cameras have an identity problem, fueled by an inherent tension between utility and privacy. As these powerful little devices have cropped up seemingly everywhere, the use of machine learning tools has automated video content analysis at a massive scale—but with increasing mass surveillance, there are currently no legally enforceable rules to limit

privacy invasions.

Security cameras can do a lot—they've become smarter and supremely more competent than their ghosts of grainy pictures past, the ofttimes "hero tool" in crime media. ("See that little blurry blue blob in the right hand corner of that densely populated corner—we got him!") Now, video surveillance can help health officials measure the fraction of people wearing masks, enable transportation departments to monitor the density and flow of vehicles, bikes, and pedestrians, and provide businesses with a better understanding of shopping behaviors. But why has privacy remained a weak afterthought?

The status quo is to retrofit video with blurred faces or black boxes. Not only does this prevent analysts from asking some genuine queries (e.g., Are people wearing masks?), it also doesn't always work; the system may miss some faces and leave them unblurred for the world to see. Dissatisfied with this status quo, researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL), in collaboration with other institutions, came up with a system to better guarantee privacy in video footage from surveillance cameras. Called "Privid," the system lets analysts submit video data queries, and adds a little bit of noise (extra data) to the end result to ensure that an individual can't be identified. The system builds on a formal definition of privacy—"differential privacy"—which allows access to aggregate statistics about private data without revealing personally identifiable information.

Typically, analysts would just have access to the entire video to do whatever they wanted with it, but Privid makes sure the video isn't a free buffet. Honest analysts can get access to the information they need, but that access is restrictive enough that malicious analysts can't do too much with it. To enable this, rather than running the code over the entire video in one shot, Privid breaks the video into small pieces and runs processing

code over each chunk. Instead of getting results back from each piece, the segments are aggregated, and that additional noise is added. (There's also information on the error bound you're going to get on your result—maybe a 2 percent error margin, given the extra noisy data added).

For example, the code might output the number of people observed in each video chunk, and the aggregation might be the "sum," to count the total number of people wearing face coverings, or the "average" to estimate the density of crowds.

Privid allows analysts to use their own deep neural networks that are commonplace for video analytics today. This gives analysts the flexibility to ask questions that the designers of Privid did not anticipate. Across a variety of videos and queries, Privid was accurate to within 79 to 99 percent of a non-private system.

"We're at a stage right now where cameras are practically ubiquitous. If there's a camera on every street corner, every place you go, and if someone could actually process all of those videos in aggregate, you can imagine that entity building a very precise timeline of when and where a person has gone," says MIT CSAIL Ph.D. student Frank Cangialosi, the lead author on a paper about Privid. "People are already worried about location privacy with GPS—video data in aggregate could capture not only your location history, but also moods, behaviors, and more at each location."

Privid introduces a new notion of "duration-based privacy," which decouples the definition of privacy from its enforcement—with obfuscation, if your privacy goal is to protect all people, the enforcement mechanism needs to do some work to find the people to protect, which it may or may not do perfectly. With this mechanism, you don't need to fully specify everything, and you're not hiding more information than

you need to.

Let's say we have a video overlooking a street. Two analysts, Alice and Bob, both claim they want to count the number of people that pass by each hour, so they submit a video processing module and ask for a sum aggregation.

The first analyst is the city planning department, which hopes to use this information to understand footfall patterns and plan sidewalks for the city. Their model counts people and outputs this count for each video chunk.

The other analyst is malicious. They hope to identify every time "Charlie" passes by the camera. Their model only looks for Charlie's face, and outputs a large number if Charlie is present (i.e., the "signal" they're trying to extract), or zero otherwise. Their hope is that the sum will be non-zero if Charlie was present.

From Privid's perspective, these two queries look identical. It's hard to reliably determine what their models might be doing internally, or what the analyst hopes to use the data for. This is where the noise comes in. Privid executes both of the queries, and adds the same amount of noise for each. In the first case, because Alice was counting all people, this noise will only have a small impact on the result, but likely won't impact the usefulness.

In the second case, since Bob was looking for a specific signal (Charlie was only visible for a few chunks), the noise is enough to prevent them from knowing if Charlie was there or not. If they see a non-zero result, it might be because Charlie was actually there, or because the model outputs "zero," but the noise made it non-zero. Privid didn't need to know anything about when or where Charlie appeared, the system just needed to know a rough upper bound on how long Charlie might appear

for, which is easier to specify than figuring out the exact locations, which prior methods rely on.

The challenge is determining how much noise to add—Privid wants to add just enough to hide everyone, but not so much that it would be useless for analysts. Adding noise to the data and insisting on queries over time windows means that your result isn't going to be as accurate as it could be, but the results are still useful while providing better privacy.

Cangialosi wrote the paper with Princeton Ph.D. student Neil Agarwal, MIT CSAIL Ph.D. student Venkat Arun, assistant professor at the University of Chicago Junchen Jiang, assistant professor at Rutgers University and former MIT CSAIL postdoc Srinivas Narayana, associate professor at Rutgers University Anand Sarwate, and assistant professor at Princeton University and Ravi Netravali of MIT. Cangialosi will present the paper at the USENIX Symposium on Networked Systems Design and Implementation Conference in April in Renton, Washington.

More information: Privid: Practical, Privacy-Preserving Video Analytics Queries, arXiv:2106.12083 [cs.CR]
doi.org/10.48550/arXiv.2106.12083