THE UNIVERSITY OF CHICAGO

Arjun Nitin Bhagoji and Emily Wenger
Department of Computer Science
University of Chicago
Contact: abhagoji@uchicago.edu,
ewenger@uchicago.edu

# Comment for the Federal Trade Commission Regulation Rule on Commercial Surveillance and Data Security

We thank the Federal Trade Commission (FTC) for the opportunity to provide feedback on the 'Commercial Surveillance ANPR, R111004' on how regulations could protect consumers from commercial surveillance practices and improve data security practices. We appreciate the Commission's commitment to rulemaking on such issues, since individual consumers are incapable of resolving these issues on their own, and companies lack incentives to do so.

We are academic researchers affiliated with the Computer Science Department at the University of Chicago. We study privacy and security issues of machine learning models, particularly facial recognition systems. Our work, and that of other researchers, has shown that: a) consumers have little to no knowledge of when their data is scraped and used in facial recognition systems [1], b) model-derived biometric artifacts (such as face templates) are often stored carelessly despite their sensitive nature [2], and c) anti-facial recognition tools can shield consumers from unwanted commercial surveillance, but have tenuous legal status and need policymaking to augment their effectiveness[3].

**Summary of response.** Our years of research on facial recognition and related topics has led us to conclude that regulatory guide-rails for facial recognition—and other commercial surveillance systems—are critical. Thus, we are grateful for this opportunity to share recommendations. We hope our insights can aid the construction of meaningful and effective regulatory guidelines. Leveraging our prior work and expertise, our responses are most pertinent to *Questions 37 and 38* in the ANPR. We focus our responses on the context of widely-used facial recognition (FR) systems that rely on machine learning models, since we are most familiar with these. Occasionally, our responses relate to other questions in the ANPR. We have made an effort to clearly state when they do.

Our response makes the following recommendations, which are discussed in the rest of this document:

1. The FTC should require registration of FR systems used for commercial surveillance and should notify consumers of their use.
2. The FTC should provide a stricter definition for "commercial" FR models.
3. The FTC should mandate that explicit consumer consent be obtained before consumer images are used in FR systems.
4. The FTC should require encrypted storage of biometric artifacts used in commercial FR models.
5. The FTC should provide legal protection and guidelines for creators and consumers of anti-FR tools.

[1] E. Wenger, S. Shan, H. Zheng, and B. Y. Zhao. Sok: Anti-facial recognition technology. In *2023 IEEE Symposium on Security and Privacy (SP)*, 2023

[2] E. Wenger, F. Falzon, J. Passananti, H. Zheng, and B. Y. Zhao. Assessing privacy risks from feature vector reconstruction attacks. *arXiv preprint arXiv:2202.05760*, 2022

[3] S. Shan, E. Wenger, J. Zhang, H. Li, H. Zheng, and B. Y. Zhao. Fawkes: Protecting Privacy Against Unauthorized Deep Learning Models. In *Proc. of USENIX Security*, 2020

6. The FTC should require existing FR models to be retrained *after* ensuring all training data was obtained with consumers' consent.

---

## Background: Machine Learning-based Facial Recognition Systems

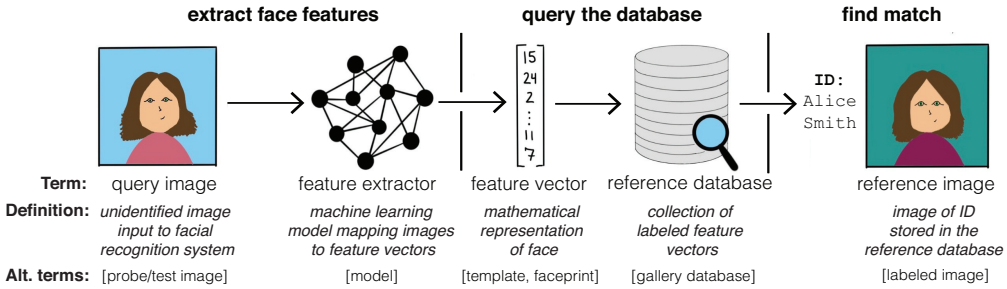


Figure 1: *The workflow of how facial recognition systems recognize a human face in an input image, along with the corresponding terminology. (a): A query image, after being submitted to the system, is passed to the feature extractor to produce a feature vector; (b): this feature vector is used to query a reference database of labeled feature vectors; (c): if the query feature vector matches a labeled feature vector in the database, the label is used to find a reference image, and the system outputs the reference image and the identity (i.e. Alice Smith in this example).*

Since our response focuses primarily on facial recognition (FR) technology, we begin by providing a broad overview of how modern FR systems that rely on machine learning work. A reader familiar with this subject may skip ahead to Observation 1, although they may still find the summary we provide here useful.

FR systems identify people by their facial characteristics, generally by comparing an unidentified human face in an image or a video against a database of facial images with known identities. While there are many design variants[4], we focus on state-of-the-art, widely adopted FR systems, which employ machine learning models called deep neural networks (DNNs) to perform recognition on digital face images. DNNs have been found to be particularly effective at extracting high-level features from large amounts of data [5].

At a high level, facial recognition engines work as follows (see Figure 1). *First*, a *query image*, i.e. a face image to be identified, is fed through a *feature extractor*, a DNN that converts the image into a *feature vector* (or a mathematical representation of the person's facial features). *Next*, this feature vector is used to query a *reference database*, a collection of face images of known identities. This query search is done by comparing the input feature vector against the reference feature vectors stored in the database to find the closest match. *Finally*, if the query search finds a reference feature vector in the database sufficiently similar to the input, the FR system declares that a match has been found and outputs the corresponding identity and the associated *reference image*.

We note the distinction between *facial recognition* systems and *facial verification* systems. Facial verification is used widely to authenticate users on mobile devices (e.g. FaceID on iPhones), by checking the similarity of a user's facial features against the stored feature vector matching the authorized user. Since facial recognition systems are used *at scale* to identify *numerous people* rather than to *authenticate a specific person*, we focus on facial recognition systems in our observations and recommendations.

## Observation 1: Biometric info is often collected without consent and used in commercial FR systems —which may also be *government* FR systems.

We first address the following sub-questions within Question 37: *"How do companies collect consumers' biometric information? What kinds of biometric information do companies collect? For what purposes do they collect and use it? Are consumers typically aware of that collection and use?"* At a high level, in the context of

[4] M. Taskiran, N. Kahraman, and C. E. Erdem. Face recognition: Past, present and future (a review). *Digital Signal Processing*, 2020

[5] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`

FR systems, we can confidently say that companies collect face images from a variety of sources for use in facial recognition systems. Commercial facial recognition engines are deployed for varied purposes —from in-store monitoring to securing facility access —and are also used by the US government. Consumers are *not typically aware of this collection and use*. We expand on these points below.

**Companies use and sell FR systems.** Commercial use of facial recognition (FR) systems has skyrocketed in recent years, fueled by algorithmic advances and the recognition of their potential[6]. Many corporations have integrated FR into their security and commerce pipelines. The most common FR use cases are enhancing store or office security (see Table 1). For example, companies like Apple, Macy's, and Lowes use FR to catch shoplifters in their stores[7]. Other companies have employed FR to monitor corporate facility access[8]. Product-based applications have also emerged, such as car companies like Subaru using FR to track driver fatigue[9] or airlines using FR to streamline passenger check-in[10].

But where do these FR systems come from? The ecosystem of FR systems is murky at best. Some companies may produce their own custom FR systems. More often than not, though, FR systems are purchased from vendors like NEC, Cognitec, Idemia, or Jemalto[11]. Little can be gleaned from online sources about how these vendors develop their systems, and if they sell only FR models themselves, or also the reference database of images (c.f. Figure 1).

**Recommendation 1**: *The FTC should require registration of facial recognition systems used for commercial surveillance and should notify consumers of their use.*

Consumers have a right to know when and where facial recognition systems are deployed, due to the sensitive nature of data collected by these systems. We recommend that **the FTC mandate that public notification be given in consumer-facing settings where FR systems are used** —whether they are used in commercial products or for aiding company operations in ways that affect consumers (e.g. in-store monitoring). Furthermore, we recommend that the FTC retain a centralized (and perhaps, public) list of all registered FR models to ensure they can be periodically audited.

**The US government uses commercial FR systems.** Entities within the US government also use facial recognition systems purchased from commercial entities[12]. Most frequently, the systems are deployed for security-related tasks like border control or criminal suspect identification[13]. The US Government Accountability Office (GAO) recently produced a report highlighting which government agencies use commercial facial recognition systems and why[14]. The report recommends concrete steps to better monitor use of such systems, but as of this writing, only the Secret Service and FDA have taken these steps[15]. We highlight government uses of commercial FR systems to emphasize that the notion of "commercial" facial recognition requires stricter definitions.

[6] See Grandview Research. Facial Recognition Market Size, Share and Trends Report, 2021 - 2028, 2021

[7] See H. Towney. The retail stores you probably shop at that use facial-recognition technology. *Business Insider*, July 2021

[8] See U. Saiidi. We went inside Alibaba's global headquarters. Here's what we saw. *CNBC*, September 2019; and M. Rogoway. Major Tech Company Using Facial Recognition to ID Workers. *The Oregonian*, March 2020

[9] See P. Lyon. Subaru Forester Is First Mainstream Model To Offer Facial Recognition Technology. *Forbes*, April 2018

[10] See A. Smith. JetBlue will test facial recognition for boarding. *CNN Business*, May 2017; and Could facial recognition be the future of airport security? Delta Air Lines is testing it out. *CBS News*, October 2021

[11] Nec; and Idemia; and Thales group; and Cognitec; and Markets and Markets. Facial recognition market worth 8.5 billion by 2025

[12] See United States Government Accountability Office. Facial Recognition Technology: Federal Law Enforcement Agencies Should Better Assess Privacy and Other Risks. *GAO-21-518*, 2021

[13] See US Customs and Border Control. Collection of Biometric Data from Aliens Upon Entry to and Departure from the United States. *USCBP-2020-0062*, 2020; C. Garvie, A. Bedoya, and J. Frankle. The perpetual lineup. *Georgetown Law Center on Privacy and Technology*, 2016; M. Mason. Biometric Breakthrough: How CBP is Meeting Its Mandate and Keeping America Safe. *U.S. Customs and Border Protection Website*

[14] United States Government Accountability Office. Facial Recognition Technology: Federal Law Enforcement Agencies Should Better Assess Privacy and Other Risks. *GAO-21-518*, 2021

[15] United States Government Accountability Office. Follow-up Recommendations from GAO-21-518. *GAO-21-518*, 2022

**Recommendation 2**: *The FTC should provide a stricter definition for "commercial" facial recognition models.*

The FTC and the US public would benefit from a clearer definition of "commercial" facial recognition. Therefore, we recommend that **the FTC more clearly define what constitutes "commercial" vs. "government" facial recognition.** Is commercial FR *any* FR system that is sold by a commercial entity (e.g. even those purchased by the government)? Or is it simply a FR *used by a company* for company-related initiatives? If the latter, then appropriate regulations should be developed to monitor government, as well as commercial, uses of FR technology.

| Commercial Use Cases | Companies |
|---|---|
| Catching shoplifters | Apple, Macy's, Lowe's[16] |
| Securing facility access | Alibaba[17], Intel[18] |
| Tracking driver behavior | Hyundai[19], Subaru[20] |
| Passenger check-in | JetBlue[21], Delta[22] |
| Virtual makeup try-on | Mac Cosmetics[23] |

Table 1: Example uses of facial recognition.

| Operator of FR system | Source of reference images |
|---|---|
| Clearview.ai | Social media photos[24] |
| PimEyes | (Public) online photos[25] |
| FBI F.A.C.E.S. | State drivers' license photos[26] |
| US Customs and Border Patrol | Passport photos[27] |

Table 2: Reported reference image sources

**Images used in FR systems come from many sources.** The definitive source of images for commercial FR models is often unknown. However, based on government reports and media articles, we outline some known sources of training, reference, and query images used by today's FR systems. The distinction between these three types of images is highlighted in Figure 1.

- **Training images** (used to train feature extractors) often come from a mix of academic training datasets[28], proprietary data, and public data scraped from social media accounts, according to the US GAO[29]. Little is known about the proprietary or public data —e.g. where it comes from, if consent was obtained before use —but it is well-known that academic face recognition datasets are often built without user consent[30].
- **Reference images**, used to create the reference database, generally come from the Internet (e.g. social media), or government databases (e.g. passport and driver license photos). A list of known reference image sources for some well-known FR operators is shown in Table 2. It is unclear if commercial FR

[16] T. Clayburn. Apple sued in nightmare case involving teen wrongly accused of shoplifting, driver's permit used by impostor, and unreliable facial-rec tech. *The Register*, May 2021; H. Towney. The retail stores you probably shop at that use facial-recognition technology. *Business Insider*, July 2021

[17] U. Saiidi. We went inside Alibaba's global headquarters. Here's what we saw. *CNBC*, September 2019

[18] M. Rogoway. Major Tech Company Using Facial Recognition to ID Workers. *The Oregonian*, March 2020

[19] In-Car Biometric Technology For Human Interaction. *Hyundai Motor Group*, *blog article*, 2020

[20] P. Lyon. Subaru Forester Is First Mainstream Model To Offer Facial Recognition Technology. *Forbes*, April 2018

[21] A. Smith. JetBlue will test facial recognition for boarding. *CNN Business*, May 2017

[22] Could facial recognition be the future of airport security? Delta Air Lines is testing it out. *CBS News*, October 2021

[23] J. J. Low. Is the beauty industry leading with AR experiences? *TechHQ*, June 2020

[24] K. Hill. The Secretive Company that May End Privacy as We Know It. *The New York Times*, January 2020

[25] Pimeyes

[26] C. Garvie, A. Bedoya, and J. Frankle. The perpetual lineup. *Georgetown Law Center on Privacy and Technology*, 2016

[27] M. Mason. Biometric Breakthrough: How CBP is Meeting Its Mandate and Keeping America Safe. *U.S. Customs and Border Protection Website*

[28] *e.g.* Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao. MS-Celeb-1M: A Dataset and Benchmark for Large-Scale Face Recognition. *arXiv preprint arXiv:1607.08221*, 2016; and Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. VGGFace2: A Dataset for Recognising Faces across Pose and Age. In *Proc. of FG*, 2018; and H.-W. Ng and S. Winkler. A data-driven approach to cleaning large face datasets. In *Proc. of ICIP*, 2014; and D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning Face Representation from Scratch. *arXiv preprint arXiv:1411.7923*, 2014

[29] United States Government Accountability Office. Facial Recognition Technology: Privacy and Accuracy Issues Related to Commercial Uses. *GAO-20-522*, 2020

[30] O. Solon. Facial recognition's 'dirty little secret': Millions of online photos scraped without consent. *NBC News*, 2019

systems provide a reference database as part of their FR "product," or if purchasers of the commercial system (e.g. another company or government entity) must create the database themselves.

- Finally, **query images** can come from both online and physical sources, including social media, police body cams, mug shots, corporate surveillance systems, state identification images, passport photos, and others[31]. After identification, query images can be fed back into the reference database, either to enhance existing feature vectors or create new ones. Effectively, this can mean that images obtained in public spaces can be used to enroll individuals in facial recognition systems.

**Recommendation 3**: *The FTC should mandate that explicit consumer consent be obtained before consumer images are used in FR systems.*

Currently, consumers cannot meaningfully choose when or if their data is included in FR systems. Simply entering a store, posting a picture online, or agreeing to an app's or website's lengthy terms-of-service contract can prompt enrollment in a system. Furthermore, consumers cannot opt-out (or even opt-in) to use of FR systems. To even the balance of power and ensure greater accountability in when and how face images are used, we recommend that the FTC mandate that **explicit, meaningful consent be obtained from consumers *before* either (1) their identities enrolled in a FR system or (2) their images are searched through a FR database.** Meaningful consent *must* include a way for consumers to opt out without having their access to services to reduced or their time unnecessarily hindered.

---

## Observation 2: When model-derived biometric artifacts are stored unencrypted, FR systems run faster, but users' privacy is endangered.

Having established the use cases and data sources for commercial FR systems, we now address the functionality of those systems. In particular, we discuss how biometric data—users' face images—are stored in FR systems today, and the privacy/security risks of those storage methods.

**FR systems rely on fast "feature vector" matching.** As Figure 1 illustrates, a key step in FR systems is the database query, in which an image is searched against stored reference images to find a match. Such queries are expedited by the use of "feature vectors," which are model-generated, numerical representations of image features. Images of the same person map to similar feature vectors. Since feature vectors are easier to compare than full-size images, fast matching is possible[32]. This fast matching allows FR systems to quickly find matches in databases containing thousands or even millions of individuals.

**Leaked feature vectors pose privacy risks.** There is a catch—feature vector matching is only fast *if the vectors are stored unencrypted*. Encryption, a way of encoding information to protect it, would make comparisons between feature vectors much slower[33]. Hence, it is common for FR operators to store feature vectors in a secure database, assuming that if the database is secure, the feature vectors will be as well. More critically, several FR operators claim that *since feature vectors are numerical, they are inherently "encrypted"*, rendering them safe even if the database is breached. [34]

[31] United States Government Accountability Office. Facial Recognition Technology: Federal Law Enforcement Agencies Should Better Assess Privacy and Other Risks. *GAO-21-518*, 2021

[32] For more info, see H. Cevikalp and H. Serhan Yavuz. Fast and accurate face recognition with image sets. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1564–1572, 2017

[33] For example, see P. Mishra, R. Lehmkuhl, A. Srinivasan, W. Zheng, and R. A. Popa. Delphi: A cryptographic inference service for neural networks. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 2505–2522, 2020

[34] For example, FaceFirst (`https://facefirst.com/trust`) says they "generate a proprietary and anonymous biometric template that is not correlated to personally identifiable information". Oosto, another leading vendor (`https://oosto.com/why-trust-us/`) states that the "Oosto system only stores mathematical vectors of the POI persons" and offers that "all facial signatures and facial images can be encrypted in transit with AES-256 bit (based on customer needs) from the camera to the local server."

This claim is patently false. Our prior work has shown that, with query-only access to a FR model, an attacker could recreate facial images from feature vectors. The reconstructed images are high-enough quality that they can then be *re-identified* by a secondary FR engine[35]. This reconstruction/reidentification attack could compromise the privacy and potentially the safety of individuals enrolled in the FR system. All it would take is a leak of the feature vector database, a not-uncommon occurrance.[36] [37]

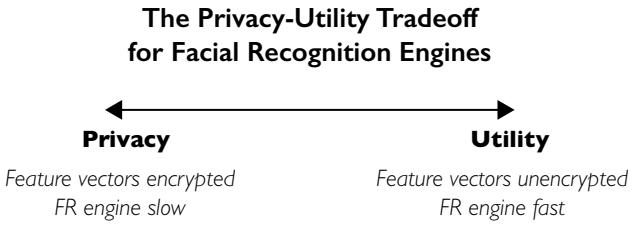


Figure 2: Illustration of the privacy-utility tradeoff for facial recognition engines and what each party (consumers, FR operators) lose/gain at each end of the continuum.

**Consumers lose in the FR privacy-utility tradeoff.** Storing unencrypted feature vectors in a database leads to privacy risks for the individuals enrolled in the FR system. However, FR providers do not currently encrypt feature vectors because it is much *faster* to work with unencrypted data. Fast matching enables FR vendors to quickly search databases that contain hundreds or even thousands of identities. This reality illustrates a phenomenon widely known in machine learning and privacy research as the "privacy-utility tradeoff," see Figure 2.

Privacy-preserving machine learning usually operates in one of two paradigms: cryptographic privacy or statistical privacy. The former requires that all operations be done on encrypted data, and does not assume that the entity collecting the data is trusted. However, operations on encrypted data are usually prohibitively expensive, particularly in the case of machine learning where there is a large amount of data and complex models. Thus, the latter model of statistical privacy has gained popularity[38] in cases where there is a trusted entity collecting the data. In this case, the privacy concerns lie with entities downstream who may try to compromise the privacy of users by performing inference over released models or statistics. To prevent this, appropriately calibrated noise is added to data byproducts such as models or statistics to prevent inference of private information. In either case, however, there is a *utility* trade-off. In the case of privacy based on cryptographic primitives, it is usually the computational cost and the inability to perform certain mathematical operations on the data. For statistical privacy, it is the loss in accuracy due to the addition of noise. For example, when computing the mean of a statistic over a population, the addition of poorly calibrated noise can make the resulting value meaningless. Thus, there is little to no incentive for entities deploying surveillance systems to use privacy-preserving versions of these systems as their performance will suffer.

**Recommendation 4**: *The FTC should require encrypted storage of biometric artifacts used in commercial facial recognition models.*

We recommend that the Commission set strict requirements for the storage and privacy of biometric artifacts derived from machine learning models, particularly feature vectors used in FR systems. In particular, we recommend that they require *encrypted* storage of such artifacts, rather than allowing companies to rely on database security to protect this information.

[35] For more on this attack, see E. Wenger, F. Falzon, J. Passananti, H. Zheng, and B. Y. Zhao. Assessing privacy risks from feature vector reconstruction attacks. *arXiv preprint arXiv:2202.05760*, 2022

[36] For example, see Biostar. Report: Data breach in biometric security platform affecting millions of users. 2021

[37] In the case of Ooosto, a curious observer could sniff the web traffic between the Oosto server and client's local server and obtain feature vectors, since they do not by default encrypt feature vectors in-transit (Footnote 18).

[38] Differential privacy, introduced by C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006, is being used in commercial settings by companies such as Apple (Apple. Differential privacy overview) to protect private data

The additional benefit of regulating the storage of biometric artifacts is that it may limit the scope of commercial surveillance. When feature vectors are protected, the facial recognition engine's search functionality will be slower, making it less feasible to search thousands of faces at once. This will force engineers to be more thoughtful with how their facial recognition engines are used and perhaps curtail some (ab)uses of facial recognition seen today.

---

**Observation 3: Anti-surveillance tools, which exploit algorithmic flaws, enhance consumers' ability to avoid surveillance but have nebulous legal status.**

The algorithms underpinning FR-enabled commercial surveillance systems have numerous documented flaws and biases. They have high error rates on underrepresented groups [39], can be trained to (mis)behave in malicious ways[40], and perform poorly on instances not encountered during training [41]. Furthermore, these algorithms are unable to take context into consideration when generating predictions [42] and are cannot pass even basic reasoning tests [43]. These observations relate to *Questions 53 and 54.*

Public documentation of these flaws—and others—in FR systems has led to pushback against their use. However, without regulation, operators of FR-enabled surveillance systems have few incentives to check for such flaws in their systems and correct them, if present. In light of this, it stands to reason that consumers should be able to choose not to be surveilled by systems relying on algorithms with dubious performance and unclear benefits. Furthermore, if consumers choose to evade dubious FR surveillance systems, they should still be able to obtain goods or services from commercial entities that engage in surveillance.

Intriguingly, as consumers and other entities develop ways to evade unwanted FR-enabled surveillance, some of the aforementioned flaws have come in handy. Anti-surveillance tools leverage known weaknesses in FR algorithms to prevent unwanted recognition. Typically, they target either FR system training[44] or deployment[45]. The main technical challenge facing these tools is that they operate in a cyberphysical space. Anti-FR tools are often physical objects (e.g. masks, t-shirts) that interact with a digital entity (e.g. FR model) after both physically and digitally-mediated changes (e.g. lighting during picture taking, varying quality after upload). Research and development of these tools is cutting-edge and ongoing.

We believe that the wide availability of anti-FR tech would empower consumers to *choose if they want to be surveilled*, even if regulations fail to meaningfully enforce a choice. These tools would allow consumers to 'opt-out' of being surveilled in a cyberphysical environment while still entering the space where surveillance is occurring. For example, a customer shopping at a grocery store which that is using FR can wear an

[39] See J. Buolamwini and T. Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91, 2018

[40] See M. Goldblum, D. Tsipras, C. Xie, X. Chen, A. Schwarzschild, D. Song, A. Madry, B. Li, and T. Goldstein. Dataset security for machine learning: Data poisoning, backdoor attacks, and defenses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022 and E. Wenger, J. Passananti, A. N. Bhagoji, Y. Yao, H. Zheng, and B. Y. Zhao. Backdoor attacks against deep learning systems in the physical world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6206–6215, June 2021

[41] Data at training and deployment time may differ due to naturally occurring data drift (see D. Hendrycks and K. Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *5th International Conference on Learning Representations, ICLR*, 2017 ) or deliberately modified inputs (see N. Carlini and D. Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017)

[42] K. Hartnett. Machine learning confronts the elephant in the room. *Quanta Magazine*, 2018

[43] L. Pandia and A. Ettinger. Sorting through the noise: Testing robustness of information processing in pre-trained language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021

[44] One such tool, Fawkes (see S. Shan, E. Wenger, J. Zhang, H. Li, H. Zheng, and B. Y. Zhao. Fawkes: Protecting Privacy Against Unauthorized Deep Learning Models. In *Proc. of USENIX Security*, 2020) recommends individuals upload images of themselves on the Web with carefully-crafted perturbations so even if they are scraped and used to train models without consent, the resultant model is unlikely to be able to identify them.

[45] Proposals for physical adversarial examples demonstrate that clothing and other accessories such as glasses (see M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proc. of CCS*, 2016) and masks can be designed to mislead facial recognition and other ML-driven surveillance algorithms.

anti-FR face mask[46] and shop without being surveilled.

Given the significance of anti-FR tools in preserving consumers' agency and privacy, regulation must be enacted to support them. In isolation, these tools are inadequate to protect consumers from surveillance. Furthermore, their use result in legal action that consumers and providers are ill-equipped to handle. For instance, repeated legal action has been taken against companies providing advertisement blocking tools by publishing houses[47]. While rulings in these cases typically side with ad-blockers, ad-block tools still have no formal legal protections. Providing proactive, protective regulation for anti-FR tools—before such legal cases arise—would benefit consumers and companies.

**Recommendation 5**: *The FTC should provide legal protection and guidelines for creators and consumers of anti-surveillance tools.*

We recommend the Commission enact regulation protecting consumers who choose to use anti-surveillance technology in spaces where they may be under surveillance. They should also clarify when and where such use violates mandates from other agencies. This will enable consumers to freely use these technologies and will draw attention to this possibility. In addition, consumers should be proactively protected against the loss of access to their accounts for potential violation of platforms' terms of service while using anti-FR tools such as Fawkes. We further recommend the Commission move to protect companies and individuals providing anti-surveillance technology, given the precedent for companies benefiting from tracking and surveillance taking legal action against those providing tools to stop it. Such actions can chill innovation in this important space.

Finally, we note that some anti-FR tools rely on a 'clean slate' approach[48]. This means that, for these tools to work well, consumers must control which images of them appear in FR model training datasets[49]. This can only occur if regulation compels commercial FR systems to be re-trained from scratch, after obtaining *consent* to use all images in their training dataset. Currently deployed models are often trained on images scraped online without user consent (see Observation 1), preventing a user from ever having a "clean slate" and using related anti-FR tools. Requiring retraining after consent is obtained would allow users the clean slate they need to use anti-FR tools against these models if they so choose.

**Recommendation 6**: *The FTC should require existing facial recognition models to be retrained* after *ensuring all training data was obtained with consumers' consent.*

The prevalence of FR systems that use data obtained without consent or the ability for consumers to use anti-FR tools to prevent identification necessitates rule-making that requires companies to start from a 'blank slate'. The Commission should enact regulation to ensure that all surveillance systems in use have been trained with data obtained with consent, and require companies to open existing ones up for audit to ensure that the data used was obtained with permission.

[46] L. Dax. Cyberdazze hyperface 1 anti recognition shield. *Zazzle*, 2022

[47] B. Williams. Five and oh ... look, another lawsuit upholds users' rights online. *AdblockPlus Blog*, 2016

[48] e.g. Fawkes—S. Shan, E. Wenger, J. Zhang, H. Li, H. Zheng, and B. Y. Zhao. Fawkes: Protecting Privacy Against Unauthorized Deep Learning Models. In *Proc. of USENIX Security*, 2020

[49] E. Radiya-Dixit and F. Tramèr. Data poisoning won't save you from facial recognition. In *ICLR*, 2022