

# Privacy of Groups in Dense Street Imagery

Matt Franchi\*  
mattfranchi@cs.cornell.edu  
Cornell University, Cornell Tech  
New York, New York, USA

Severin Engelmann  
severin.engelmann@cornell.edu  
Cornell University, Cornell Tech  
New York, New York, USA

Hauke Sandhaus\*  
hgs52@cornell.edu  
Cornell University, Cornell Tech  
New York, New York, USA

Wendy Ju  
wendyju@cornell.edu  
Jacobs Technion-Cornell Institute,  
Cornell Tech  
New York, New York, USA

Madiha Zahrah Choksi  
mc2376@cornell.edu  
Cornell University, Cornell Tech  
New York, New York, USA

Helen Nissenbaum  
hn288@cornell.edu  
Cornell University, Cornell Tech  
New York, New York, USA



**Figure 1: AI-inferred group membership in a dataset of more than 25 million facially de-identified dashcam images from NYC in 2023. A project website is available at [dsi.tech.cornell.edu](https://dsi.tech.cornell.edu).**

## ABSTRACT

Spatially and temporally dense street imagery (DSI) datasets have grown unbounded. In 2024, individual companies possessed around 3 *trillion* unique images of public streets. DSI data streams are only set to grow as companies like Lyft and Waymo use DSI to train autonomous vehicle algorithms and analyze collisions. Academic researchers leverage DSI to explore novel approaches to urban analysis. Despite good-faith efforts by DSI providers to protect individual privacy through blurring faces and license plates, these measures fail to address broader privacy concerns. In this work, we find that increased data density and advancements in artificial intelligence enable harmful group membership inferences from supposedly anonymized data. We perform a penetration test to

demonstrate how easily sensitive group affiliations can be inferred from obfuscated pedestrians in 25,232,608 dashcam images taken in New York City. We develop a typology of identifiable groups within DSI and analyze privacy implications through the lens of contextual integrity. Finally, we discuss actionable recommendations for researchers working with data from DSI providers.

## CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy**; *Penetration testing*; • **Computing methodologies** → *Object recognition*.

## KEYWORDS

privacy, dense street imagery, group privacy, contextual integrity, computer vision, surveillance, penetration testing, auditing

## ACM Reference Format:

Matt Franchi, Hauke Sandhaus, Madiha Zahrah Choksi, Severin Engelmann, Wendy Ju, and Helen Nissenbaum. 2025. Privacy of Groups in Dense Street Imagery. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3715275.3732185>

\*The authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

FAccT '25, June 23–26, 2025, Athens, Greece

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1482-5/2025/06...\$15.00

<https://doi.org/10.1145/3715275.3732185>

## 1 INTRODUCTION

Dense Street Imagery (DSI) represents a major breakthrough in vehicle hardware, imaging technology, and networking, enabling dynamic, real-time depictions of locations worldwide. Unlike traditional static snapshots, such as those provided by Google Street View, DSI capitalizes on innovations in *temporal* density—achieved through networked dashcams [42] and advanced driver-assist systems [59] to deliver fresh, continuous imagery at an unparalleled frequency. Researchers have found many beneficial applications of DSI; these include tracking transient weather events [43, 103], documenting sidewalk scaffolding [95], analyzing vehicle placement patterns [45], and assessing dynamic street conditions [31, 44]. By transforming mobile cameras into a distributed sensing network, DSI offers researchers new ways to understand the rapidly changing physical and social landscapes of urban environments in ways that are more flexible and adaptable than possible with location-fixed sensing systems.

The shift from periodic to near-continuous image capture has made it possible to monitor people at much shorter intervals; cycles that were once difficult or even impossible to observe. Unlike Google Street View (GSV), which faced several privacy controversies in the early 2000s, DSI has so far avoided similar public scrutiny. In the 2014 Canadian case *Pia Grillo v. Google*, the plaintiff sued Google for invading her privacy after Google Street View (GSV) published an image of her outside her home in Quebec [21, 47, 65]. Although GSV had blurred her face, it failed to obscure her license plate and home address, and the photo revealed part of her upper body. The court found that these visible details could allow others to identify her, despite the facial blurring. The case highlights how peripheral information such as license plates and addresses can undermine anonymity, even when facial features are obscured.

Today, an expansive infrastructure has emerged in which companies like Mobileye manage vast datasets (reportedly 200 petabytes of data) [77], comprising over 3 trillion images from cameras mounted on consumer vehicles [56]). Lyft, a leading ride-sharing company, also recently gained access to *200 million miles* of driving imagery [56]. For reference, Google reported its Street View had 220 billion images and 10 million miles of footage in 2022. DSI has, without much notice, assembled a data moat more than ten times the size. Recent research by Sandhaus et al. [93] reveals that autonomous vehicle companies possess vast amounts of DSI data; while they are reluctant to share it openly, they have sophisticated internal methods to remotely retrieve data from their fleets. How can individuals escape the surveillance potential inherent to DSI? Increasingly, it seems that to opt out, one must opt out of public space [53]. Privacy defenses for sensitive objects in DSI include the blurring of faces, license plates, and other user-requested objects. Such practices have precedents in earlier technologies like Google Street View [6]. Google states: “We have developed cutting-edge face and license plate blurring technology that is designed to blur identifiable faces and license plates within Google-contributed imagery in Street View” [52]; this seems to be the commercially-standard protection standard [46]<sup>1</sup>. Privacy-preserving mechanisms

for static street view technologies, including blurring faces, license plates, and other user-requested objects, are inadequate and notably blunt [20], and obfuscation failure modes are noted and exist [20] – see Appendix C for details on specific failure modes. The inevitable increase in *temporal* density further undermines these established mechanisms [42]. Ultimately, in DSI, objects of interest are *traceable*. Additionally, artificial intelligence (AI)’s inferential capabilities make it possible to generate detailed information without direct collection, for example, using vision models to analyze visual data and identify individuals’ clothing types, styles, and accessories in public spaces in near real-time [18, 26]. This allows identification of group affiliations (e.g., demonstrators, religious congregations, professionals) based on attire and accessories. Even without explicit group markers, physical proximity to others displaying group affiliation can signal group association [89]. In public spaces, computer vision models can infer protected attributes like gender through pose estimation [63] and disabilities using proxies like wheelchairs [92]. Moreover, cross analyses with mobility data or activity data, such as pings of individuals’ cell phone location across time, further motivate privacy risks in DSI. A landmark study analyzing mobility data for 1.5 million individuals over 15 months revealed that just four spatiotemporal points were sufficient to uniquely identify 95% of individuals in the dataset [32]. These results have been replicated in subsequent studies (e.g., [33]) and highlight the significant surveillance potential of DSI. Why? Data of this nature is highly sensitive due to the uniqueness of human behavior [53]. DSI – with its visual dimension – reveals a heightened sensitivity, enabling AI to make inferences about appearance and behaviors.

This work pioneers the exploration of group identifiability in public street imagery, leveraging a real-world dataset of 25,232,608 unique dashcam images captured in New York City (visualized in Figure S3) provided for research evaluation by Nexar, Inc., a dashcam manufacturer and smart-mapping startup. We note that this work does *not* address the potential identifiability of individuals. Our study intersects penetration testing, group privacy, and contextual integrity to investigate DSI’s implications for society. We begin with a penetration test of a real-world DSI dataset to demonstrate how face-deidentified imagery can be circumvented with ease, revealing artifacts that can lead to group privacy harms. The results of our penetration test motivate our downstream research questions.

Next, we present related work, including an overview of DSI-producing technologies, relevant privacy theory, and examples of inferences produced by computer vision models. We then demarcate information flows within DSI, using the framework of contextual integrity. Finally, we discuss and synthesize findings, document harms, and offer recommendations for DSI data sharing and use within academia.

## 2 PENTESTING A DATASET OF DSI FOR RISKS OF PRIVACY HARMS

We begin by demonstrating how efficiently DSI can be integrated into an application capable of inferring membership in a *group*, all while preserving individual privacy through de-identification.

<sup>1</sup>We note that at the time of writing, Nexar, the provider of our experimental dataset, goes beyond the popular commercial standard by blurring entire pedestrian figures instead of just faces (see Figure S1 for an example). However, this practice reduces

the visual fidelity and utility of certain pedestrian-dense images, presenting an open problem.

How do we define ‘de-identification’? In the context of DSI, de-identification refers to the visual blocking of elements considered sensitive, such as blurring all faces in a dataset. DSI data providers disproportionately rely on *facial blurring* as a privacy measure because of its simplicity and extendability across diverse contexts [20]. Building on this, we add that individual, pedestrian-level obfuscation is not enough privacy protection. We examine this assertion through an approach borrowed from computer security literature, called *penetration testing*, or ‘pentesting’ [16]. Penetration testing typically involves breaching a system to assess the difficulty of doing so. However, Bishop [16] emphasizes that it should also include a thorough analysis of threats and potential attackers, an approach that aligns with our later examination of DSI information flows through the lens of contextual integrity. Rather than assuming the absence of privacy harms, our work actively identifies them [16].

## 2.1 Adversarial Methodology

In our penetration test, we define an adversary as an ‘authority,’ such as law enforcement or government officials. Our experiments utilize a comprehensive dataset of 25,232,608 images collected across New York City (detailed sampling methodology in Appendix B). From an adversary’s perspective, we identify several methods using DSI imagery—enabled by recent advances in artificial intelligence—that could potentially harm a group or its members. These methods are outlined in Table 1 and intended to be illustrative rather than exhaustive.

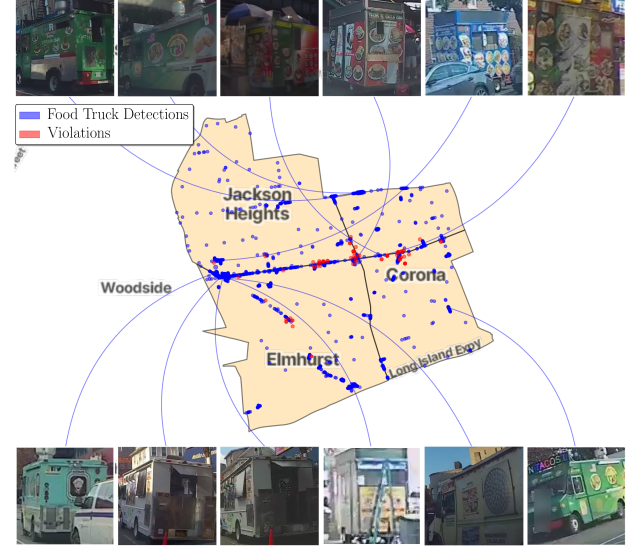
*Experiment 1.* The first experiment starts with the *zero-shot* method described in Table 1. To source training data, we conducted a zero-shot image classification task on 500,000 randomly sampled images, leveraging vision-language models (VLMs) [87, 109]. Specifically, we used the VLM Cambrian-13B [102] to answer the prompt: “Is there a food truck in this image?”, receiving yes or no answers. Next, we manually validated the classified positives through human annotation<sup>2</sup>. We also report standard model performance metrics in Section B.1.4. Finally, we trained a series of lightweight YOLOv11 ([62]) object detection models and selected the most performant. These models, capable of real-time and distributed inference, were chosen to illustrate the ease with which imagery can be transformed into spatiotemporal distributions of inferred group members. Lastly, we estimated the spatiotemporal distribution of each group in the entire dataset of 25,232,608 images by running inferences on each image with the trained YOLO model. We provide more information on the training of the YOLO model in Section B.1.2.

*Experiment 2.* The second experiment more directly encompasses the *zero-shot* method from Table 1. Similarly, we run a zero-shot image classification task on 500,000 randomly sampled images using the same VLM (Cambrian-13B). For this task, we asked Cambrian: “Is there a bike rider with a box on their back in this image?”<sup>3</sup>. Then, we took the Cambrian model output as ground-truth and created a

<sup>2</sup>This task was carried out by a team of human annotators, including two authors of this paper, both with extensive experience observing the cultural norms and street activity of New York City.

<sup>3</sup>We experimented with several different prompts on a small sample of randomly sampled images and found that Cambrian has little predictive power on domain-specific terms like ‘food delivery worker’ or ‘Uber Eats driver’. Consequently, we prompted for the flagship equipment that food delivery workers wear while biking around the city: food storage boxes strapped onto the back of a bike.

detection heatmap (see Figure 3). We draw parallels between this approach and the biased, partially inaccurate machine learning models that have been deployed in algorithmic policing endeavors [14, 71, 91]. In the following, we provide information on mobile food vending and food delivery in New York City.



**Figure 2: A map showing reported vending violations against food truck detections in Jackson Heights, Queens. For higher precision, we choose a confidence threshold of 0.7, which yields a precision of 0.90 and a recall of 0.50 on the test set.**

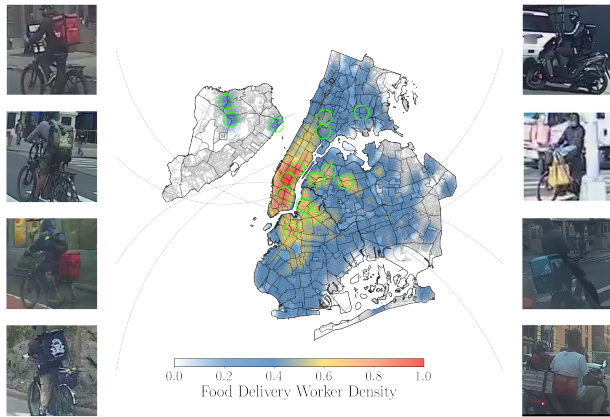
We also generate inferences on several other group identities using the zero-shot method from Table 1, including nurses, religious-presenting individuals, dog owners, police officers, and protesters, all shown in Figure 1.

**2.1.1 Pentest: Mobile Food Vending and Delivery in New York City.** As of October 2024, Hunter College’s Food Policy Center reports that the NYC Department of Health has issued around 4,600 food cart permits and about 500 food truck permits [84]. Cart permits are in high demand: Thousands of individuals are currently on the NYC Department of Health permit waiting list [99]. The current population of food delivery workers in New York City is estimated to be between 65,000 and 122,000, many of whom are immigrants [2, 9]. Given their vulnerability in real-world scenarios, we focus our exploration on mobile vendors and delivery workers, a group frequently targeted by adversaries. In April 2024, a group of vendors in Brooklyn protested a surge in ticketing, alleging unfair targeting by the NYPD [99]. Just a few months later, in July, police detained a vendor who collapsed while handcuffed and required hospitalization [99]. In 2023, the NYPD took street vendors to criminal court almost six times as frequently as in 2019, issuing over 1,200 criminal summonses [28]. Food delivery workers are also targeted by various groups. As early as 2021, The New York Times reported on the theft of workers’ e-bikes (valued at up to 3,000 USD), highlighting the need for self-organization [41, 74, 86]. More recently, the NYPD has



Method	Description	Example
Zero-shot	Inferring images with a foundation model or machine learning model, and treating its output as ground truth, without any task-specific labeled data.	Using a web-trained model to identify group instances, e.g., by describing visual characteristics like "umbrellas" or "LED signs."
Supervised	Human generation of training and validation labels, training a model on these annotations, and treating its output as ground truth.	Crowdworkers annotate images with labels like "protesters" or "food vendors," train a model on these labels, and then classify new images.
Unsupervised	Generating image embeddings from a model, clustering them, and manually identifying labels for the various clusters.	Cluster image embeddings, label visual clusters (e.g., "colorful trucks"), and apply these labels to all images in the cluster.
Known-event matching	Retrieving images occurring at the same location and time as a known event, and using the image's contents as additional context for an adversarial task.	Manually matching images to a known protest time and location to interpret the scenes and identify groups.
Geofencing	Retrieving images within a geographic region of interest, and using the contents of those images as additional context for an adversarial task.	Focusing on a city block known for street vendors, labeling initial images to define key features, and then applying those features to identify similar vendors in the area.

**Table 1: DSI group identification and retrieval methods. We perform experiments based on zero-shot methods.**



**Figure 3: Using zero-shot image retrieval, we queried Cambrian for the prompt “Is there a bike rider with a box on their back in this image?”. An authority may readily use this to create a strategic map for deployment zones optimal for monitoring food delivery worker hotspots, as depicted. The computed hotspots correspond to the average ‘lunch rush’ period (10AM-2PM) and can be easily computed over each day within our dataset. From a ground-truth annotation of 500 random positive detections, we estimate precision at 0.70.**

cracked down on illegal mopeds, with much of the focus directed at delivery workers [57].

The 23,000 street vendors in New York City are a particularly vulnerable group with the majority of them vending as their primary source of income. As previously mentioned, these are predominantly immigrants (96%), and most of them operate in legally gray areas [94]. Indeed, around 75% of mobile food vendors have no permit [94]. While providing a valuable service to New Yorkers by offering fast, convenient, and affordable goods, these small businesses

operate in a precarious environment, their continued existence reliant on the challenges faced by enforcement agencies. While exact demographic information on food delivery personnel is opaque, we know that food delivery apps like Uber Eats and DoorDash rarely comply with governmental requests for information [2].

Our pentest investigates whether adversaries, such as police, can gain cheap, easy, and more direct access to information about where and when vendors operate using DSI datasets, even when individuals are de-identified. Such information poses a significant threat to the livelihood of these vendors, who already face minimal job security and ongoing concerns about policing [24]. Thus, significant threats to privacy result from the identification of vendors at the group and community level.

**2.1.2 Pentest Execution.** Having acquired data representing inferred detections of mobile food vending and food delivery workers, we now behave as an adversary might, and attempt to measure the spatiotemporal distribution of these groups in New York City. We labeled the training data needed for the object detection model in 4 hours. We trained a high-performance, convergent model in 44 hours. Using this model, we were able to infer 192 images per second on a single GPU, processing the entire dataset of 25,232,608 images in just 36 hours. Under an optimal confidence threshold of 0.205, the model asserted 196,183 images depicting food trucks. We show a zoom-in of the Jackson Heights area of Queens in New York City (see Figure 4 in Supplement), a known hotspot for unlicensed vending [3].

**2.1.3 Pentest Findings.** From Experiment 1 we find that out of all studied food truck vending violations<sup>4</sup>, the median distance to the nearest high-confidence model detection is only 127 feet and gets as near as 36 feet. We find a clear visual overlap between known food truck vending violations and our high confidence detections, shown

<sup>4</sup>We assemble a list of all food truck vending violations during our dataset’s coverage period from NYC OpenData, specifically the NYC Office of Administrative Trials and Hearings case status dataset [85].



in Figure 2. This means that, for areas with high concentrations of food truck vending violations, there is more than ample imagery with which an adversary could pursue remote inspections, similar to the methodology in [95].

From the detections of food delivery workers in Experiment 2, we are also able to easily create targeted deployment zones for the in-the-wild surveillance of food delivery workers, shown in Figure 3. These results demonstrate the ease of developing a useful attack tool against groups with DSI, even under commercially-standard de-identification.

We have shown the *ease* of crafting useful attack tools. The question remains as to *what* research applications of DSI are ethical to pursue. We tackle this question in the paper’s latter sections, using the contextual integrity framework. We now present a background and related work section, and then move to our analysis of information flows in DSI, followed by discussion.

### 3 RELATED WORK

This section describes related and contextual work, including an overview of DSI-producing technologies, the inferential power of computer vision models, group privacy theory, and the framework of contextual integrity.

#### 3.1 DSI-producing Technologies

Earlier sensing technologies like Google Street View (GSV) [6] have been used to computationally characterize longer-term societal processes, including gentrification [60, 101], street safety [78], health outcomes [79], shade from street trees [101], and demographic distributions [48]. People and people-descriptive objects are visible in these images. Advances in tools that collect street view imagery have attracted considerable scrutiny, notably as facial features [50], vehicle license plates [35], and individual homes [35] become characterizable. As empirical research using GSV has evolved, so has research focusing on pedestrian obfuscation [39, 67, 83, 104] and commentary on the ethical use of GSV [10, 58].

Street view imagery has been used in attempts to characterize short-term processes, such as counting public pedestrians [22, 110] and alcohol consumption [29]. However, the temporal variability of GSV, averaging 7 years on a controlled, short-route surveying study [66], means it is unsuited for capturing short-term, real-time or near real-time phenomena related to groups. Audit works at FAccT have shown the insuitability of GSV and similar technologies for measuring abstract concepts like ‘livability’ [5], particularly when human annotators are involved [111].

We qualify DSI-producing technologies as those that produce spatially representative data at a frequency sufficient to capture short-term phenomena including people and their behaviors. DSI-producing technologies include networks of static traffic cameras [34, 97] and collections of mobile cameras (“dashcams”), either on private [42, 45] or public [90, 93] fleets. The sensing capability of a DSI-producing technology can be evaluated quantitatively as the fraction of all possible space-time pairs with some number of images. DSI can be deployed to fingerprint phenomena at very high granularity [42]. As an example, DSI permits the continuous monitoring of foot traffic patterns on a local sidewalk network at 15-minute increments [42].

In the following, we provide background on the key privacy concepts we apply for our analysis of DSI. To account for the complexity of privacy vulnerabilities in DSI, we ground our analysis along three core concepts: inferential models, group privacy, and contextual integrity (CI).

#### 3.2 The Inferential Power of Vision Models

Ascendant inferential capabilities of AI models, particularly computer vision models, pose a significant threat to group privacy in DSI. AI’s capacity to draw inferences enables the extraction of specific information of interest without directly collecting it. Computer vision AI ‘makes sense’ of a sea of visual data [15] by drawing inferences from it: vision models in robotics [23], self-driving vehicles [61], and emotion recognition systems [55] analyze the semantics of images and videos—millions of which are generated daily across different digital socio-technical systems, including the DSI. AI inferences present significant challenges to conceptions of privacy, both in theory and in data protection practice. The privacy conceptions most susceptible to erosion by AI’s inferential power are likely those grounded *exclusively* in categorical distinctions – such as classifying data as sensitive or non-sensitive – while at the same time uncritically accepting AI-generated inferences as inherently valid [37]. Critical data scientists, particularly members of the FAccT community, have demonstrated the adverse impacts of invalid inferential models, especially due to biased misrepresentations that result in improper and unfair predictive descriptions of individuals they cannot understand, correct, or control (e.g., [36, 38, 49, 51, 98, 105, 106]).

Although blurring may prevent the processing of facial data from individuals, such approaches do not fully protect against the inferential power of computer vision models in identifying group-relevant attributes. DSI serve both as a training ground for model development and as a deployment environment for trained models. Inferential models can leverage the semantics inscribed in urban centers, buildings, and public squares, which manifest in corresponding scripts of activities, behaviors, and roles. People may gather for protests in public squares, work in specialized buildings such as hospitals or construction sites, or prepare to engage in religious practices outside designated spaces such as temples, churches, synagogues, or mosques. Models can detect such environmental cues and infer sensitive information about facially obfuscated individuals by analyzing group membership proxies such as clothing, accessories, and behaviors. Computer vision models effectively analyze visual data, enabling the inference of clothing types, styles, and the presence of accessories such as bags, hats, glasses, and jewelry [18, 26]. This capability facilitates real-time identification of individuals belonging to specific groups, including demonstrators, religious congregations, or professionals (e.g., doctors and nurses), based on distinctive attire and accessories. Even when individuals do not explicitly display group affiliation through their own clothing or accessories, their physical proximity to those who do can result in their association with the group—a phenomenon often referred to as the “lookalike” effect [89]. In public spaces, computer vision models can also infer protected attributes such as gender through pose estimation [63], or physical and mental disabilities through proxies such as wheelchair or white cane use [92]. When

deployed on the DSI, AI's purported inferential power generates unprecedented privacy challenges, turning public spaces into arenas for real-time, automated meaning-making that facial or body obfuscation alone cannot prevent.

### 3.3 Privacy as Appropriate Information Flow vs. Privacy as Preference

In motivating our choice to utilize the contextual integrity framework in our analysis of information flows in DSI, we now detail two prominent takes on privacy: privacy defined by group preferences (aligning with the theory of group privacy [100]), and privacy defined by appropriate information flows (aligning with the framework of contextual integrity [80]).

**3.3.1 Group Privacy.** Group privacy has been the subject of scholarly discussion since the late 1990s, when researchers began examining how new information technologies classify individuals according to shared attributes rather than treating them as isolated subjects. In his prescient work, Vedder [108] introduced the concept of “categorical privacy” to address how data mining techniques lead to the “deindividualization of the person,” where judgments about individuals are based on group characteristics rather than individual merits. This early Dutch work on group privacy emphasized that privacy concerns extend beyond individual data subjects to entire classes of people, anticipating the challenges posed by modern big data analytics.

Building on this foundation, Taylor et al. [100] define group privacy as the collective ability of a group to control its personal and shared information. In practice, this involves protecting both individuals' information within the group and safeguarding shared information pertaining to the collective entity itself. This definition recognizes that groups, as distinct entities, have privacy interests beyond those of individual members, considering the shared norms and values that shape collective privacy practices [40].

Bloustein and Pallone [17] further developed the concept of group privacy as protecting confidential information shared among two or more individuals against external parties. They outline the “right to huddle,” referring to a group's ability to gather and communicate confidentially within their own boundaries, enabling groups to maintain trust, collaboration, and collective decision-making without undue external surveillance or interference. Within our paper, we specifically focus on normative groups observable in DSI, which we discuss in detail in Table 2.

Loi and Christen [70] distinguish between two concepts of group privacy: “what happens in Vegas stays in Vegas” privacy, concerning confidential information shared within a group, and “inferential privacy,” dealing with inferences about groups defined by shared features. Our pentest results directly implicate this second form of privacy, showing how readily group memberships can be inferred from seemingly anonymized DSI data. Mantelero [73] further argues that in the context of big data analytics, privacy and data protection should be considered collective rights rather than purely individual ones.

van der Sloot [107] explores whether groups should have a right to protect their group interest in privacy, noting that while privacy rights have historically focused on individuals, contemporary technological paradigms like big data present threats that materialize at

group levels rather than individual ones. As Asgarinia [7] argues, the traditional focus on individual privacy rights fails to address the vulnerabilities of “clustered groups” designed by algorithms, where information about the group can be used for harmful purposes even when individual members remain anonymous. These perspectives highlight the inadequacy of individual-focused privacy frameworks in addressing collective privacy challenges posed by modern data analytics.

As our penetration test demonstrated, traditional approaches to privacy that focus on individual anonymization fail to prevent group-based privacy violations in DSI. While individual privacy frameworks might emphasize control and consent, they prove inadequate in addressing the collective, inferential privacy challenges posed by dense spatial-temporal imagery and AI analysis. This limitation points to the need for a more holistic framework that can address both individual and group privacy concerns in the context of DSI.

**3.3.2 Contextual Integrity (CI).** The inferential reality of computer vision AI models and real-time DSI produce privacy vulnerabilities for groups in public spaces. For the purposes of our work, Nissenbaum's theory of *Privacy as Contextual Integrity (CI)* helps distinguish legitimate from illegitimate information flows according to contextual norms for such groups [27]. Drawing on social theory, social philosophy, and law, CI conceives of social life as comprising distinct social domains or *contexts*, such as commerce, education, finance, healthcare, civic life, family, and friends [80]. A CI context is ultimately defined by its ends, aims, or goals, which further determine its role in society at large, as well as its values, be it equality, justice, or individual autonomy, among others. As such, in a healthcare context, for example, the goal or aim is to cure and prevent illness, alleviate pain, and there is a commitment to ethical values such as equity and patient autonomy. The precise composition of ends and values may differ across societies, and may even be open to political contestation, e.g., in an education context, it is open to debate whether the goals are to enlighten or train, to teach rote skills or encourage creativity, or to generate workers as opposed to enable a responsible citizenry. CI shifts away from notions of privacy as information control or secrecy, and conceives of privacy as the appropriate flow of information: flow that conforms with contextual informational norms. Contextual informational norms define acceptable data practices and may range from implicit and weak—social disapproval of friends betraying confidences—to explicit and embodied—laws protecting journalists refusing to name sources or requiring physicians to maintain the confidentiality of health data. A complete statement of a contextual informational norm provides values for five parameters: data subject, data sender and data recipient (collectively referred to as *actors*), information type (topic or attribute), and transmission principle (the conditions under which information flows) [80–82].

Actors (subject, sender, recipient) are labeled according to contextual capacities or *roles*, such as physicians, educators, or political figures. Information types are defined according to contextual ontologies, such as an educator's reports about a student's learning progress in an educational context. Transmission principles are the conditions or constraints under which a particular information type flows from senders to recipients. Transmission principles include

Table 2: Typology of Identifiable Groups

Group Type	Description	Examples
<b>Self Organized*</b>	Groups formed voluntarily by members who share a common purpose, set of values, or specific goals. These groups often emerge organically through shared interests or collective aspirations, and function independently of external mandates or authority.	Protesters organizing for a cause, local community action groups, religious (i.e. church or temple) groups.
<b>Role-based*</b>	Groups composed of individuals acting within defined roles tied to their social, professional, or communal responsibilities. Membership is typically based on one's job, function, or societal duty rather than personal traits or voluntary affiliation.	Nurses, police, construction workers, tourists, school children.
<b>Clusters</b>	Spatially proximate groups of individuals who are temporarily assembled or gathered in a shared physical location, often without preexisting social connections or enduring relationships. Such groups are usually context-dependent and formed by situational proximity rather than shared purpose.	People waiting at a train station, individuals standing in line, commuters at a crosswalk.
<b>Attribute-based</b>	Groups organized around intrinsic, often immutable characteristics or shared traits of individuals. Membership in these groups is typically defined by demographics or identity markers, which may influence societal perceptions.	Age groups (e.g., seniors, children), gender-specific communities, racial or ethnic groups, individuals with disabilities.

\* Indicates normative groups.

confidentiality, reciprocity, consent or mandated by law, among others. CI (and therefore privacy) is achieved or preserved if all information flows within a particular context align with entrenched informational norms. Hence, to determine the appropriateness of an information flow, one must determine all five parameters characterizing such flow.

Unlike privacy-as-preference approaches, which focus primarily on individual control and consent, contextual integrity offers a more comprehensive framework for analyzing the group privacy harms revealed in our penetration test. CI enables us to evaluate the appropriateness of information flows in DSI by considering not just who is depicted, but how that imagery is collected, processed, and shared within specific contexts. This makes CI particularly suited to addressing the complex privacy challenges posed by DSI, where individual de-identification proves insufficient to protect group privacy. In the following section, we apply the CI framework to demarcate appropriate and inappropriate information flows in DSI based on our empirical findings.

#### 4 APPLYING CI: DEMARCATING INFORMATION FLOWS IN DSI

Motivated by the group identifiability risks posed in commercial DSI data [30, 69], even under intense de-identification, we provide recommendations rooted in a more holistic approach. Under the contextual integrity framework [13, 80], information flows consist of five components: a subject, a sender, a recipient, an information type(s), and a transmission principle. In applying the CI framework, we dissect information flows within DSI and outline the typical values assigned to each of the framework's five components.

The **subjects** in DSI are *groups* of pedestrians. We delineate the typology of identifiable groups in DSI in Table 2. In many cases, the **sender** is the *data provider*. While a vehicle driver is also complicit in the act of the data capture, it is the data provider who makes the decision of when to take an image, how many images to take, and

how many images to upload for downstream transfer. Alternatively, an adversarial machine learning model may also assume the role of data provider if it sends its generated outputs to a recipient. Subsequent analysis, such as drawing inferences on top of a given dataset, create a novel information flow, and typically involve a new sender, such as an academic researcher or organization.

The **recipient** is variable. In most research projects that use DSI to date, the recipient is a *research group*. DSI can also have commercial uses, in which case the recipient is a *private company*, or a *public sector agency*. In DSI, **information types** are *photographs* with attached geospatial telemetry data. This combination creates a record of a group's (or group member's) location and the time they were situated there. DSI involves a **transmission principle** where the subjects are not required to give consent and have no right to revoke the transmission unless they preempt the information flow via requesting an obfuscation of their appearance in the dataset.<sup>5</sup> As we demonstrate in Figure 4, the contextual integrity framework provides a structured approach to evaluate whether information flows in DSI respect privacy norms by examining the complete five-parameter tuple rather than isolated elements.

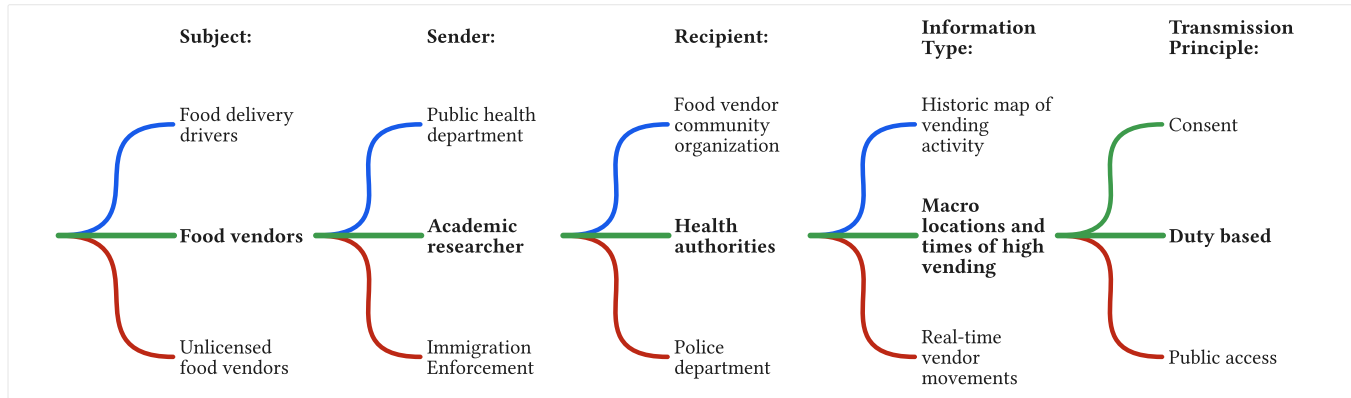
##### 4.1 An Inappropriate Data Flow in DSI

According to CI [81], data flows that breach socially accepted norms are considered 'inappropriate.' DSI, in combination with automated image retrieval through algorithmic methods (described in Table 1), radically disrupts many contextual information norms, similar to the advent of big data technologies [68].

As an extension of our penetration test, we consider food delivery workers as a data subject. Prior to the introduction of street view and DSI technologies, encounters of food delivery workers and authorities in urban environments were ephemeral. Records of

<sup>5</sup>Google supports identifiable content blurring in Street View [1], but not removal. The European Union's GDPR framework advances the right to removal through its articulation of the "right to be forgotten." [88].





**Figure 4: Contextual Integrity Analysis of a DSI Information Flow.** Changing a single parameter in an information flow can transform it from appropriate (green,  $\rightarrow$ ) to inappropriate (red,  $\rightarrow$ ). Contextual integrity requires evaluating fully specified information flows to avoid ambiguous cases (blue,  $\rightarrow$ ).

food delivery were noticed by the local community or public, in journalism, media, and writing, or within individually distributed photographs. In these settings, food vendors could reasonably react to the presence of authority, and communicate in person.

In this context, we conceptualize the city as the recipient (or adversary) of the data, with the data provider acting as the sender. The data itself consists of images enriched with computational meta-data, including precise latitude and longitude coordinates, as well as the exact date and time of capture. Crucially, this transmission is sufficient to track a vendor's activities over time or issue fines for operating without a license. The underlying transmission principle assumes that vendors cannot revoke the transfer of this data, nor are data providers required to inform them that they have been recorded.

A proponent of DSI surveillance might argue that DSI providers must give the police persistent access to food vendors' vending locations and movement throughout the city upon request. While this information flow may, at face value, appear morally justifiable, it causes direct privacy harms to food vendors. Such surveillance exposes vendors to risks of increased fines, job loss, or even the inability to continue their work, threatening their livelihood and way of life. This impact is particularly severe for undocumented immigrants, for whom food vending is not only a direct lifeline to sustain their families [74]. Disrupting this flow has moral implications such as economic survival, access to low-cost food for others, and respecting the right to work without unnecessary or harmful interference.

## 4.2 Inappropriate Flows in Other Groups

We highlight examples of inappropriate information flows that arise as DSI technologies become more pervasive. These examples, some of which are illustrated in our dataset (Figure 1), offer a glimpse into the many ways DSI and inferential models can violate established privacy norms, posing risks to the interests and values of individuals and groups in public spaces. Using CI, we trace the following components of each information flow: (1) the subject, (2) the sender, (3) the recipient, (4) the information type(s), and

(5) the transmission principle. We specifically locate groups in the DSI imagery that pose contextual integrity *harms* from merely being present in a geo-temporally tagged image. That is, there are documented, real-world instances of a group being targeted based on their public situation in space and time. Below we describe several complete flows for groups from the typology, involving plausible harms to those groups (this information is also presented in a table in the appendix Table 2).

*Protesters and demonstrators:* Protesters' meeting locations and times, captured in geo-tagged images by DSI providers, may be surreptitiously shared with political groups opposed to their cause. This information flow enables adversarial groups to target and disrupt peaceful assemblies, undermining the protesters' right to organize and express dissent. The resulting harm includes exposure to retaliation, suppression of free speech, and the erosion of democratic principles. *Workers and employees:* In professional settings, shift patterns of nurses inferred from DSI images captured outside healthcare facilities could be sold to exploitative employers or staffing agencies seeking to take advantage of their availability. Such data flows enable adversaries to target nurses with aggressive recruitment tactics, pressure them into accepting poorly compensated shifts at inconvenient hours, or manipulate them into working under unsafe or undesirable conditions. This undermines nurses' autonomy in the workplace and their ability to make independent decisions about their labor. *Pedestrians:* Geo-tagged imagery of people gathered at busy crosswalks during peak hours may be sold by predatory advertisers seeking to exploit behavioral patterns. For example, advertisers might use this data to push high-pressure marketing campaigns for products like payday loans or fast food, directly targeting the commuters' mobile devices in certain locations or interactive billboards. *Religious groups:* Muslims or Hasidic Jews, photographed in locations perceived as incongruent with their practices (e.g., near entertainment districts), risk having these images and their associated metadata shared with hate groups, potentially exposing them to targeted harassment or discrimination. This flow facilitates targeted harassment, stigmatization, and even violence against the group, violating societal norms of religious freedom.

### 4.3 Appropriate Data Flows in DSI

While many data flows enabled by DSI can lead to privacy violations, this technology also has beneficial uses, and privacy can simultaneously be protected when contextual norms are respected. We briefly outline two examples of appropriate information flows that align with societal expectations and provide public benefit.

*Urban planning for pedestrians:* City planners can use DSI to study pedestrian movement patterns at crosswalks and intersections to optimize traffic signal timing and improve infrastructure. Transit authorities can analyze DSI, showing commuter congregation patterns at bus stops and train stations to adjust service frequency and capacity.

*Crisis support for homeless:* Non-governmental organizations receive DSI to study the temporal movement of encampments. This allows crisis teams to provide on-demand support to those communities in need.

Figure 4 illustrates how contextual integrity operates in the scenario described throughout this analysis. The center pathway represents an appropriate flow: academic researchers sharing macro-level vending patterns with health authorities under duty-based principles. However, altering a single parameter, such as changing the recipient to law enforcement or modifying the information type to real-time movements, transforms the flow into an inappropriate one that violates privacy norms or an ambiguous one. The transmission principle outlines conditions under which data is obtained, used and reused, for example, the same vending location data that enables public health planning becomes problematic when made publicly accessible without restrictions.

These examples demonstrate how DSI can be deployed ethically when the information flow of DSI and its inferences are fully specified according to CI, as well as when: (1) The purpose serves a clear public benefit, (2) Usage is limited to the legitimate use purpose, (3) Access is restricted to appropriate and previously agreed parties, (4) Affected groups are given agency in how their information is collected and used, (5) Focus is restricted to those groups by limiting inferrable information.

The distinction between appropriate and inappropriate information flows hinges not only on the type of data collected or who the data subject is, but on the full specification of the flow across all contextual integrity parameters. To ensure DSI technologies are used responsibly, robust safeguards and ethical guidelines must be implemented to protect privacy while enabling beneficial applications.

## 5 DISCUSSION & CONCLUDING REMARKS

We conclude by returning to the targeted surveillance of mobile food vendors in New York City. Combined with findings from our penetration test, we demonstrated that DSI can effectively be used to identify contextual information about group distributions and facilitate harm against group members. The object detection model generated (photograph, place, time) tuples for tens of thousands of food trucks from any given set of images. At present, the most realistic threat model for these outputs is largely *curation*. Authorities, however, with additional information on food truck congregations and outlying locations, can ramp up targeting in those areas, likely resulting in increased fines and summonses [99].

### 5.1 Issue of Reducing Privacy to Anonymity

This work challenges the assumption that anonymizing individuals within DSI sufficiently protects privacy. Through contextual integrity analysis, we identify group vulnerabilities created by inferential model development, illustrated in Table S2. Despite being promoted as effective protection against privacy harms [25, 75], anonymity fails against AI’s inferential capabilities. Our penetration test shows facial blurring provides no protection for food delivery workers against adversaries analyzing group membership. Authorities accessing DSI could organize targeted enforcement leading to loss of income, imprisonment, and other negative consequences.

The increasing density of street imagery, combined with advances in image retrieval techniques, makes it possible to circumvent anonymity in practice. As Barocas and Nissenbaum [12] explains: “Even where strong guarantees of anonymity can be achieved, common applications of big data undermine the values that anonymity traditionally protected. Even when individuals are not ‘identifiable,’ they may still be ‘reachable’ and subject to consequential inferences and predictions made on that basis.” Anonymity alone is insufficient to safeguard against the broader harms enabled by DSI and inferential AI models. CI’s theory of privacy does not deem all information flows involving ‘sensitive’ data as inappropriate. Instead, it emphasizes the need to evaluate fully defined information flows, which include the subject, sender, receiver, information type, and transmission principle [96]. This nuanced approach ensures that privacy judgments are context-specific and grounded in the norms governing the particular scenario. For example, a research or citizen advocacy group with access to the same New York City food delivery worker distribution, specifically tied to the use of that data for a legitimate purpose—such as surveying the state of food delivery in the city—would likely not be considered an inappropriate information flow. The primary use of DSI in this example is positioned as delivering social benefits, such as enhancing city planning and improving citizen safety through the understanding of group behaviors. Given these intended benefits, CI is the chosen privacy framework to evaluate and guide the responsible use of such DSI datasets. Full removal of all traces of human activity in DSI, such as by blocking all humans, their attachments, and their vehicles, would undermine the utility of these tools.

### 5.2 DSI’s Threat to Public Space Itself

Beyond harms to specific groups identified in our penetration test, contextual integrity theory highlights a fundamental concern: DSI threatens the nature and value of public space itself as a social resource [72, 107]. Public spaces have historically functioned as domains where contextual norms allow for spontaneity, economic opportunity, and democratic participation—values now threatened by surveillance infrastructures.

For food vendors and delivery workers in our study, this transformation is particularly consequential. What historically served as accessible venues for entrepreneurship—street corners and public thoroughfares—become sites of enforcement vulnerability when continuously documented through DSI. As our penetration test demonstrates, the ability to track food vendors’ locations and movement patterns fundamentally alters the social contract that has governed public space use. When vendors’ presence in a particular

location becomes algorithmically flagged (as shown in Figure 2), "perfect enforcement" becomes a possibility. Unlike before, where food vendors in New York City were able to establish themselves gradually, eventually gaining sufficient community support for advocacy groups to campaign effectively for policy reforms [94].

Similarly, for protesters, religious communities, and other groups identified in our typology (Table 2), DSI fundamentally reconstructs public space from a domain of relative freedom to one of persistent visibility. As Ben Green asserts in *The Smart Enough City*, the unfettered use of technologies like DSI are moving society towards a state where to avoid being tracked, you must take the quixotic step of opting out of public space [53]. The chilling effect on public assembly, worship, or everyday activities represents not merely a privacy harm to these groups, but a diminishment of public space's societal value. This perspective suggests that CI can help us understand how DSI may alter the implicit social agreements governing public spaces. Our findings indicate that the increasing presence of DSI, combined with AI analysis capabilities, could shift what Lane et al. [68] describe as "reasonable expectations" of contextual privacy in public spaces. As our food vendor case study suggests, these changes may disproportionately affect vulnerable groups who rely on public spaces for essential activities.

### 5.3 Recommendations

Motivated by the penetration test and the theoretical group privacy framework under contextual integrity, we suggest establishing responsible data practices particularly directed at research institutions and DSI providers. Additionally, we advocate for more nuance in technical approaches to privacy protection in DSI.

*Require DSI dataset usage approvals.* We call academic institutions to establish oversight bodies, similar to the Research Ethics Board proposed in the Menlo Report [11], to evaluate the ethical implications of research involving DSI datasets, particularly when such work falls outside the traditional scope of Institutional Review Boards (IRBs). Twelve years after the Menlo Report, we notice that our own institution lacks provisions for ethical review outside of research that deals directly with *human subjects*. In fact, prior work that utilizes DSI has been deemed IRB exempt, even when studying an innately societal phenomenon like police deployments [45]. As DSI and other sensitive datasets that *depict* individuals without their notice or consent are increasingly shared with researchers, we recommend that universities apply greater scrutiny to the projects that use them.

*Establishing DSI data usage norms and promises.* As DSI imagery becomes more widely available and the cost of associated analytics continues to decline (see our list of potential algorithmic group identification methods in Table 1), companies that provide access to these datasets should take responsibility for ensuring their ethical use by researchers, governments, and corporations. Unconditional sharing, without legal repercussions, will inevitably cause privacy harm to groups. We propose that DSI providers adopt a more practical approach to ensure that data sharing is restricted solely to legitimate purposes. Sharing and reuse should only occur under a new *transmission principle*: purpose-limited, privacy risk-assessed,

and with usage and reuse documented. To help achieve this, *researchers* working with DSI need to develop frameworks, databases, and a centralized system to track use agreements and ensure accountability in the sharing process. This system could be established as a "Data Use Agreement Database" where each use case is logged with respect to its purpose, risk assessment, and compliance with privacy protections.

*Study Societal Norms.* A crucial aspect of understanding proper information flows in DSI involves examining societal expectations regarding the depiction and inferrability of groups in such datasets, in line with work studying the public perception of DSI-producing technologies like CCTV [76], dashcams [8, 54], and smart glasses [64]. To acquire this knowledge, researchers should conduct factorial vignette studies [4] to build evidence that helps demarcate appropriate and inappropriate flows [19] in this novel technology. We propose research that puts a strong focus on potential privacy harms to groups.

*Develop contextual obfuscation tools for DSI.* There is a pressing need for more nuanced privacy protection tools tailored to DSI. As our pentest demonstrates, good faith privacy protections like facial de-identification are insufficient in protecting group privacy and, in certain instances, individual privacy too. A more robust approach such as blurring entire bodies and vehicles within a DSI frame may offer a short-term solution, but it falls short in two important respects. First, they substantially reduce the image's visual quality and utility. Second, they do not go far enough, as other surrounding factors, like body attachments and color information from blurred rectangles, allow for contextual identification (see Table S1). In Figure S1, we show how full-body blurring reduces DSI frames' visual quality and utility, while still leaking indicators permitting the inference of a farmer's market event happening. Further, other objects not attached to people can pose privacy threats to groups. This is shown in the supplement in Figure 2, where we document how inferences of food trucks, despite the blurring of license plates, pose similar privacy risks to food vendors.

In summary, we argue that image blurring techniques must move beyond generic object suppression and begin to treat privacy as a matter of contextually appropriate information flow. This requires the development of contextually aware obfuscation tools [20]. Rather than indiscriminately blurring objects, such tools should assess what information about individuals, or groups, might be inferred from contextual cues, and adapt the level of de-identification accordingly.

**5.3.1 Future Work.** In our work, we assume that an adversary is human. However, in contexts where fully automated enforcement targeting a group is implemented, it is imperative to evaluate the performance of leading Vision-Language Models (VLMs) in tasks such as de-identifying and clustering groups. This profiling could yield critical insights into the capabilities and limitations of automated systems in enforcing group-based surveillance or interventions. Further, there is a need for downstream downstream efforts in academia that promote proactive ethical decision-making in research, especially when working with DSI and other data streams that enable group-level measurement and identification, such as cell phone mobility data.



## ENDMATTER STATEMENTS

### Author Contributions

All authors discussed the results and implications and commented on the manuscript, but contributed at different stages.

**Matt Franchi:** Ideation, Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing - Original draft, Review & Editing, Visualization, Discussion.

**Hauke Sandhaus:** Conceptualization, Methodology, Theory, Investigation, Writing - Original draft, Review & Editing, Data curation, Discussion.

**Madiha Zahrah Choksi:** Theory, Writing - Original draft, Group typology.

**Severin Engelmann:** Writing - Original Draft, Inferences, Review & Editing.

**Wendy Ju:** Supervision, Resources.

**Helen Nissenbaum:** Supervision, Theoretical oversight.

### Acknowledgments

We thank Hal Triedman, Ricky Takkar, and Tom Ristenpart for formative feedback on a research proposal for this project. We are grateful for generative discussions with our colleagues Ilan Mandel, Gabriel Agostini, Sidhika Balachandar, and Emma Pierson at various stages of our research. We also thank the Digital Life Initiative Reading Group at Cornell Tech, and the Urban Tech Hub at Cornell Tech for longstanding research support. Finally, we thank Nexar, Inc., for collaboration, support, and data access.

### Positionality

Our interdisciplinary research team brings together diverse perspectives and expertise that shaped this work. The author team includes computer scientists with backgrounds in computer vision, machine learning, design, and human-computer interaction; interdisciplinary scholars with expertise in privacy law and data ethics; and philosophers focused on privacy theory and normative technology ethics.

All of our authors are based at a technology research institution in New York City, giving us firsthand experience with the urban environment we study. This positioning has informed our understanding of how dense street imagery technologies operate in practice and their implications for urban residents. We acknowledge that our institutional affiliations may influence access to resources, datasets, and industry partnerships that facilitated this research.

### Ethical Considerations

In our penetration testing experiment, where we assume the role of adversaries, we are committed to ensuring that no harm arises to any individuals depicted or inferable in our data. As a preliminary measure, all data used in this study was collected during 2023 and is over a year old at the time of publication. Our study exclusively focuses on groups that could face civil repercussions as a result of adversarial detection, deliberately avoiding cases where detection might lead to criminal consequences. To prevent any potential misuse of the models developed during this research, we delete all model weights and checkpoints prior to the project's conclusion.

We acknowledge our positionality as researchers and the interdisciplinary nature of our team. Our collaboration is emblematic of a balanced approach to both technical rigor and ethical considerations. Our methodology reflects a commitment to responsible data handling, focusing on the broader implications of DSI without exposing individuals or their specific contexts.

### Adverse Impact

We acknowledge and reflect on the fact that this research may cause readers to experience heightened anxiety over the state of increasing surveillance in our society, especially as we shed light on DSI, a relatively unknown technological product.

## REFERENCES

- [1] [n. d.]. Blur or remove 360 imagery & Photo Paths - Computer - Google Maps Help. <https://support.google.com/maps/answer/7011973?hl=en&co=GENIE.Platform%3DDesktop>
- [2] 2024. New Report Shows Mayor Adams, Commissioner Mayuga Deliver for Delivery Workers by Significantly Boos. <http://www.nyc.gov/office-of-the-mayor/news/539-24/new-report-shows-mayor-adams-commissioner-mayuga-deliver-delivery-workers-significantly>
- [3] 2024. Vendors rally against NYPD crackdown, call for more licensing in New York City. <https://abc7ny.com/street-vendors-rally-against-nypd-crackdown-set-to-hold-march-call-for-more-licensing-in-new-york-city/14687737/> Section: crime-safety.
- [4] Herman Aguinis and Kyle J Bradley. 2014. Best practice recommendations for designing and implementing experimental vignette methodology studies. *Organizational research methods* 17, 4 (25 oct 2014), 351–371. <https://doi.org/10.1177/1094428114547952>
- [5] Tim Alpherts, Sennay Ghebream, Yen-Chia Hsu, and Nanne Van Noord. 2024. Perceptive Visual Urban Analytics is Not (Yet) Suitable for Municipalities. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Rio de Janeiro Brazil, 1341–1354. <https://doi.org/10.1145/3630106.3658976>
- [6] Dragomir Anguelov, Carole Dulong, Daniel Filip, Christian Frueh, Stéphane Lafon, Richard Lyon, Abhijit Ogale, Luc Vincent, and Josh Weaver. 2010. Google Street View: Capturing the World at Street Level. *Computer* 43, 6 (June 2010), 32–38. <https://doi.org/10.1109/MC.2010.170> Conference Name: Computer.
- [7] Haleh Asgarinia. 2024. Limiting access to certain anonymous information: From the group right to privacy to the principle of protecting the vulnerable. *The Journal of value inquiry* (23 apr 2024), 1–27. <https://doi.org/10.1007/s10790-024-09980-x>
- [8] Dewan Mehrab Ashrafi. 2024. Technology Takes the Wheel: Unveiling the Drivers of Car Dashcam Adoption. *International Journal of Innovation and Technology Management* 21, 03 (May 2024), 2450024. <https://doi.org/10.1142/S021987702450024X> Publisher: World Scientific Publishing Co..
- [9] Cedar Attanasio. 2024. Getting Paid Now More Complex for NYC Food Delivery Workers. <https://www.tnnews.com/articles/new-york-food-delivery-workers>
- [10] Michael D. M. Bader, Stephen J. Mooney, Blake Bennett, and Andrew G. Rundle. 2017. The Promise, Practicalities, and Perils of Virtually Auditing Neighborhoods Using Google Street View. *The ANNALS of the American Academy of Political and Social Science* 669, 1 (Jan. 2017), 18–40. <https://doi.org/10.1177/0002716216681488> Publisher: SAGE Publications Inc.
- [11] Michael Bailey, David Dittrich, Erin Kennaally, and Doug Maughan. 2012. The Menlo Report. *IEEE Security & Privacy* 10, 2 (March 2012), 71–75. <https://doi.org/10.1109/MSP.2012.52> Conference Name: IEEE Security & Privacy.
- [12] Solon Barocas and Helen Nissenbaum. 2014. Big Data's End Run around Anonymity and Consent. In *Privacy, Big Data, and the Public Good: Frameworks for Engagement*, Helen Nissenbaum, Julia Lane, Stefan Bender, and Victoria Stodden (Eds.). Cambridge University Press, Cambridge, 44–75. <https://doi.org/10.1017/CBO9781107590205.004>
- [13] Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. 2006. Privacy and contextual integrity: Framework and applications. In *2006 IEEE symposium on security and privacy (S&P'06)*. IEEE, 15–pp. <https://ieeexplore.ieee.org/abstract/document/1624011>
- [14] Lyria Bennett Moses and Janet Chan. 2018. Algorithmic prediction in policing: assumptions, evaluation, and accountability. *Policing and Society* 28, 7 (Sept. 2018), 806–822. <https://doi.org/10.1080/10439463.2016.1253695>
- [15] Dulari Bhatt, Chirag Patel, Hardik Talsania, Jigar Patel, Rasmika Vaghela, Sharnil Pandya, Kirit Modi, and Hemant Ghayvat. 2021. CNN variants for computer vision: History, architecture, application, challenges and future scope. *Electronics* 10, 20 (2021), 2470.
- [16] Matt Bishop. 2007. About Penetration Testing. *IEEE Security & Privacy* 5, 6 (Nov. 2007), 84–87. <https://doi.org/10.1109/MSP.2007.159> Conference Name: IEEE Security & Privacy.
- [17] Edward J. Bloustein and Nathaniel J. Pallone. 2017. *Individual and Group Privacy*. Routledge, New York. <https://doi.org/10.4324/9781351319966>
- [18] Lukas Bossard, Matthias Dantone, Christian Leistner, Christian Wengert, Till Quack, and Luc Van Gool. 2013. Apparel classification with style. In *Computer Vision—ACCV 2012: 11th Asian Conference on Computer Vision, Daejeon, Korea, November 5–9, 2012, Revised Selected Papers, Part IV* 11. Springer, 321–335.
- [19] August Bourgeois, Laurens Vandercruysse, and Nanouk Verhulst. 2024. Understanding contextual expectations for sharing wearables' data: Insights from a vignette study. *Computers in human behavior reports* 15, 100443 (1 aug 2024), 100443. <https://doi.org/10.1016/j.chbr.2024.100443>
- [20] Mark Burdon, Tegan Cohen, Josh Buckley, and Michael Milford. 2024. From object obfuscation to contextually-dependent identification: enhancing automated privacy protection in street-level image platforms (SLIPs). *Information & Communications Technology Law* 33, 2 (May 2024), 198–221. <https://doi.org/10.1080/13600834.2024.2321052> Publisher: Routledge \_eprint: <https://doi.org/10.1080/13600834.2024.2321052>
- [21] Mark Burdon and Alissa McKillop. 2014. The Google street view Wi-Fi scandal and its repercussions for privacy regulation. *Monash University Law Review* 39, 3 (Jan. 2014), 702–738. <https://doi.org/10.3316/informit.376209506308923> Publisher: Monash University - Faculty of Business and Economics.
- [22] Richard Campanella. 2017. People-Mapping Through Google Street View. *Places Journal* (Nov. 2017). <https://doi.org/10.22269/171114>
- [23] Filippo Cavallo, Francesco Semeraro, Laura Fiorini, Gergely Magyar, Peter Sinčák, and Paolo Dario. 2018. Emotion modelling for social robotics applications: a review. *Journal of Bionic Engineering* 15 (2018), 185–203.
- [24] Urban Justice Center. 2023. Street Vendor Project. <https://www.streetvendor.org/what-is-svp>
- [25] David L. Chaum. 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* 24, 2 (Feb. 1981), 84–90. <https://doi.org/10.1145/358549.358563>
- [26] Wen-Huang Cheng, Sijie Song, Chieh-Yun Chen, Shintami Chusnul Hidayati, and Jiaying Liu. 2021. Fashion meets computer vision: A survey. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–41.
- [27] Madiha Zahrah Choksi, Ero Balso, Frauke Kreuter, and Helen Nissenbaum. 2024. Privacy for Groups Online: Context Matters. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (2024), 1–23.
- [28] Haidee Chu. 2024. NYPD Dragging Many More Vendors to Criminal Court, Data Shows. <http://www.thecity.nyc/2024/02/05/nypd-vendors-criminal-summonses-court-spike/>
- [29] Chris Clews, Roza Brajkovich-Payne, Emily Dwight, Ayob Ahmad Fauzul, Madeleine Burton, Olivia Carleton, Julie Cook, Charlotte Derolles, Ruby Faulkner, Mary Furniss, Anahera Herewini, Daymen Huband, Nerissa Jones, Cho Wool Kim, Alice Li, Jacky Lu, James Stanley, Nick Wilson, and George Thomson. 2016. Alcohol in urban streetscapes: a comparison of the use of Google Street View and on-street observation. *BMC Public Health* 16, 1 (May 2016), 442. <https://doi.org/10.1186/s12889-016-3115-9>
- [30] Etienne Corvee, Slawomir Bak, and François Bremond. 2012. People detection and re-identification for multi surveillance cameras. <https://inria.hal.science/hal-00656108>
- [31] Bahar Dadashova, Chiara Silvestri Dobrovolny, Mahmood Tabesh, Safety through Disruption (Safe-D) University Transportation Center (UTC), and Texas A&M Transportation Institute. 2021. *Detecting Pavement Distresses Using Crowd-sourced Dashcam Camera Images*. Technical Report TTTITI- Student – 07. <https://rosap.nhtl.bts.gov/view/dot/60311>
- [32] Yves-Alexandre De Montjoye, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports* 3, 1 (2013), 1–5. <https://doi.org/10.1038/srep01376> Publisher: Nature Publishing Group.
- [33] Yves-Alexandre De Montjoye, Laura Radaelli, Vivek Kumar Singh, and Alex “Sandy” Pentland. 2015. Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science* 347, 6221 (Jan. 2015), 536–539. <https://doi.org/10.1126/science.1256297>
- [34] Bryce J. Dietrich and Melissa L. Sands. 2023. Seeing racial avoidance on New York City streets. *Nature Human Behaviour* 7, 8 (Aug. 2023), 1275–1281. <https://doi.org/10.1038/s41562-023-01589-7> Publisher: Nature Publishing Group.
- [35] Sarah Elwood and Agnieszka Leszczynski. 2011. Privacy, reconsidered: New representations, data practices, and the geoweb. *Geoforum* 42, 1 (Jan. 2011), 6–15. <https://doi.org/10.1016/j.geoforum.2010.08.003>
- [36] Severin Engelmänn, Madiha Zahrah Choksi, Angelina Wang, and Casey Fiesler. 2024. Visions of a discipline: Analyzing introductory AI courses on YouTube. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 2400–2420.
- [37] Severin Engelmänn and Helen Nissenbaum. 2025. Countering Privacy Nihilism. In *Conceptions of Data Protection and Privacy: Legal and Philosophical Perspectives*, Elisa Orrù and Ralf Poscher (Eds.). Hart Publishing.
- [38] Severin Engelmänn, Chiara Ullstein, Orestis Papakyriakopoulos, and Jens Grossklags. 2022. What people think AI should infer from faces. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 128–141.
- [39] Arturo Flores and Serge Belongie. 2010. Removing pedestrians from Google street view images. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*. 53–58. <https://doi.org/10.1109/CVPRW.2010.5543255> ISSN: 2160-7516.
- [40] Luciano Floridi. 2017. Group Privacy - A Defense and an Interpretation. <https://doi.org/10.2139/ssrn.3854483>
- [41] Jessica Formoso. 2024. Bronx bodega delivery workers targeted in string of robberies. <https://www.fox5ny.com/news/bronx-bodega-delivery-workers-targeted-string-robberies> Publisher: FOX 5 New York.
- [42] Matt Franchi, Debargha Dey, and Wendy Ju. 2024. Towards Instrumented Fingerprinting of Urban Traffic: A Novel Methodology using Distributed Mobile Point-of-View Cameras. In *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '24)*. Association for Computing Machinery, New York, NY, USA, 53–62. <https://doi.org/10.1145/3640792.3675740>

- [43] Matt Franchi, Nikhil Garg, Wendy Ju, and Emma Pierson. 2025. Bayesian Modeling of Zero-Shot Classifications for Urban Flood Detection. <https://doi.org/10.48550/arXiv.2503.14754> arXiv:2503.14754 [cs].
- [44] Matthew Franchi, Maria Teresa Parreira, Fanjun Bu, and Wendy Ju. 2025. The Robotability Score: Enabling Harmonious Robot Navigation on Urban Streets. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–17. <https://doi.org/10.1145/3706598.3714009>
- [45] Matt Franchi, J.D. Zamfirescu-Pereira, Wendy Ju, and Emma Pierson. 2023. Detecting disparities in police deployments using dashcam data. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAcCT '23)*. Association for Computing Machinery, New York, NY, USA, 534–544. <https://doi.org/10.1145/3593013.3594020>
- [46] Andrea Frome, German Cheung, Ahmad Abdulkader, Marco Zennaro, Bo Wu, Alessandro Bissacco, Hartwig Adam, Hartmut Neven, and Luc Vincent. 2009. Large-scale privacy protection in Google Street View. In *2009 IEEE 12th International Conference on Computer Vision*. 2373–2380. <https://doi.org/10.1109/ICCV.2009.5459413> ISSN: 2380-7504.
- [47] Patrick Gallo and Houssain Kettani. 2020. On Privacy Issues with Google Street View CLEAR Conference Computer Science Academic Papers. *South Dakota Law Review* 65, 3 (2020), 608–622. <https://heinonline.org/HOL/P?h=hein.journals/sdlr65&i=666>
- [48] Timnit Gebru, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. 2017. Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States. *Proceedings of the National Academy of Sciences* 114, 50 (Dec. 2017), 13108–13113. <https://doi.org/10.1073/pnas.1700035114>
- [49] R Stuart Geiger, Kevin Yu, Yanlai Yang, Mindy Dai, Jie Qiu, Rebekah Tang, and Jenny Huang. 2020. Garbage in, garbage out? Do machine learning application papers in social computing report where human-labeled training data comes from?. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 325–336.
- [50] Roger C Geissler. [n. d.]. Private Eyes Watching You: Google Street View and the Right to an Inviolable Personality. *HASTINGS LAW JOURNAL* 63 ([n. d.]).
- [51] Jake Goldenfein. 2019. The profiling potential of computer vision and the challenge of computational empiricism. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 110–119.
- [52] Google. 2025. Google-Contributed Street View Imagery Policy. <https://www.google.com/streetview/policy/>
- [53] Ben Green. 2019. *The Smart Enough City: Putting Technology in Its Place to Reclaim Our Urban Future*. MIT Press. Google-Books-ID: avGRdWAAQBAJ.
- [54] Tim Gruchmann and Amer Jazaïry. 2025. Big brother is watching you: Examining truck drivers' acceptance of road-facing dashcams. *Transportation Research Part F: Traffic Psychology and Behaviour* 111 (May 2025), 316–330. <https://doi.org/10.1016/j.trf.2025.03.015>
- [55] Aya Hassouneh, AM Mutawa, and M Murugappan. 2020. Development of a real-time emotion recognition system using facial expressions and EEG based on machine learning and deep neural network methods. *Informatics in Medicine Unlocked* 20 (2020), 100372.
- [56] Andrew J. Hawkins. 2024. Lyft is also partnering with robotaxi companies. *The Verge* (Nov. 2024). <https://www.theverge.com/2024/11/6/24289475/lyft-may-mobility-mobility-nexar-autonomous-robotaxi>
- [57] Mark Healy. 2024. How the NYPD's Scooter Crackdown Beat Down Food Delivery Workers. <https://www.curbed.com/article/moped-crackdown-nypd-seizure-delivery-workers-erie-basin.html>
- [58] Marco Helbich, Matthew Danish, S. M. Labib, and Britta Ricker. 2024. To use or not to use proprietary street view images in (health and place) research? That is the question. *Health & Place* 87 (May 2024), 103244. <https://doi.org/10.1016/j.healthplace.2024.103244>
- [59] Sam Hind and Alex Gekker. 2024. Automotive parasitism: Examining Mobileye's 'car-agnostic' platformisation. *New Media & Society* 26, 7 (July 2024), 3707–3727. <https://doi.org/10.1177/14614448221104209> Publisher: SAGE Publications.
- [60] Lazar Ilic, Michael Sawada, and Amaury Zarzelli. 2019. Deep mapping gentrification in a large Canadian city using deep learning and Google Street View. *PLoS one* 14, 3 (2019), e0212814. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0212814> Publisher: Public Library of Science San Francisco, CA USA.
- [61] Joel Janai, Fatma Güney, Aseem Behl, Andreas Geiger, et al. 2020. Computer vision for autonomous vehicles: Problems, datasets and state of the art. *Foundations and Trends® in Computer Graphics and Vision* 12, 1–3 (2020), 1–308.
- [62] Glenn Jocher, Jing Qiu, and Ayush Chaurasia. 2023. Ultralytics YOLO. <https://github.com/ultralytics/ultralytics> original-date: 2022-09-11T16:39:45Z.
- [63] Dimitrios Kastaniotis, Ilias Theodorakopoulos, George Economou, and Spiros Fotopoulos. 2013. Gait-based gender recognition using pose information for real time applications. In *2013 18th international conference on digital signal processing (DSP)*. IEEE, 1–6.
- [64] Fareed Kaviani, Ben Lyall, and Sjaan Koppel. 2024. Exploring social perceptions of everyday smartglass use in Australia. *PLOS ONE* 19, 11 (Nov. 2024), e0313001. <https://doi.org/10.1371/journal.pone.0313001> Publisher: Public Library of Science.
- [65] Kanako Kawaguchi and Yukiko Kawaguchi. 2012. What Does Google Street View Bring about? -Privacy, Discomfort and The Problem of Paradoxical Others-. *Contemporary and Applied Philosophy* 4 (Aug. 2012), 19–34. <https://doi.org/10.14989/180276> Accepted: 2014-01-15T04:33:11Z.
- [66] Junghwan Kim and Kee Moon Jang. 2023. An examination of the spatial coverage and temporal variability of Google Street View (GSV) images in small- and medium-sized cities: A people-based approach. *Computers, Environment and Urban Systems* 102 (June 2023), 101956. <https://doi.org/10.1016/j.compenurbysys.2023.101956>
- [67] Anil Kunchala, Melanie Bourroche, and Bianca Schoen-Phelan. 2023. Towards A Framework for Privacy-Preserving Pedestrian Analysis. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Waikoloa, HI, USA, 4359–4369. <https://doi.org/10.1109/WACV56688.2023.00435>
- [68] Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum. 2014. *Privacy, big data, and the public good: Frameworks for engagement*. Cambridge University Press, Cambridge, England.
- [69] Michael J. V. Leach, Rolf Baxter, Neil M. Robertson, and Ed P. Sparks. 2014. Detecting Social Groups in Crowded Surveillance Videos Using Visual Attention. 461–467. [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_workshops\\_2014/W14/html/Leach\\_Detecting\\_Social\\_Groups\\_2014\\_CVPR\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_workshops_2014/W14/html/Leach_Detecting_Social_Groups_2014_CVPR_paper.html)
- [70] Michele Loi and Markus Christen. 2020. Two concepts of group privacy. *Philosophy & technology* 33, 2 (June 2020), 207–224. <https://doi.org/10.1007/s13347-019-00351-0>
- [71] Kristian Lum and William Isaac. 2016. To predict and serve? *Significance* 13, 5 (2016), 14–19. <https://doi.org/10.1111/j.1740-9713.2016.00960.x> Publisher: Oxford University Press.
- [72] Kevin Macnish. 2012. Unblinking eyes: the ethics of automating surveillance. *Ethics and information technology* 14, 2 (7 jun 2012), 151–167. <https://doi.org/10.1007/s10676-012-9291-0>
- [73] Alessandro Mantelero. 2017. From group privacy to collective privacy: Towards a new dimension of privacy and data protection in the big data era. In *Group Privacy*. Springer International Publishing, Cham, 139–158. [https://doi.org/10.1007/978-3-319-46608-8\\_8](https://doi.org/10.1007/978-3-319-46608-8_8)
- [74] Coral Murphy Marcos. 2021. As Bike Thefts Jump, Delivery Workers Band Together for Safety. *The New York Times* (Oct. 2021). <https://www.nytimes.com/2021/10/12/business/delivery-workers-thefts-neighborhood-watch.html>
- [75] Steve Matthews. 2010. Anonymity and the Social Self. *American Philosophical Quarterly* 47, 4 (2010), 351–363. <https://www.jstor.org/stable/25734161> Publisher: [North American Philosophical Publications, University of Illinois Press].
- [76] Lorraine Mazerolle, David Hurley, and Mitchell Chamlin. 2002. Social Behavior in Public Space: An Analysis of Behavioral Adaptations to CCTV. *Security Journal* 15, 3 (July 2002), 59–75. <https://doi.org/10.1057/palgrave.sj.8340118>
- [77] Mobileye. 2022. Mobileye's Self-Driving Secret? 200PB of Data | Mobileye Blog. <https://www.mobileye.com/blog/mobileye-cs-2022-self-driving-secret-data/>
- [78] Nikhil Naik, Jade Philipoom, Ramesh Raskar, and César Hidalgo. 2014. Streetscore-predicting the perceived safety of one million streetscapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 779–785. [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_workshops\\_2014/W20/html/Naik\\_Streetscore\\_-\\_Predicting\\_2014\\_CVPR\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_workshops_2014/W20/html/Naik_Streetscore_-_Predicting_2014_CVPR_paper.html)
- [79] Quynh C. Nguyen, Sahil Khanna, Pallavi Dwivedi, Dina Huang, Yuru Huang, Tolga Tasdizen, Kimberly D. Brunisholz, Feifei Li, Wyatt Gorman, and Thu T. Nguyen. 2019. Using Google Street View to examine associations between built environment characteristics and US health outcomes. *Preventive medicine reports* 14 (2019), 100859. <https://www.sciencedirect.com/science/article/pii/S2211335519300440> Publisher: Elsevier.
- [80] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119. [https://heinonline.org/hol/cgi-bin/get\\_pdf.cgi?handle=hein.journals/washlr79&section=16](https://heinonline.org/hol/cgi-bin/get_pdf.cgi?handle=hein.journals/washlr79&section=16) Publisher: HeinOnline.
- [81] Helen Nissenbaum. 2011. A Contextual Approach to Privacy Online. *Daedalus* 140, 4 (Oct. 2011), 32–48. [https://doi.org/10.1162/DAED\\_a\\_00113](https://doi.org/10.1162/DAED_a_00113)
- [82] Helen Nissenbaum. 2019. Contextual Integrity Up and Down the Data Food Chain. *Theoretical Inquiries in Law* 20, 1 (Jan. 2019), 221–256. <https://doi.org/10.1515/til-2019-0008> Publisher: De Gruyter.
- [83] Angelo Nodari, Marco Vanetti, and Ignazio Gallo. 2012. Digital privacy: Replacing pedestrians from Google Street View images. In *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*. 2889–2893. [https://ieeexplore.ieee.org/abstract/document/6460769?casa\\_token=21svJksOsrUAAAAA:k1y3DyIhtAKZyWjhl58\\_4UbM6LICrAc270qW\\_BISSIn1r1evSrnvCSwneVvEqLhcdV\\_13Wq](https://ieeexplore.ieee.org/abstract/document/6460769?casa_token=21svJksOsrUAAAAA:k1y3DyIhtAKZyWjhl58_4UbM6LICrAc270qW_BISSIn1r1evSrnvCSwneVvEqLhcdV_13Wq) ISSN: 1051-4651.
- [84] nycadmin. 2014. New York City Food by the Numbers: Food Trucks. <https://www.nycfoodpolicy.org/new-york-city-food-numbers-food-trucks/>
- [85] NYCOpenData. 2025. OATH Hearings Division Case Status. <https://data.cityofnewyork.us/City-Government/OATH-Hearings-Division-Case-Status/jz4z-kudi>




- [86] Rommel H. Ojeda, Ambar Reyes, April Xu, CEDAR ATTANASIO Documented, and Associated Press • • 2024. Newly arrived migrants encounter hazards of food delivery on NYC streets: robbers. <https://www.nbcnewyork.com/news/local/safety-migrants-encounter-hazards-food-delivery-nyc-robbers/5609522/>
- [87] Fei Pan, Sangryul Jeon, Brian Wang, Frank Mckenna, and Stella X. Yu. 2024. Zero-Shot Building Attribute Extraction From Large-Scale Vision and Language Models. 8647–8656. [https://openaccess.thecvf.com/content/WACV2024/html/Pan\\_Zero-Shot\\_Building\\_Attribute\\_Extraction\\_From\\_Large-Scale\\_Vision\\_and\\_Language\\_Models\\_WACV\\_2024\\_paper.html](https://openaccess.thecvf.com/content/WACV2024/html/Pan_Zero-Shot_Building_Attribute_Extraction_From_Large-Scale_Vision_and_Language_Models_WACV_2024_paper.html)
- [88] Eugenia Politou, Efthimios Alepis, and Constantinos Patsakis. 2018. Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions. *Journal of Cybersecurity* 4, 1 (Jan. 2018). <https://doi.org/10.1093/cybersec/tyy001>
- [89] Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. 2017. Knock knock, who's there? Membership inference on aggregate location data. *arXiv preprint arXiv:1708.06145* (2017).
- [90] Keith A. Redmill, Ekim Yurtsever, Rabi G. Mishalani, Benjamin Coifman, and Mark R. McCord. 2023. Automated Traffic Surveillance Using Existing Cameras on Transit Buses. *Sensors* 23, 11 (Jan. 2023), 5086. <https://doi.org/10.3390/s23115086> Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.
- [91] Kate Robertson, Cynthia Khoo, and Yolanda Song. 2020. To surveil and predict: A human rights analysis of algorithmic policing in Canada. (2020). <https://citizenlab.ca/wp-content/uploads/2021/01/To-Surveil-and-Predict1.1.pdf>
- [92] Paul D Rosero-Montalvo, Diego Hernn Peluffo-Ordóñez, Vivian Felix Lopez Batista, Jorge Serrano, and Edwin A Rosero. 2018. Intelligent system for identification of wheelchair user's posture using machine learning techniques. *IEEE Sensors Journal* 19, 5 (2018), 1936–1942.
- [93] Hauke Sandhaus, Angel Wang, Qian Yang, and Wendy Ju. 2024. My Precious Crash Data: Barriers and Opportunities in Encouraging Autonomous Driving Companies to Share Safety-Critical Data. *arXiv:2504.17792 [cs.HC]* <https://arxiv.org/abs/2504.17792> To appear in Proc. ACM Hum.-Comput. Interact., Computer-Supported Cooperative Work & Social Computing (CSCW), 2025.
- [94] Shamier Settle and David Dyssegaard Kallick. 2024. *Street Vendors of New York*. Technical Report. Immigration Research Initiative. <https://immresearch.org/publications/street-vendors-of-new-york/>
- [95] Shapira, Dorin, Franchi, Matthew, and Ju, Wendy. 2024. Fingerprinting New York City's Scaffolding Problem with Longitudinal Dashcam Data.
- [96] Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2018. Analyzing Privacy Policies Using Contextual Integrity Annotations. <https://doi.org/10.48550/arXiv.1809.02236> *arXiv:1809.02236 [cs]*.
- [97] Corey Snyder and Minh Do. 2019. STREETS: A Novel Camera Network Dataset for Traffic Flow. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2019/hash/ee389847678a3a9d1ce9e4ca69200d06-Abstract.html>
- [98] Luke Stark and Jesse Hoey. 2021. The ethics of emotion in artificial intelligence systems. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 782–793.
- [99] Dodai Stewart and Juan Arredondo. 2024. Can New York City Street Vendors Survive a Police Crackdown? *The New York Times* (Sept. 2024). <https://www.nytimes.com/2024/09/16/nyregion/street-wars-vendors.html>
- [100] Linnet Taylor, Luciano Floridi, and Bart van der Sloot (Eds.). 2017. *Group privacy: New challenges of data technologies* (1 ed.). Springer International Publishing, Cham, Switzerland. <https://doi.org/10.1007/978-3-319-46608-8>
- [101] William Thackway, Matthew Ng, Chyi-Lin Lee, and Christopher Pettit. 2023. Implementing a deep-learning model using Google street view to combine social and physical indicators of gentrification. *Computers, Environment and Urban Systems* 102 (2023), 101970. <https://www.sciencedirect.com/science/article/pii/S0198971523000339> Publisher: Elsevier.
- [102] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. 2024. Cambrian-1: A Fully Open, Vision-Centric Exploration of Multimodal LLMs. <http://arxiv.org/abs/2406.16860> *arXiv:2406.16860*
- [103] Dai Quoc Tran, Minsoo Park, Yuntae Jeon, Jinyeong Bak, and Seunghee Park. 2022. Forest-Fire Response System Using Deep-Learning-Based Approaches With CCTV Images and Weather Data. *IEEE Access* 10 (2022), 66061–66071. <https://doi.org/10.1109/ACCESS.2022.3184707> Conference Name: IEEE Access.
- [104] Ries Uittenbogaard, Clint Sebastian, Julien Vijverberg, Bas Boom, Dariu M. Gavrilă, and Peter H. N. de With. 2019. Privacy Protection in Street-View Panoramas Using Depth and Multi-View Imagery. 10581–10590. [https://openaccess.thecvf.com/content\\_CVPR\\_2019/html/Uittenbogaard\\_Privacy\\_Protection\\_in\\_Street-View\\_Panoramas\\_Using\\_Depth\\_and\\_Multi-View\\_Imagery\\_CVPR\\_2019\\_paper.html](https://openaccess.thecvf.com/content_CVPR_2019/html/Uittenbogaard_Privacy_Protection_in_Street-View_Panoramas_Using_Depth_and_Multi-View_Imagery_CVPR_2019_paper.html)
- [105] Chiara Ullstein, Severin Engelmann, Orestis Papakyriakopoulos, Michel Ho-hendanner, and Jens Grossklags. 2022. AI-competent individuals and laypeople tend to oppose facial analysis AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–12.
- [106] Chiara Ullstein, Severin Engelmann, Orestis Papakyriakopoulos, Yuko Ikkatai, Naira Paola Arnez-Jordan, Rose Caleno, Brian Mboya, Shuichiro Higuma, Tilman Hartwig, Hiromi Yokoyama, et al. 2024. Attitudes Toward Facial Analysis AI: A Cross-National Study Comparing Argentina, Kenya, Japan, and the USA. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 2273–2301.
- [107] Bart van der Sloot. 2017. Do groups have a right to protect their group interest in privacy and should they? Peeling the onion of rights and interests protected under article 8 ECHR. In *Group Privacy*. Springer International Publishing, Cham, 197–224. [https://doi.org/10.1007/978-3-319-46608-8\\_11](https://doi.org/10.1007/978-3-319-46608-8_11)
- [108] Anton Vedder. 1999. KDD: The challenge to individualism. *Ethics and information technology* 1, 4 (Dec. 1999), 275–281. <https://doi.org/10.1023/a:1010016102284>
- [109] Jiahe Wang, Jiale Huang, Bingzhao Cai, Yifan Cao, Xin Yun, and Shangfei Wang. 2024. Zero-shot Compound Expression Recognition with Visual Language Model at the 6th ABAW Challenge. <https://doi.org/10.48550/arXiv.2403.11450> *arXiv:2403.11450*.
- [110] Li Yin, Qimin Cheng, Zhenxin Wang, and Zhenfeng Shao. 2015. 'Big data' for pedestrian volume: Exploring the use of Google Street View images for pedestrian counts. *Applied Geography* 63 (Sept. 2015), 337–345. <https://doi.org/10.1016/j.apgeog.2015.07.010>
- [111] J.D. Zamfirescu-Pereira, Jerry Chen, Emily Wen, Allison Koenecke, Nikhil Garg, and Emma Pierson. 2022. Trucks Don't Mean Trump: Diagnosing Human Error in Image Analysis. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 799–813. <https://doi.org/10.1145/3531146.3533145>

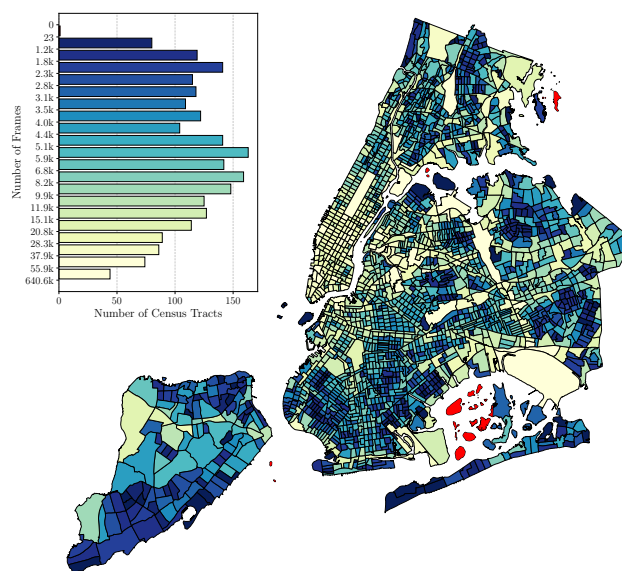
## A SUPPLEMENTAL FIGURES



**Figure S1:** An example of a DSI image processed by Nexar's current pedestrian obfuscation algorithm, showing that pedestrians' entire bodies and faces are obscured by a blurred rectangular box. While individual identities are rendered nearly impossible to infer, environmental cues still allow for the inference of the scene being a farmer's market.



**Figure S2:** An example of group membership inference, even under full-body pedestrian obfuscation. Due to the high-visibility vest worn by this NYPD traffic officer, a group membership inference can be made solely from neon-green color, black pants, and situation on the corner of a traffic intersection.



**Figure S3:** Chloropleth map of the census tracts of NYC, colored by number of dashcam images acquired in a tract. Counts can be referenced in the accompanying frequency histogram.

## B FURTHER DETAILS ON ATTACK TOOL IMPLEMENTATIONS

### B.1 Experimental Design

We design two experiments to demonstrate the privacy threats that commercially standard DSI imagery can pose to groups when combined with image retrieval and processing using machine learning (ML).

**B.1.1 Sourcing DSI.** We developed software to systematically extract dashcam images, organizing them by date and capture location. Between August 11, 2023, and January 10, 2024, we collect a total of 25,232,608 images, ensuring comprehensive geographic and temporal coverage.<sup>6</sup> Throughout the sampling period, the data provider maintained an assertion that sampled imagery was driven by (1) crafting a representative sample and (2) replacing stale imagery with fresh imagery. An important and key limitation of this dataset is that we can not independently verify that images were randomly sampled from the network of cameras, nor can we access statistics about the members of the camera network. We only have access to downstream imagery.

To illustrate the temporal density of the dashcam dataset, we group the images into 15-minute intervals that encompass all of the times that the scraper was fully operational. This produces 6,144 intervals. Out of these intervals, only 4, produce no new imagery; when investigating, we find that this corresponds to the clocks being set back an hour at the end of daylight savings time, on November 5 2023. The mean 15-minute interval produces 3,444 new images, across an area 350.3 square miles (for reference, New York City and its water areas encompass about 469 square miles).<sup>7</sup>

We offer an important caveat regarding the rigor of our data validation and training process. Our objective is investigatory rather than focused on developing a highly accurate model. Several reasons inform this approach: primarily, achieving high accuracy would likely necessitate the use of crowdsourced human annotators [45], whom we are unwilling to expose to sensitive imagery. Second, our core goal is to demonstrate that group privacy threats are both real and present in DSI data. For this purpose, somewhat imprecise distributions are sufficient to support our findings.

**B.1.2 Crafting Data.** We query approximately 500,000 randomly-sampled images with Cambrian-13B, a leading open-source vision language model [102], in 2.3 days, asking the model for each image, 'Does this image show a food truck?'. Of our subset, 2,903 are inferred positive, and 557,602 are inferred negative. We then randomly sample 2,000 images from the set of predicted positive images for human annotation. Two authors distributed the labeling process among themselves. We randomly sample 50,000 images from the set of predicted negative images to use as background images in our model. Then, from this set of 52,000 images, we craft a 60-20-20 training-validation-test split; we will use these splits in training a

more lightweight object detection model downstream in the penetration test.

**B.1.3 Data Validation.** Cambrian-13B is a time and compute-intensive model; we require a RTX A6000 GPU with 48GB of ram to load and infer images with the model, and inference takes around 2 seconds per image. That said, it has empirically useful zero-shot accuracy. We evaluate Cambrian's ability to classify images as containing a food truck. To evaluate the precision of Cambrian on this task, we randomly sample 2000 images from the set of images classified positive, and manually annotate them. Of these 2000 images, 1496 contain food trucks (true positives), and 645 do not (false positives), yielding a true positive rate (TPR) of 0.70. For false positives, we manually annotate 'decoys' in the image, or objects that we infer Cambrian mistook for a food truck<sup>8</sup>. To evaluate the recall of Cambrian on this task, we randomly sample 200 images from the set of images classified negative, and manually annotate them. Of these 200 images, 3 depict food trucks, yielding a false negative rate (FNR) of 0.015<sup>9</sup>. This seems small, but the FNR will magnify at the scale of our dataset; we estimate that Cambrian-13B missed 377,136 images with food trucks. Nonetheless, Cambrian is effective at demarcating positives and negatives; effectively, an image that Cambrian predicts as having a food truck will depict a food truck 70% of the time, and an image that Cambrian predicts as not having a food truck will have a food truck only 1.5% of the time.

**B.1.4 Model Validation.** The trained model fits well to the manually-labeled positive images from Cambrian. Figure S4 shows the Receiver Operating Characteristic (ROC) and Precision-Recall (PR) curves for the most performant decoy-enabled model and most performance decoy-disabled model, as evaluated on the test set. Both models achieve an average precision (AP) of 0.78. As the decoy-enabled model achieves higher AUC (0.96 vs. 0.94), we select it for inference on the entire dataset, or "deployment".

**B.1.5 Lightweight Object Detection Models.** We train You-Only-Look-Once (YOLO, specifically YOLOv11 [62]) object detection models on our crafted dataset to induce the capability to identify group members in the entire set of dashcam images. We use the this architecture as it is a standard and principled tool used for object detection tasks in urban scenes ([45], [42], [95]). We train one model per object of interest, and report standard performance metrics from the data splits described in the previous section.

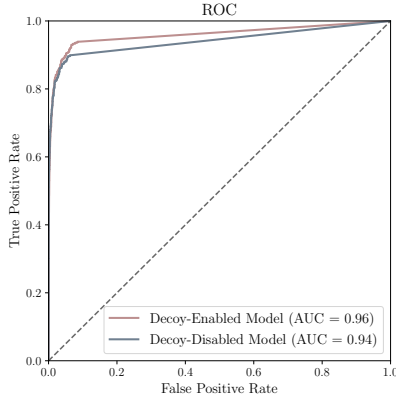
<sup>6</sup>We provide additional context regarding our sampling process, which was notably complex. On October 1, 2024, the data provider overhauled the API used for downloading frames and metadata, rendering our tool inoperable. We adapted to these changes and resumed data collection on October 20, 2024. Alongside the API overhaul, the data provider announced a reduction in the daily volume of uploaded imagery, attributing the change to operating cost constraints.

<sup>7</sup>Here, we calculate the area spanned by all images in a 15-minute interval by computing the convex hull of the image subset.

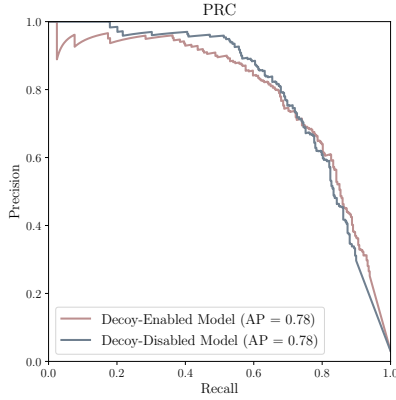
<sup>8</sup>False positives detected by Cambrian stem from two sources. Visual Confusions: Objects that look like food trucks, such as NYC dining sheds with LED signs. Language Confusion: Trucks featuring food images that lead to the incorrect assumption they are food trucks.

<sup>9</sup>It is worth noting that two of the three of these images are partially or almost fully blocked by an full-body obfuscated pedestrian





(a) Receiver operating characteristic (ROC) curve.



(b) Precision-recall (PR) curve.

**Figure S4: Standard performance curves for our food truck object detection model.**

## C TAXONOMY OF DE-IDENTIFICATION IN DSI

Burdon et al. [20] identify four key failures in the standard de-identification approach of blurring faces, license plates, and other identifiable objects: false negatives, false positives, the ‘Streisand effect,’ and contextual identification. False negatives and false positives are common concepts in the machine learning literature, referring to instances where de-identification fails to blur identifiable objects (false negatives) or unnecessarily obscures non-identifiable objects (false positives).

False positives do not necessarily raise privacy concerns, and more so impact the quality of the product’s imagery. We introduce a new failure mode that emerges from considering *group privacy*, instead of only the privacy of individuals. We call this failure mode ‘group membership inference.’ In this failure mode, a collection of de-identified objects of the same type can be clustered using computational methods, leaking the privacy guarantees of the de-identification. We offer a concrete example to illustrate this point: consider construction workers, who frequently wear high-visibility

green vests. Even with the most stringent de-identification method, blurring the entire pedestrian body, a clustering algorithm could detect and group these vests. This would effectively compromise the privacy guarantees for the group of construction workers, despite individual de-identification. We summarize DSI’s de-identification failure modes in Table S1.

Mode of Failure	Description of object de-identification issue
False Negatives	System fails to accurately detect and blur features of an object before images are made publicly available.
False Positives	System detects and blurs an object which is not required to be blurred.
The "Streisand" Effect	System blurs/blocks an object, and the act of blurring paradoxically draws attention to the object which is intended to be concealed.
Contextual Identification	System accurately detects and blurs/redacts features, but the object is identifiable from contextual indicators.
Membership inference	System blurs, blocks or de-identifies several objects of the same class in a visually-similar fashion, allowing for clusters to be generated that leak privacy and enable group membership inference.

Table S1: De-identification failure modes in dense street imagery. The top segment of the table is inherited from Burdon et al. [20].

D   EXAMPLES OF GROUP-BASED PRIVACY VIOLATIONS IN DSI UNDER CI

<b>Protesters</b> <b>Information:</b> Meeting locations and times (geo-tagged). <b>Recipients:</b> Political groups opposed to the protester’s cause. <b>Transmission Principle:</b> Information about meeting locations and times is shared outside the original context of trusted participants and disseminated to opposition political groups, violating expectations of confidentiality and purpose limitation. <b>Harms:</b> Retaliation, disruption of assemblies, suppression of free speech.
<b>Nurses</b> <b>Information:</b> Shift patterns derived from DSI near healthcare facilities. <b>Recipients:</b> Malicious employers or staffing agencies. <b>Transmission Principle:</b> Shift pattern data is aggregated and shared with malicious employers or staffing agencies without the consent of the individuals, violating principles of data minimization and appropriate recipient access. <b>Harms:</b> Exploitation of labor patterns, unsafe working conditions, reduced autonomy.
<b>Commuters</b> <b>Information:</b> Behavioral patterns at busy crosswalks during peak hours. <b>Recipients:</b> Predatory advertisers. <b>Transmission Principle:</b> Behavioral patterns shared with predatory advertisers violates expectations of anonymity and proportionality in the collection and use of public data. <b>Harms:</b> Exploitation through targeted marketing (e.g., payday loans, fast food).
<b>Religious groups</b> <b>Information:</b> Images and metadata in cultural or religious locations. <b>Recipients:</b> Hate groups. <b>Transmission Principle:</b> Sensitive imagery and metadata from religious locations shared with hate groups violate principles of contextual sensitivity, trust, and non-maleficence in handling sensitive personal data. <b>Harms:</b> Harassment, stigmatization, violence, violation of religious freedoms.

Table S2: Examples of Group-Based Privacy Violations in DSI under CI