Preventing Surveillance Cities: Developing a Set of Fundamental Privacy Provisions

Wajeeha Ahmad[1] and Elizabeth Dethy[1]

[1]Massachusetts Institute of Technology, Internet Policy Research Initiative, CSAIL

Corresponding author: wajeeha@mit.edu

Keywords: Privacy, Smart Sensors, Smart City, Surveillance, Public Policy

Executive Summary: We propose a set of five fundamental privacy provisions that are essential to mitigate the harms associated with exploiting smart city data. By doing so, we aim to address the existing gaps between the technical solutions proposed in previous works and the government regulations deemed necessary to protect privacy in prior literature but left unspecified in terms of their specific goals. Our policy proposals include differentiating personally identifiable smart city data from de-identified data, creating a warrant requirement for personally identifiable smart city data, limiting the sharing of personally identifiable information collected by smart city sensors, adopting data minimization requirements, and introducing private and public enforcement mechanisms. Taken together, these provisions can lay the foundation for creating a robust, privacy-protective response to the threats posed by unregulated access to smart city data. In order to prevent the emergence of surveillance cities, we encourage states and local governments to implement these fundamental privacy provisions in their specific jurisdictions.

I. Introduction

Smart cities are sprouting up all across the world. While many cities already have cameras monitoring traffic intersections and public areas, in new smart city projects, the camera feeds are augmented by data from weather reports, shipping movements, license plate readers, gunshot-detection microphones, and facial recognition software. In other cases, cities are experimenting with sensors under roads to track how people navigate and interact with certain landmarks, and deploying street lights with built-in motion-sensors and cameras that detect pedestrians and turn off when no one is around. With numerous data streams collected from new sensors, unregulated access to smart city data by both authorities and private actors can result in chilling effects on the freedoms of speech, movement and association. The availability of large volumes of such data could also be misused to restrict dissent and democratic rights.

However, in most places around the world, existing laws and regulations do not adequately address privacy concerns regarding smart city data containing personally identifiable information (Weber and Žarko 2019). To prevent the exploitation of smart city data for pernicious targeting, cities must limit lawful access to personal smart city data unless probable cause to obtain the data exists. Additionally, since individuals cannot opt out of this data collection, their personal data should not be shared with or sold to third parties. Restricting the number of actors with access to this data and placing limitations on data retention can also reduce the risk of exploitation by malicious actors. Bolstering local privacy laws by addressing these risks can ensure that smart cities do not inadvertently become surveillance cities.

II. Benefits of smart city technology: improved efficiency, cost savings and public safety

The proliferation of internet-connected sensors has allowed cities to develop real-time monitoring systems for energy distribution, traffic control, and crime prevention, among numerous other uses. Smart cities leverage such sensors to improve citizens' quality of life, conserve city resources, and spur economic development. Many cities are embarking on their own smart city initiatives to address the issues of mobility, safety, and development. For the purposes of this paper, we define smart city sensor data as any data collected, transmitted, or stored by a device installed in a public location or embedded in civic systems and infrastructure. This may include audio, photo-optical, thermal, aerosol, or electronic data.

Smart city technology offers incredible cost-saving potential. This is because data collected from smart sensors can be used to improve resource allocation for delivering essential public services such as water and lighting, resulting in millions of dollars of energy savings over several years. According to one estimate, five years from now, 75 smart cities worldwide with an estimated 10 million inhabitants each, can save over an aggregated $5 trillion per year based on savings from smart utilities, buildings, transportation, manufacturing and other sectors (Bonte 2017). It will become increasingly more important to take advantage of these efficiency gains as cities continue to grow at an exponential rate. According to the latest estimate from the United Nations Population Division, More than half of the world's population currently lives in urban areas, and approximately 68 percent of the world's population will live in an urban environment by 2050 (United Nations Population Division 2018). Growing populations moving to urban centers will make smart cities a necessity to efficiently manage public resources, and remove excessive burdens to the climate, energy, environment, and living conditions.

In addition to improving efficiency and cost savings, smart city technology is also used to address public safety concerns. Ranging from the deployment of *ShotSpotter* technology for swiftly responding to gunshots in public streets to mining traffic data and social media posts for tracking virus outbreaks, numerous companies have partnered with cities to offer solutions for economic security, public health management and counter-terrorism. Given the increasing adoption of smart city technology in cities across the world, there is a need to ensure that citizens can benefit from smart cities without sacrificing their privacy and fundamental rights.

III. Unfettered law enforcement access to smart city data threatens individual privacy

i. Constant surveillance leads to chilling effects on the freedoms of expression, movement and association

The deployment of smart city technology could have a chilling effect on citizens' constitutionally protected freedoms. Passive and indiscriminate data collection generates an enormous amount of data, allowing authorities to engage in the mass retention of personal data collected by smart sensors. Automated bike counters can track thousands of bikes each day or tens of millions of trips per year, using only one type of sensor. Similarly, populous cities using automated license plate readers collect data on millions of cars travelling on highways each day. Despite being presented as innocuous measures to reduce crime or increase efficiency, this information can be used to create databases of individuals' activities, revealing sensitive information about personal associations, values and health conditions. Additionally, a recognition of this data collection can cause individuals to monitor their behavior for fear of isolation, retribution or raising suspicion.

Unfettered access to smart city data provides authorities with a proverbial time machine, allowing them to go back in time to watch, listen and investigate any event that takes place in public. In one case, a private security company provided footage to the Baltimore Police Department that connected aerial video from its wide-angle camera-equipped planes with ongoing police reports (Reel 2016). This allowed law enforcement to map the movements of all citizens — criminals and innocents alike. Listening devices, such as San Diego's *Shotspotters* that enable law enforcement to quickly detect and respond to gunshots, are also capable of recording conversations, inadvertently infringing on the privacy of individuals within range of the devices (Hill 2016). This collection of private information in public spaces can be used to create personal profiles of all individuals in a city. Furthermore, the lack of existing limitations on storage duration could allow for the creation of lifelong individual profiles.

Individuals change behavior when made aware of pervasive surveillance. An analysis of Wikipedia traffic following the popular 2013 Snowden revelations revealed a 20 percent decline in page views on terrorism-related articles, including those mentioning "al-Qaeda," "car bomb" or "Taliban" (Penney 2016). Similarly, the behavior of an entire population, even in a democratic society, could be changed by awareness of surveillance in a smart city. This is because people who believe that they are being watched engage in far more conformist behavior than those who believe they are acting without monitoring (Richards 2013, Panagopoulos 2011). Behavior monitoring in public spaces can deter people from exercising their fundamental civil rights, including the rights to free and anonymous expression. Publicly-deployed pervasive recording devices may also cause individuals to self-censor speech. One researcher found that surveillance can lead to a "spiral of silence," wherein "individuals, motivated by fear of isolation, continuously monitor their environments to assess whether their beliefs align with or contradict majority opinion" (Stoycheff 2016). This is particularly dangerous because it can silence minority opinions (Penney 2017).

The privacy risks associated with smart city data are not limited to the data directly tied to an individual. Anonymized and aggregated datasets still pose privacy risks because it is technically feasible to re-identify individuals from such datasets. In 2007, researchers from the University of Texas at Austin re-identified anonymized individuals from the Netflix Prize dataset. The records from the Netflix dataset were combined using statistical techniques with records from IMDB to reveal the identities of the anonymized Netflix records (Narayanan and Shmatikov 2008). Further research has explored the ease with which records from numerous sources can be combined to de-anonymize datasets and identify individuals (Narayanan and Shmatikov 2019). Therefore, if smart city data is aggregated with other internal or publicly-accessible databases, various panoptic scenarios are possible, leading to a world where all aspects of a citizen's public life are captured, stored and accessible to various government and private entities indefinitely.

ii. Misuse of smart city data can threaten dissent and democratic rights

The tools of mass monitoring used to create smart cities, such as cameras, geo-tracking, pattern recognition, and predictive analytics, can be misused to restrict dissent and democratic rights. Using surveillance to suppress political dissent is already a reality in many countries. There exist several recent examples of smart city projects being used to meet political goals in a manner that violates citizen's essential rights to freedom of expression and association (Privacy International 2017). In Beijing, the government uses a *grid management system,* where city and individual household data is gathered in a central database and analyzed with artificial intelligence to detect and respond to trends of social unrest. The Rio Operation Centre also sheds light on the danger of creating an increasingly policed public space, where control rooms designed to tackle natural disasters, also allowed city authorities to violently suppress the pre-World Cup protests. Even in the U.S., the New York City Police Department drove cars, equipped with automatic license plate readers, past mosques to report on every attendee (Crump 2014). In the U.K., a 90-year old pensioner was placed on a hotlist merely for attending political demonstrations to make sketches. Smart sensors granting authorities unregulated access to personal data dramatically increase the potential of data misuse for political reprisals, blackmail, or even voyeurism by automating mass tracking previously done by individual police officers (Gurman 2016).

IV. Smart cities expose individuals to privacy harms by third parties

i. Individuals cannot opt out of smart city data publicly collected and shared with third parties

Public data collection removes an individual's choice in providing personal information to third parties. Smart city data, such as individual data collected about movements through a city, may be valuable to for-profit entities, such as advertisers. While the current state of commercial technology requires an individual to be in the frame of the same camera for tracking to occur, recent research has shown that a network of cameras can be used to automatically track an individual within sight of any camera (Chu and Hwang 2014). Thus, with advancements in technology, a network of city cameras can track individuals in all places where cameras are deployed. When equipped with facial recognition, cameras deployed to obtain information on general pedestrian traffic flow can also allow advertisers to specifically target an individual by knowing their identity, activity patterns, and the places they visit. Although individuals regularly share personal data voluntarily in exchange for services, they have a choice in what they share and can use privacy control mechanisms to manage the extent of data sharing. No options, however, are available to individuals who would like to opt out of the data collection and storage features of the smart city.

ii Malicious actors can exploit the availability of large volumes of personal smart city data

Long-term personal data storage by numerous entities presents multiple points of attack for malicious actors to gain access to sensitive information. Since 2005, hackers have stolen sensitive data from major retailers, such as Target and Home Depot. As citizens walk around in a smart city dense with sensors tracking their activities, there is a risk that nefarious entities with access to such a rich dataset may use it to track specific individuals. With sensors becoming more widespread, this may include increasingly sensitive personal information. Already, underground sensors are used to monitor pathways, track customers' paths, and generate off-grid electricity to power street lights. These non-intrusive vibration sensors can monitor footsteps to detect the individual occupants of public spaces and analyze gait patterns to obtain pedestrians' health information (Shirakawaa et al. 2013, Lam et al. 2016). In the absence of limitations on data retention and storage, millions of city residents and visitors remain vulnerable to breaches of personal data.

V. Examining the need for policy interventions to protect smart city privacy

i. Proposed technical solutions are insufficient in the absence of clear policy guidelines

Several papers have provided general technical recommendations on mitigating the privacy risks associated with deploying smart city technology. Such proposals tend to focus on one of two areas of risk mitigation: regulatory requirements to govern appropriate use of smart city data or technical requirements to prevent unlawful use of smart city data. These two categories have also been differentiated as process-oriented privacy protections and data-oriented privacy protections (Eckhoff and Wagner 2018). Process-oriented privacy protections according to the authors, focus on establishing best practices for deploying and using smart city technologies. This includes testing the devices, auditing access to the data by different individuals, and transparency with communities. Enforcing provisions of process-oriented privacy protections would likely require government regulation. Data-oriented privacy protections, however, are enforceable at the technical level. Technologies such as homomorphic encryption and zero-knowledge proofs can be used to prevent, at a technical level, nefarious actors from access private data. Once implemented, these technologies do not require government regulation to enforce.

Most recommendations to address the privacy concerns of smart city technologies fall into the latter category, focusing on the technical means to protect private data, rather than establishing guidelines regulating government access and use of the data. Recommendations that do attempt to address the role government should play in regulating smart city technologies generally do not explicitly define or limit the role of government in making use of smart city data.

In a survey of current smart city uses, Cui et. al. identified the necessary steps for governments to consider in preserving smart city privacy (Cui et. al. 2018). These include the government being held responsible for carefully considering which data is open and who has the right to access the data as well protecting data and model development. Although this survey identifies areas for the government to regulate smart city privacy protections, it does not identify concrete steps and criteria for governments to consider in developing regulations. Khatoun and Zeadally focus on technical solutions to address privacy challenges and argue that government intervention in the form of legislation is important to guarantee the implementation of their proposed strategies (Khatoun and Zeadally 2017). While these suggested technical recommendations allude to the necessity of government regulations to protect privacy, they do not propose any specific guidelines for cities to consider when adopting smart city technologies.

Finch and Tene argue for the need to introduce general principles to address smart city privacy protections (Finch and Tene 2014). They mention four needed developments: greater access to data for citizens, data featurization, de-identification, and enhanced transparency. Increasing data access for citizens and enabling citizens to benefit from their own data will engender consumer trust. They also contend that deidentification is a valuable privacy protection even in spite of work that shows it is possible to re-identify individuals in de-identified datasets (Narayanan and Shmatikov 2008).

Despite these general recommendations, there remain significant gaps in the adoption of smart city technology by cities across the world and corresponding legislation or policies to mitigate the harms associated with exploiting personal smart city data. Rather than proposing specific technical regulations, which would inevitably require updating as technology changes, this paper proposes a set of essential privacy provisions for municipalities considering the adoption of various forms of smart city technology.

ii. Existing privacy and legislative frameworks lack effective policies to prevent harms

It is evident that privacy mechanisms, policies and legal standards governing smart city data have not kept pace with technology. Consideration of privacy protection appears to be largely missing from the solution offerings of major private companies partnering with cities around the world to deploy smart city technologies. For instance, a policy framework on responsible data use for a smart city project in Toronto makes general references about embedding data privacy into all deployed systems, but does not establish comprehensive access and use limitations for smart city data (Sidewalk Labs 2018). This further illustrates that local jurisdictions implementing smart city solutions must take the lead in concurrently adopting a set of clear privacy protections to minimize current and future damages to their citizens.

Even among the most privacy-conscious jurisdictions in the US, there is an absence of effective legislation to mitigate smart city privacy concerns. California is considered a pacesetter for privacy protections, and is hailed as having adopted the “nation’s best privacy law” (Leno and Anderson 2018), the California Electronic Communications Privacy Act (“CalECPA”) in 2016. CalECPA establishes stricter privacy protections for electronic data than its federal namesake. Under the law, authorities may access electronic device information with consent from the device’s “authorized possessor.” CalECPA’s warrant requirement applies to compelling an entity other than the device’s authorized possessor or owner to provide electronic communication or device information. However, this leaves law enforcement access to smart city sensor data unrestricted. Smart city devices are owned by private companies or government agencies rather than individuals. Thus, under CalECPA, these device-owning entities can consent to law enforcements’ access of an individual’s data, rather than the individual himself/herself. Moreover, existing California legislation also does not prohibit the sharing or sale of personal information, except for health and financial information.

In many cases, the constitutional protections regarding smart city data are either unclear or non-existing. In the United States, two Supreme Court cases indicate long-term surveillance using electronic data may constitute an unlawful search and seizure under the Fourth Amendment. *United States v. Jones* did not clearly establish whether long-term G.P.S. tracking constituted an unlawful search and seizure. Justice Sotomayor’s concurrence indicated that she believed such tracking goes against an individual’s expectation of privacy. A similar case regarding unwarranted law enforcement use of cell site location information, *Carpenter v. United States*, held that a warrant is required for police to access cell site location information from a cell phone company. The data collected in *Jones* and *Carpenter* is similar but not equivalent to smart city data, thus any decision of the Court may not be binding on establishing limitations of law enforcement access and use of such data. The lack of clear and consistent Constitutional protections for personal smart city data in many parts of the world adds to the urgency of why cities must adopt a set of fundamental guidelines to prevent the exploitation of smart city data and the erosion of essential liberties.

VI. Proposed Policies: Adopting a fundamental set of privacy provisions

The lack of privacy protection for individuals navigating smart cities can be addressed by adopting a set of fundamental privacy provisions. We argue for five key policy provisions: 1) differentiating personally identifiable data from de-identified data; 2) creating a warrant requirement for personal smart city data; 3) prohibiting the sharing of personally identifiable information collected by smart city sensors; 4) adopting data minimization requirements; and 5) introducing private and public enforcement mechanisms. We urge local jurisdictions and municipalities to adopt these policies and guidelines into their respective legislative frameworks to ensure that their citizens can continue to benefit from smart city technology while being protected from the threats they pose.

i. Differentiating personally identifiable data from de-identified data

Many, if not all, smart city sensors collect sensitive information about individuals, even if they were not necessarily designed to do so. For example, video cameras, installed to monitor traffic patterns, will inevitably capture faces, license plates numbers, and other sensitive information. In order to differentiate between data that is used to assess general city problems and trends, and that which can be used to specifically target an individual, it is important to distinguish between two separate categories of data: raw data that contains personally identifiable information and processed data that has been de-identified to remove personally identifiable information.

Personally identifiable information refers to any data that identifies an individual – either by itself or by linking to information in the public domain. In the context of a smart city, this information can include items such as a photograph of someone's face, the MAC address of a smartphone, or a vehicle registration number. We refer to smart city data that contains such information as "personal smart city data." This information is clearly the most sensitive, as it is easy to pick out an individual from a dataset with such identifiers present. On the other hand, de-identified data has had personally identifiable information and any indirect identifiers removed or manipulated to break the linkage to real world identities. As technology continues to progress, new types of information can be used to identify individuals – from their gait, to the shape of their ears, or even what WiFi networks they pass every day on their way to work. Consequently, our understanding of what data may be personally identifiable is likely to change with technological advancements. However, the distinction between personally identifiable data and de-identified data remains important since it can be used to invoke different access and use requirements, allowing for a balance between individual privacy protections and legitimate uses of aggregated city data.

ii. Acquisition and use of smart city data

Personal smart city data and de-identified data should be given different access and use limitations. In general, while there are benefits to making certain categories of aggregated and anonymized data publicly available, a warrant should be required for any access to personal information. Some exceptions to this requirement should be provided in order to allow law enforcement to use smart city sensors to keep citizens safe. Additionally, one limitation placed on the usage of de-identified data is that law enforcement cannot use this information to attempt to identify an individual.

Absent a specified exception, for law enforcement to obtain personal smart city data, probable cause must exist and a warrant must be obtained. Specifically, the warrant must describe the information it is seeking, a time frame, a geographic scope, and which smart city sensors have that information. This meets both the standard warrant requirements of particularity, but also aims to limit the ability of law enforcement to obtain large amounts of personal data unrelated to the criminal investigation at hand. Other laws such as the California Electronic Communications Privacy Act ("CalECPA") provide similar requirements to access the contents of electronic communications. Given the rapid proliferation of smart cities, it is important that such procedural restrictions be extended to access personal smart city data unless an exception applies.

We identify three key exceptions to the warrant requirement that allow law enforcement to optimize efficiency while maintaining protections for individuals. These exist in cases of: (1) exigent circumstances; (2) limited short-term surveillance; and (3) minor traffic violations. Emergencies, such as a high-speed chase or an active kidnapping, require law enforcement to act quickly. Since warrants can often not be obtained swiftly enough to address such concerns, authorities should continue to have the ability to act in the best interest of the public under exigent circumstances. Additionally, allowing law enforcement to request data about a suspicious individual within a 24-hour timeframe, so long as they have not made such a request about that individual within the past 90-days, allows law enforcement to obtain a small amount of information about an individual without a warrant. Due to the constrained time frame, however, they are restricted in their ability to assess trends or habits, or engage in long-term digital surveillance without a warrant. Lastly, red-light cameras and speeding sensors are already in use across numerous countries. Smart city regulation should not attempt to prohibit the use of such devices since they serve a legitimate purpose in public safety with minimal, if any, implications for citizen's privacy.

In this work, we do not argue for general access or use restrictions on de-identified data. There is great social and economic value in identifying long-term trends and aggregated behaviors and it is important to maximize the utility of such data without sacrificing individual privacy. Therefore, cities, law enforcement, and third parties may use de-identified datasets indiscriminately for their legitimate tasks, such as creating crime heat maps, or identifying problematic intersections. However, the legitimate use of de-identified data should not be used as excuse by government entities or law enforcement to sidestep the warrant requirement by attempting to re-identify an individual.

Since it is possible to sometimes make inferences about private individual information using de-identified aggregate data, we recommend both the use of technical safeguards and legal restrictions on using de-identified data to attempt to identify individuals. When de-identified smart city data is publicly released, technical precautions should include the use of differential privacy mechanisms that add random "noise" to database query results to prevent making inferences about individual records while enabling the detection of overall trends in the data (Dwork 2008). It may still be possible to re-identify individuals despite removing all personal identifying information from a dataset (Narayanan and Shmatikov 2008). Therefore, technical safeguards should be accompanied with a legal restriction on attempting to use de-identified data in conjunction with other available data to identify specific individuals. Such a legal restriction should, however, have exceptions to allow for good-faith research into finding and responsibly disclosing potential security vulnerabilities associated with smart city technologies. There exists precedence for allowing security research that may otherwise breach legal restrictions in the exemptions adopted by the Library of Congress to the provision of the Digital Millennium Copyright Act (Federal Register 2015). Another example can be found in changes to the Wassenaar Arrangement, which imposes export control requirements on hacking tools, but allow researchers to engage in vulnerability incident response (Waterman 2017). Permissive use of the de-identified data, with these limitations, balances the needs of governments and citizens alike.

iii. Prohibition on sharing personally identifiable smart city data

Despite emerging concerns regarding the sale of private information to third parties in the aftermath of recent privacy scandals such as Cambridge Analytica, the sharing of smart city data is largely being ignored across the world. Currently, there are few or no restrictions limiting the selling or sharing of this information by data controllers, which are entities that lawfully collect smart city data. This is particularly concerning provided that citizens have no ability to "opt-out" short of never stepping foot in the city. To strengthen individual privacy rights with regards to private actors, entities that partner with city municipalities should be prohibited from sharing or selling personal smart city data.

In order to promote public-private partnerships, smart city regulations should not place unnecessary use restrictions on entities deploying smart city sensors. Such entities should retain the ability to sell or share smart city data that has been aggregated or anonymized, as discussed above. However, protections must be provided for individuals in order to balance the interests of private entities with the legitimate concerns of citizens.

iv. Data minimization requirements

Smart city sensors collect data on every person in the city. Because of this, incorporating data minimization procedures is necessary to reduce the likelihood of data falling into the wrong hands. For private actors, all personally identifiable information must be de-identified or deleted within two years of collection, unless retained pursuant to a lawful preservation order or in accordance with other legal evidentiary requirements.

Law enforcement may retain personal smart city data, obtained by a duly issued warrant, for the duration of their investigation, or, if the investigation results in a criminal conviction – for as long as the individual remains incarcerated in connection with the case. Once that period of time expires, law enforcement should be required to permanently delete all copies of smart city data within a specified period of time.

This data minimization requirement limits the impact of a catastrophic data breach. An attacker will not be able to find personal records for an individual stretching back more than two years. This also establishes a standard for all individuals in the city that data collected on them will not be kept indefinitely.

v. Enforcement of privacy protections

We have identified policy recommendations to protect citizens' privacy in a smart city. For such policies to be effective, they should be accompanied by the establishment of robust enforcement mechanisms to protect against potential abuses by both private actors and government agencies. This could include the creation of a private right of action for any individual whose data is sold or used in violation of the aforementioned policies – along with statutory, punitive, and actual damages. Such a two-pronged approach can ensure that individual privacy is protected while allowing cities to maximize efficiency and improve public safety.

Three key protections could help ensure that government entities will comply with the aforementioned provisions. First, the State Attorney General should have the power and authority to compel any government entity to comply with the Act. Second, evidence obtained in violation of the outlined policies should not be admissible in a criminal proceeding. Our recommendations include comprehensive exceptions to the warrant requirement, such as in exigent circumstances and for short term surveillance. This provision protects individual civil rights by disincentivizing violations of this policy. Lastly, there should exist a private right of action for data controllers to challenge a warrant if they believe it is facially insufficient or overly broad. Without this provision, data controllers who collect and store smart city data, would lack standing to challenge a warrant and would have to turn over data without an opportunity for judicial review. These three protections strike the balance between allowing law enforcement to keep citizens safe but also protect individual civil rights in a criminal trial.

Enforcement of the aforementioned recommendation should also empower individual citizens to be in direct control of their smart city data. This should include a notice requirement – any individual whose personal smart city data is obtained pursuant to a warrant must be notified of the warrant and the data obtained. To that end, an individual should also be able to challenge a warrant for their smart city data. Secondly, individual citizens should have a private right of action if any personal smart city data is sold, shared, disclosed, or used in violation of the bill. This should be accompanied by equitable and declaratory relief, statutory and punitive damages, and reasonable attorney's fees. Statutory damages, in addition to actual damages, would create a strong monetary incentive for data controllers to comply with these provisions. The attorney's fees provisions incentivize a plaintiff's attorney to take possible violations and ensure that citizens are afforded their full civil rights.

Finally, we aim to draw a fine line to avoid exposing data controllers to unlimited civil liability. This can be achieved by creating a two-year statute of limitations for an individual to bring suit against a data controller. We also recommend providing good faith defenses for data controllers that reasonably act to comply with a valid request for information.

VII. Balancing privacy protections with improving efficiency and public safety

Our proposed policy recommendations for governing smart city data aim to both protect individual privacy while allowing the two biggest stakeholders, government entities and data controllers, to benefit from the collected data. This ensures that government entities and data controllers can use de-identified data to improve efficiency and safety while simultaneously protecting individual civil rights.

Government agencies should be required to apply similar restrictions to smart city data that already exist for electronic communications in many US states. The three broad warrant exceptions, discussed above, strike the necessary balance between privacy and security. Law enforcement can continue to respond to imminent threats while safeguarding individual civil rights. Although we propose creating an additional level of scrutiny for accessing personal smart city data, this would not impede the vital work of various government and law enforcement agencies since such entities are free to use the de-identified data for improving municipal efficiency and safety without jeopardizing citizens' right to privacy.

The same balance must be struck for data controllers: they can use the data for beneficial purposes but face common-sense civil liabilities for violating the aforementioned privacy provisions. Data controllers should be protected by a two-year statute of limitations and good faith defenses for compliance.

Data controllers should be held responsible for ensuring that de-identified data cannot be reasonably re-identified and have the ability to challenge an overbroad warrant. We aim to strike a crucial balance – allowing data controllers to indiscriminately collect and analyze public data, while ensuring an individual's expectation of privacy in public. The common-sense requirements for data controllers protect them from excess liability while incentivizing compliance and protection for personal smart city data.

Our proposed policy recommendations thread the needle – they allow cities to benefit from smart city data while preventing the creation of surveillance cities. Any additional regulations will pose a further burden to relevant stakeholders, either government entities or data controllers. Yet, the exceptions and limitations on civil liability that we have proposed aim to both ensure individual privacy and allow stakeholders to benefit from the data they collect.

VIII. Conclusion

The deployment of smart city technology in cities across the world provides an unprecedented opportunity to optimize resource management and improve living conditions via data-driven decision-making. The mass deployment of sensors that could be used as surveillance devices, however, present serious privacy risks to individuals who cannot opt-out of such indiscriminate data collection. Proposed technical solutions, while helpful in mitigating certain privacy risks, provide insufficient protection in the absence of fundamental policy-based privacy provisions. Furthermore, existing legislative frameworks and privacy policies being considered by companies fail to adequately address these risks. Cities across the world must take action to confront the gaps in their current laws to safeguard the privacy of their citizens. Adopting the set of fundamental privacy provisions outlined in this paper will provide a necessary first step in ensuring that citizens reap the benefits of smart cities without inadvertently enabling the creation of surveillance cities.

References

Bonte, Dominique. 2017. *Smart Cities and Cost Savings*. ABI Research, October 19, 2017.

Chu, Chen-Te and Hwang, Jenq-Neug. 2014. "Fully Unsupervised Learning of Camera Link Models for Tracking Humans Across Non-overlapping Cameras." *IEEE Transactions on Circuits and Systems for Video Technology* 24, No. 6, June 2014: 979–94. https://doi.org/10.1109/TCSVT.2014.2302516.

Crump, Catherine. 2014. "A Cop May Be Following You Everywhere." *CNN*, October 6, 2014. Accessed September 27, 2018. http://www.cnn.com/2014/10/06/opinion/crump-police-surveillance/index.html.

Cui, Lei, Gang Xie, Youyang Qu, Longxiang Gao, and Yunyun Yang. 2018. "Security and Privacy in Smart Cities: Challenges and Opportunities." *IEEE Access*, 1–1. https://doi.org/10.1109/ACCESS.2018.2853985.

Dwork, Cynthia. 2008. "Differential Privacy: A Survey of Results." In *Theory and Applications of Models of Computation*. Edited by Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li, 4978:1–19. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-79228-4_1.

Eckhoff, David and Wagner, Isabel. 2018. "Privacy in the Smart City: Applications, Technologies, Challenges, and Solutions." *IEEE Communications Surveys Tutorials* 20, no. 1 (First quarter 2018): 489–516. https://doi.org/10.1109/COMST.2017.2748998.

Federal Register. 2015. "Exemption to Prohibition on Circumvention of Copyright Protection Systems for Access Control Technologies." *Copyright Office, Library of Congress*, October 28, 2015. https://www.federalregister.gov/documents/2015/10/28/2015-27212/exemption-to-prohibition-on-circumvention-of-copyright-protection-systems-for-access-control.

Finch, Kelsey and Tene, Omer. 2014. "Welcome to the Metropticon: Protecting Privacy in a Hyperconnected Town." *Fordham Urban Law*

Journal 41, 1581. Available at: https://ir.lawnet.fordham.edu/ulj/vol41/iss5/4.

Gurman, Sadie. "Across US, Police Officers Abuse Confidential Databases." *Associated Press*, September 28, 2016. https://www.apnews.com/699236946e3140659fff8a2362e16f43.

Hill, Christie. 2016. "The Color of Surveillance in San Diego." *Medium*, October 20, 2016. https://medium.com/@SDACLU/the-color-of-surveillance-in-san-diego-4dce43abe67c.

Khatoun, Rida and Zeadally, Sherali. 2017. "Cybersecurity and Privacy Solutions in Smart Cities." *IEEE Communications Magazine* 55, no. 3 (March 2017): 51–59. https://doi.org/10.1109/MCOM.2017.1600297CM.

Lam, Mike, Mostafa Mirshekari, Shijia Pan, Pei Zhang and Hae Young Noh. 2016. "Robust Occupant Detection Through Step-Induced Floor Vibration by Incorporating Structural Characteristics." *Dynamics of Coupled Structures*, Volume 4. https://doi.org/10.1007/978-3-319-29763-7_35.

Leno, Mark and Anderson, Joel. 2018. "California Electronic Communications Privacy Act (CalECPA) - SB 178." *ACLU Northern California*, April, 2018. https://www.aclunc.org/our-work/legislation/calecpa.

Narayanan, Arvind and Shmatikov, Vitaly. 2008. "Robust De-Anonymization of Large Sparse Datasets." *Proceedings of the 2008 IEEE Symposium on Security and Privacy* 111, 121.

Narayanan, Arvind and Shmatikov, Vitaly. 2019. "Robust De-Anonymization of Large Sparse Datasets: A Decade Later." May 21, 2019. Available at: http://randomwalker.info/publications/de-anonymization-retrospective.pdf

Panagopoulos, Costas. 2011. "Social Pressure, Surveillance and Community Size: Evidence from Field Experiments on Voter Turnout." *Electoral Studies* 30:2, 353–57.

Penney, Jon. 2016. "Chilling Effects: Online Surveillance and Wikipedia Use." *Berkeley Technology Law Journal*, 31 (1) 117.

Penney, Jon. 2017. "Internet Surveillance, Regulation, and Chilling Effects Online: A Comparative Case Study." *Internet Policy Review*, Volume 6, Issue 2. DOI: 10.14763/2017.2.692.

Privacy International. 2017. "Smart Cities: Utopian Vision, Dystopian Reality". October 2017. Accessed August 18, 2018. http://privacyinternational.org/report/638/smart-cities-utopian-vision-dystopian-reality.

Reel, Monte. "Secret Cameras Record Baltimore's Every Move from Above." *Bloomberg*, August 23, 2016. https://www.bloomberg.com/features/2016-baltimore-secret-surveillance/.

Richards, Neil M. 2013. "The Dangers of Surveillance." *Harvard Law Review* 126, Volume 1934.

Shirakawaa, Tomohiro, Moto Kamiura, AkihiroTakagi, and Hiroshi Sato. 2013. "Analysis for the Gait Patterns of Healthy Subjects During March." *Procedia Computer Science* 24: 167-174. https://www.sciencedirect.com/science/article/pii/S1877050913011824.

Sidewalk Labs. 2018. "Responsible Data Use Policy Framework." *Waterfront Toronto,* Version 0.2, May 1, 2018. Accessed October 3, 2018. https://sidewalktoronto.ca/wp-content/uploads/2018/05/Sidewalk-Toronto-Responsible_Data_Use_Framework_V0.2.pdf

Stoycheff, Elizabeth. 2016. "Under Surveillance: Examining Facebook's Spiral of Silence Effects in the Wake of NSA Internet Monitoring." *Journalism & Mass Communication Quarterly*. https://doi.org/10.1177/1077699016630255.

United Nations Department of Economic and Social Affairs. 2018. *World Urbanization Prospects: The 2018 Revision*. United Nations Population Division.

Waterman, Shaun. 2017. "The Wassenaar Arrangement's Latest Language Is Making Security Researchers Very Happy." *Cyberscoop*, December 20, 2017. https://www.cyberscoop.com/wassenaar-arrangement-cybersecurity-katie-moussouris/.

Weber, Mario, and Ivana Podnar Žarko. 2019. "A Regulatory View on Smart City Services." Sensors (Basel, Switzerland) 19, no. 2 (January 21, 2019). https://doi.org/10.3390/s19020415.

Wajeeha Ahmad is a recent graduate from the Technology and Policy Program at the Massachusetts Institute of Technology (MIT). She conducted her graduate research with the Internet Policy Research Initiative and the Advanced Network Architecture group within the Computer Science and AI Lab at MIT. Prior to graduate school, she received her Bachelors of Science in Mathematics with Computer Science from MIT and worked in research and industry roles at the MIT Senseable City Lab, the World Bank and Google X among others.

Elizabeth Dethy is a recent graduate from the Masters of Engineering program in Computer Science at the Massachusetts Institute of Technology. She conducted her graduate research with the Internet Policy Research Initiative within the Computer Science and AI Lab at MIT. Prior to graduate school, she worked as a software engineer and consultant for several years. She received her Bachelors of Science in Computer Science and Electrical Engineering from MIT.

Acknowledgements

The authors would like to thank Professor Daniel Weitzner, director of the Internet Policy Research Initiative at MIT, and Professor Alvaro Bedoya, director of the Center on Privacy & Technology at Georgetown Law, for insightful comments during the preparation of this paper.