Privacy in a Smart City

Martin Peterson[a] & Barbro Fröding[b]

[a] Texas A&M University, Department of Philosophy, martinpeterson@tamu.edu
[b] KTH Royal Institute of Technology, Department of Philosophy and History of Technology, barbro.froding@abe.kth.se

Early View publication date: 28 May 2024
DOI: http://dx.doi.org/10.5324/eip.v18i1.5572

An increasing number of cities around the globe aim to become Smart Cities through the implementation of a range of Information and Communication technologies (ICT), AI applications and cloud-based IoT solutions. While the underlying motivations vary, all such transitions require large amounts of data. In this paper, we articulate and defend two claims about privacy in a Smart City. Our first claim is that some level of systematic data collection and processing is ethically permissible. However, there is an upper limit for what is permissible: We suggest that it is never permissible to collect and process data that significantly undermine people's autonomy. Our second claim specifies when such minor privacy infringements that do not significantly undermine people's autonomy are permissible: We suggest that the only policies legitimized by the first claim are those that promote the collective good.

Keywords: smart city, privacy, autonomy, natural right, collective good

Introduction

Imagine that the interview for your dream job is scheduled to start imminently. To get to the company's headquarters, you have to cross a busy street. The light turns red, but you dash across and make it to the interview on time. Your decision to jaywalk did not cause any harm or put anyone at risk and you soon forget this minor wrongdoing.

However, had this happened in Rongcheng in Northeast China, your behavior would have put a dent in your social credit record (Nast 2021). In Rongcheng all citizens start with 1000 social credits, and cameras scattered around the city record questionable and rewardable behavior and adjust people's social credit scores accordingly. Low social credit rating has a tangible negative impact on citizens' quality of life. Consequences include, for example, not being eligible for certain jobs or schools. Good social credit rating, on the other hand, brings rewards, such as reduced energy bills, lower or no required deposit when booking hotels or renting bikes, and more favorable interest rates for loans offered by state-owned banks.

Rongcheng is not unique. There are numerous similar social credit systems in China (Nast 2021; Yang 2022). While specifics vary, the core idea of these systems is that citizens have social credits that can be increased or reduced as a result of desirable or undesirable behavior.

The aim of this article is to discuss the notion of privacy in a Smart City. The social credit system in Rongcheng illustrates an extreme version of surveillance which most Western cities would not engage in, but it is nonetheless *one* way of using Smart City technology. Other examples, to be explored in this article, include mapping people's movements in public spaces to further public safety, health, or sustainable transportation. These technologies pose a very real threat to privacy, but it is hard to deny that some degree of surveillance can boost the collective good, either by triggering behavioral changes or by helping city planners shape the city in ways that enable people to make better choices. This raises the question of what, if any, level of systematic data collection and processing is ethically acceptable in a Smart City.

In what follows, we will defend the following two claims about privacy in a Smart City:

1. Some level of systematic data collection and processing is ethically permissible. This sometimes infringes on people's privacy, but there is an upper limit for what is permissible: We suggest that it is never permissible to collect and process data that significantly undermines people's autonomy.
2. Not all privacy infringements significantly undermine people's autonomy. Our second claim specifies when such minor privacy infringements are permissible: We suggest that the only policies legitimized by the first claim are those that promote *the collective good.*

The first claim pertains to what type and quantity of data may be collected. The second claim pertains to the reason why such systematic data collection is sometimes permissible. Our principles are lexically ordered, meaning that if a policy violates the first principle it is impermissible even if it stands to promote the common good. We do not believe that these claims are particularly controversial, and we deliberately use the vague terms "somewhat" and "significantly" as the purpose of this paper is not to establish the *precise threshold* for when privacy intrusions are permissible.[1] The modest goal of this paper is merely to identify the relevant *dimensions* for determining when a privacy intrusion is ethically permissible. The main contribution of this article is thus that it connects a series of deeply held philosophical intuitions with real-world examples of privacy-invasive policies.

For the purposes of this paper, we will adopt a straightforward but admittedly somewhat imprecise working definition of privacy. As we use the term, "privacy" refers to a plausibly legitimate claim to freedom from interference, intrusion, and control. In a similar vein, "autonomy" refers to a person's ability to self-govern herself and make uncoerced decisions.

What is a Smart City?

The concept of a Smart City refers to the use of certain technologies, such as AI and cloud-based interconnected CCTV, with the intention to make cities function "better" in certain respects. However, no consensus as yet exists on how, exactly, the term "Smart City" should be defined. Some scholars focus primarily on the role of *technology* (see e.g. Deren, JianJun & Yuan 2015), while others emphasize a combination of technical *and* social factors. Kitchin, for example, prefers the latter type of hybrid definition. Drawing on the work of Allwinkle and Cruickshank (2011), Kitchin writes that a Smart City

> encompasses policies related to human capital, education, economic development and governance and how they can be enhanced by ICT. In this scenario, networked infrastructures are enabling technologies, the undergirding platform for innovation and creativity that facilitates social, environmental, economic, and cultural development. (Kitchin 2014: 2)

In another paper, Kitchin (2015: 131) adds that "a smart city is one whose economy is increasingly driven by technically inspired innovation, creativity and entrepreneurship, enacted by smart people". Other notable efforts to articulate definitions of what a Smart City is have been proposed by Caragliu et al. (2011) and Albino et al. (2015). More recently, Ziosi et al. (2022) have offered a four-component framework through which they analyze a series of ethical concerns about Smart Cities.

Since the focus of this article is on ethical aspects of Smart Cities, a definition of Smart City that articulates the combination of technology and the social good (e.g. life quality improvements) is the better fit. In what follows we will use Kitchin's (2014) hybrid definition, which emphasizes the combination of technical and social aspects, as our working definition.

Any city that could be referred to as "smart" requires large amounts of data. Indeed, data is central for the "production of sophisticated data analytics for understanding, monitoring, regulating and planning the city" (Kitchin 2014: 12). Data needs to be collected, processed and used to, for example, train and improve AI applications that map and predict people's behavior and choices. Some data is open and freely accessible, but in other instances, data is personal and therefore sensitive, and hence subject to legislation like the General Data Protection Directive (GDPR). [2]

Examples of data that are useful for developing a Smart City include information about mobility and traffic (cars, bikes, public transport etc.), how the city is used by people (finding patterns of how people move around and what they do, e.g. parks, culture, sport), air quality, garbage collection, energy consumption, heat leakage, and other environmental aspects, data concerning health, data generated by the schools, geodata, etc. (Sourbati & Behrendt 2021; Csukás and Szabó 2021). More controversial potential effects include increased security and more efficient crime prevention (Laufs et al. 2020). The data can be collected in various ways, for example via "digital cameras, sensors, transponders, GPS, kiosks, meters, personal devices, appliances, social networks, and machine-readable objects" (Ryan & Gregory 2019: 5).

However, while the development and use of various technological solutions stand to make cities more sustainable (socially, environmentally, and economically), there is also a risk that the introduction of new technologies may have negative effects. When introduced unquestioningly, these technological solutions might incur the danger of exacerbating pre-existing problematic aspects of old solutions as well as generating new ones (Ziosi et al. 2022: 19).

Consider, for example, how AI and machine learning can both solidify existing inequities and bias as well as introduce new ones.[3] Or how advanced surveillance technologies introduced to increase security and government services can undermine citizen privacy by undercutting their autonomy.

As noted earlier, our focus here is on ethical issues related to privacy, but it is worth keeping in mind that the incessant need for data in Smart Cities gives rise to a wide range of ethical challenges. These include ethical issues related to transparency, trust and trustworthiness, fairness and bias, responsibility and accountability, anticipatory governance and nudging as well as inclusion. (Hollands 2008; Kitchin 2014; Ryan & Gregory 2019; Ziosi et al. 2022). In addition, numerous value conflicts will likely arise. This includes conflicts between trust and privacy; efficiency and privacy; and security and transparency.

Two approaches to privacy

In *Understanding Privacy*, Daniel J. Solove (2008) argues that no single definition of privacy can capture all the nuances of this multifaceted concept. Consider, for instance, the notion of privacy adopted by the U.S. Supreme Court in *Roe v. Wade* (1973), the precedent that guaranteed the right to abortion in the United States for nearly fifty years (until a new ruling in 2022 left the issue to states). In *Roe v. Wade*, the court defined privacy as a woman's absolute right to control of her body. A notion of privacy that focuses on absolute control of something is very different from the suggestion that privacy is an important but non-absolute moral value that can be balanced against other moral values. Solove accounts for these radically different notions of privacy by invoking Wittgenstein's (1953) classic concept of family resemblance: Many slightly different but somewhat similar conceptions of privacy exist, but there is no set of necessary and sufficient conditions that covers them all. Each notion is related to some part or aspect of the others, but without any common core that is shared by all.

In contrast to Solove's broad approach to privacy, Helen Nissenbaum (2004) limits her discussion to information technologies that trigger privacy concerns. Her account is not necessarily at odds with Solove's thesis that privacy should be understood in terms of family resemblance, but notably, she is critical of the influential view that the function of privacy is to restrict access to intimate, sensitive, or confidential information, or to curtail intrusions into spheres considered to be private or personal. Nissenbaum points out that we sometimes ought to be concerned about privacy-intrusive information technologies even when the information is *not* intimate, sensitive, confidential, private, or personal. She makes her point by way of the following example:

> Online profiling is troubling, even when the information gathered is not sensitive (excludes credit card information, for example) and when it takes place on the public Web [the Internet]. According to the framework, therefore, it seems that public surveillance is determined not to be a privacy problem. Because this conclusion is at odds with the intuition and judgment of many people, it warrants more than simple dismissal. (Nissenbaum 2004: 134)

The upshot of Nissenbaum's analysis is that traditional accounts of privacy are too blunt and therefore unable to capture important nuances. To rectify this shortcoming, Nissenbaum proposes an alternative, *contextualized* notion of privacy:

> One of the key ways contextual integrity differs from other theoretical approaches to privacy is that it recognizes a richer, more comprehensive set of relevant parameters. […] the crucial issue is not whether the information is private or public, gathered from private or public settings, but whether the action breaches contextual integrity. The use of credit cards and the emergence of information brokers, along with a host of technical developments, however, have altered patterns of availability and flow in well-known ways. (Nissenbaum 2004: 151 – 52)

We agree with Nissenbaum that it is important to consider the context in which a potential privacy violation occurs. What it is permissible to do in one context may not be permissible in another. We also find the "integrity" part of Nissenbaum's contextual integrity approach plausible. Private information should, at least in many cases, remain within the domain it is obtained. Medical information about a patient should typically not be shared outside the medical domain, in particular not with the government or the patient's employer.

However, while Nissenbaum explicitly conceives of privacy as a *right*, we propose that it is more fruitful to conceive of privacy as a *moral value*. In the literature, it is common to define a right as something absolute that cannot be trumped by competing moral considerations, whereas values are less rigid and can be trumped by other factors. On this view, values can be balanced against other types of moral considerations, but rights (of the absolutist kind we have in mind) are rigid in the sense that they do not permit trade-offs. Locke (1689/1823), Kant (1785/1914) and Nozick (1974) are examples of influential thinkers who advocate absolutist rights of this type; Locke and Nozick (unlike Kant) also claim that rights are "natural" in the sense that they exist independently of social processes and legal norms.[4] Non-absolutist notions of rights, which we will set aside here, have been proposed by Rawls (1999) for example and, depending on how he is read, Ross (1930).

If you have a right to privacy in the absolutist sense, and your privacy is violated, then you have been wronged regardless of whatever mitigating circumstances might apply. But if we think of privacy as a moral value existing alongside other comparable values – such as wellbeing, justice, or safety – then a measure that limits a person's privacy can sometimes be permissible if it leads to sufficient gains in wellbeing, justice, or safety.

It is helpful to elaborate on the (admittedly somewhat oversimplified) contrast between privacy conceived as a right and privacy conceived as a value, by considering Locke's (1689/1823) claim that rights are "natural" in the sense that we are born with our rights; they are not given to us by society or moral conventions. On this view, rights cannot justly be restricted by the government or by any kind of social policies or conventions: "The natural liberty of man is to be free from any superior power on earth, and not to be under the will or legislative authority of man, but to have only the law of Nature for his rule" (Locke 1689/1823: 114). However, Locke does not claim that we have a natural right to privacy, nor does he claim that we lack such a right. In fact it is an open question whether Locke thought that such a right could be derived from his theory, as it is essentially a theory of ownership. We own our bodies, as well as the products we can create by mixing our labor with things not previously owned by anyone, but information about ourselves may not always be covered by this notion of ownership. If you visit a public space where a camera records your presence, then nothing that belongs to you has been taken from you, on the Lockean account. Simply by being present in a public space you consent to letting others know that you were there. The fact that cameras systematically recorded your presence makes no difference, at least not if the cameras were clearly visible or you were informed about the cameras in advance. If you know that you are being monitored by a camera and you voluntarily choose to visit a place, then you consent to being recorded. If you do not like the idea of being recorded by a camera in a public space, the Lockean account advises you to stay home behind the walls of your own property. On this view, society has no obligation to provide citizens with public spaces that protect their privacy for the same reason as society has no obligation to provide us with anything but law and order.

Nissenbaum does not express any support for Locke's account of rights. Although she never explicitly says so, it is reasonable to infer that she conceives of privacy as a socially constructed right. The *locus classicus* for the constructivist approach is Brandeis and Warren's (1890) seminal article in the *Harvard Law Review*, in which they argue that all persons in the U.S. have a legal right to privacy. Brandeis and Warren justify their right to privacy by reminding the reader of how "new" technologies such as photography and newspapers impact people's lives. A photo of a person taken at the wrong moment, in the wrong situation, can have unpleasant consequences. The function of a socially constructed right to privacy is to prevent this and other similar unwelcome events from occurring.

We prefer to think of privacy as a value rather than a right for at least two reasons. The first is that the value-based approach coheres better with some considered moral intuitions we take to be important. Consider, for instance, contact tracing during a pandemic. To claim that contact tracing would be morally impermissible because people have an inviolable right to privacy would not be a logically contradictory position, but it does not square well with our considered intuitions. In our view, it is plausible to think that at least *some* privacy intrusive measures, such as contact tracing, would be morally justified for limiting the spread of a dangerous virus. By thinking of privacy as one value amongst others, it is easier to explain why some (but not all) privacy invasive measure are permissible – this is so because the value of privacy is being balanced against other considerations.

Our second reason for conceptualizing privacy as a value rather than an absolute right is that even if people consent to sharing their personal information it can still be morally impermissible to use this information. If privacy is an absolute right, the fact that rational and well-informed people consent to having their personal information collected is strong evidence that no right has been violated. (Admittedly, some rights cannot be waived even by consenting, rational, and well-informed adults. In 2006 the German cannibal Armin Meiwes was convicted of killing and eating a consenting volunteer. Although the victim's consent was documented on video, the court did not find that he could waive his right to life.[5]) However, under *ordinary* circumstances, an individual who consents to share private information in a Smart City waives the right to privacy. This makes it hard to address what we take to be legitimate privacy concerns. Consider, for instance, a future society in which it is common knowledge that face recognition technologies are widely used by the government in public spaces and matched with medical records. If you choose to visit a public space in such a society, you implicitly consent to sharing your medical records with the government. If privacy is a right, this would not violate your right to privacy, but if privacy is a value, we can explain why the way your medical information is used might be morally impermissible. The explanation is that privacy is a value, not an absolute right that is nullified when people consent, and this value has to be balanced against other values. Although Nissenbaum does not use this vocabulary, her emphasis on contextual integrity arguably captures a similar intuition.

The connection between autonomy and privacy: an example

In this section we argue that a person cannot be fully autonomous without having a certain degree of privacy, meaning that privacy is a necessary condition for autonomy. If you, for instance, know that the government is using face recognition technology in public spaces to monitor people's whereabouts, this will in some cases limit your autonomy. Even if you are an ordinary law-abiding citizen, you might rationally choose to change your behavior if this technology is used. Suppose, for instance, that face recognition technologies are used to keep track of people entering and exiting liquor stores. Although it is not illegal or immoral to buy liquor, the mere fact that you are recorded when you do so might alter your behavior. If so, your autonomy has been reduced by this technology, and every technology that limits a person's autonomy reduces her privacy in a morally relevant sense.

To further illustrate the relationship between privacy and autonomy, it is helpful to consider a real-world example. We have purposely selected a case in which the privacy intrusion is obviously unethical. This allows us to explain the connection between privacy and autonomy without first taking a stand on whether some privacy intrusions, if any, are morally permissible.

In November 2020, democratic elections were held in Myanmar. However, the day before the newly elected members of parliament were to be sworn in, the military initiated a coup d'état. As a result, extensive political unrest broke out. To crush the protesters the military used drones, hacked cellphones and deployed

tracing software to "…track the protesters live location and listen in to their conversations" (Beech 2021). These actions were obviously unethical, but that is not the reason for introducing this example. Our point is that the privacy intrusions orchestrated by the military also restricted peoples' autonomy. Many citizens in Myanmar refrained from exercising their democratic right to express their true views about the unlawful military coup because they knew that they were being watched. The behavior of the government in Myanmar had a negative effect on peoples' autonomy in a non-trivial sense. The systematic use of drones, cellphone hacking, and tracing software made it less attractive for citizens with dissenting views to exercise their autonomy.[6] Options that protesters otherwise would have chosen were no longer available to them, as they knew they were being monitored and risked being severely punished.

There is no doubt that the government's use of smart technology made many citizens in Myanmar afraid of exercising their democratic rights. And rightly so, it *was* risky to express dissident views because of how the government used its smart technology. People did not merely feel and believe their autonomy was being limited, their autonomy was in fact curtailed. If we consider autonomy to be a moral value that governments are typically not permitted to curtail, the slogan *No autonomy without privacy* neatly explains why the government of Myanmar violated people's autonomy through not respecting their privacy. Admittedly, the Myanmar coup is an extreme example, but the general point can be applied to more mundane examples. Our point is simply that autonomous decision-making requires some degree of privacy, and this explains why privacy intrusions that *significantly* restrict people's autonomy are profoundly morally controversial.

Individual autonomy vs. the collective good in the Smart City

We will now turn to discuss our second claim, namely the idea that it is sometimes ethically acceptable to enforce privacy invasive policies that limit people's privacy *somewhat*, but not significantly, if this promotes the collective good. As mentioned earlier, we purposely use the vague terms "somewhat" and "significantly", as the purpose of this article is not to establish the precise threshold for when exactly a privacy intrusion is permissible. We take it to be uncontroversial that *some* violation of people's privacy could be permissible if this would contribute to a significant public good. In such cases, privacy is not an absolute right but rather a moral value that has to be balanced against other values as suggested earlier. To support this claim, we will consider three cases in which we think it would be permissible to violate people's privacy in order to promote a greater good. The greater good comes in the areas of safety, public health and substantial environmental gain.

Example 1: The public safety case

Our first example is the "Camera in Beeld" (Camera in picture) system used by the Dutch police since 2016 for connecting public and private CCTV cameras into a nationwide system for crime prevention (Hofmans 2021). Business owners and private individuals with CCTV cameras are encouraged to give the police permission

to access images showing areas such as streets and parking lots that are captured by privately owned cameras. No one is required to sign up, the system keeps an electronic register with the location of the cameras registered by businesses and homeowners on a voluntary basis. The police cannot livestream the information on their screens, but videos and pictures can be accessed retroactively after an incident has occurred. As of July 2021, about 215,000 business owners and 44,000 home owners have joined the Camera in Beeld project. These privately owned cameras are connected to 19,419 public cameras, which enables the police to, for instance, retroactively watch a burglar who ran away from a crime scene even if the burglar avoided public spaces. The Dutch law that regulates the Camera in Beeld project gives the police permission to store the images for 28 days, but so far this rule has been frequently violated. Pictures captured by public cameras have been reported to be stored for an average of 38 days (Snijders et al. 2020).

There are, as far as we know, no systematic surveys of the public's acceptance of the Camera in Beeld system. Some people might be surprised by the lack of public resistance, given that few would deny that the widespread use of private and public CCTV cameras connected in a nationwide system limits peoples' privacy. It is, for instance, not illegal or immoral to walk home drunk from a bar, but the Camera in Beeld system makes it possible to retroactively follow your behavior on your way home if the police suspect you of having committed a crime.

Does the Camera in Beeld system reduce people's autonomy? According to our working definition, this depends on the extent to which the system impacts our ability to self-govern ourselves and make uncoerced decisions. Self-governance and coercion come in degrees. It seems fair to say that to some (rather small) extent, you are coerced to behave in a certain way if you know that your movements are recorded by cameras. The cameras limit your privacy, and this impacts your degree of autonomy. However, the impact on autonomy is relatively small. The Camera in Beeld system does not limit autonomy so much that an outright ban is morally required. On the contrary, the privacy intrusion might be morally permissible, because the Camera in Beeld system also promotes the public good. In many cases, privacy must be balanced against the public good, and the Camera in Beeld system seems to be an example in which the public good might very well win over privacy because the privacy intrusion is not significantly large.

Example 2: The Covid-19 Pandemic

During the Covid-19 pandemic, several cities around the globe analyzed data from smartphones to assess to what extent people complied with restrictions (Grantz et al. 2020; Kraemer et al. 2021). In Sweden, the Public Health Authority collaborated with Telia Crowd Insights to measure compliance with its "work-from-home-if-you-can" policy. Telia Crowd Insights gathered data from mobile phone networks to map how people moved around in Stockholm and to what extent they used public transport (Frick 2020). By comparing different time periods (i.e. before the recommendation was issued and afterwards), the authorities gained an important tool for understanding the impact their recommendations had on behavior and how compliant people were. This information informed the decision not to introduce

further restrictions. Importantly, data was anonymized and aggregated before being processed and thus it was not possible to trace any information to specific individuals. Further, the type of data collected is not classified as sensitive data under the GDPR. However, there was no consent and no possibility to opt-out.

Was people's autonomy violated? As before, this depends on whether tracking of (anonymized) smartphones impacted people's ability to self-govern themselves and make uncoerced decisions. As far as we can tell, this was not the case. However, it might nevertheless be the case that people (per our working definition of privacy) had a potentially legitimate claim to freedom from intrusion and control in this case. If so, their privacy might have been violated to some (presumably small) degree. Not every violation of privacy leads to a violation of autonomy; we are merely claiming that you cannot be fully autonomous without having a certain degree of privacy, as privacy is a necessary condition for autonomy. However, even if this was a minor violation of privacy, it seems clear that legitimate concerns about the public good trumped legitimate, but less weighty, concerns about privacy.

Example 3: Environmental Sustainability in Smart Cities

Stockholm, like many other cities, has adopted goals for reducing greenhouse gases and pledged to be fossil-free by 2040. The city plans to introduce a city-wide fossil-free fleet of buses and shuttles to achieve this goal. However, it will not be possible to reach the goal by merely *introducing* such vehicles; uptake must also be boosted. To effectively lower the thresholds that stop some people from choosing public transport, the city would need to get a better handle on how, when, and for what purposes people travel around in the city. One efficient way to collect such data points is to track smartphones. In addition to planning public transport, this data would also allow the city to even out traffic flows and plan necessary road work more efficiently.

Our ethical analysis of this case is similar to that of the second. Tracking anonymized smartphones does not significantly limit people's autonomy; if there is some potentially legitimate concern about privacy left to address, this would be a case in which concerns about the public good trump limited concerns about privacy. However, if smartphone data had not been anonymized, the conclusion would of course be different.

Note that our analyses of the three examples are based on the assumption that all data is handled by a trustworthy party (the municipality). In the second and third examples (but not the first), this is done under GDPR, which ensures that data is anonymized, aggregated, not shared in real-time, and not shared with third parties. Further, there is no risk that any individual is punished or otherwise restricted by the authorities as a consequence of the information gathered. The collected information is only used for making policy decisions. In this type of situation, it would, we think, be hard to argue that this is a substantial violation of privacy. According to our first principle, *some* level of systematic data collection and processing is ethically permissible given that it does not significantly undermine people's autonomy.

We of course recognize that the authorities might have other reasons to tread carefully and not routinely harvest data without consent or information. One such

reason is that doing so could undermine trust, even if the data are not privacy sensitive or substantially limit autonomy. Even though our examples do not entail significant intrusions, it could be argued that the practices described constituted *some form* of privacy intrusion. Not because the authorities sought to track any specific person or group, nor because of the overarching motive but simply because the implicit agreement between citizens and the state in a democracy is that law-abiding citizens are not tracked without consent while they go about their business in the public sphere. To suddenly do so jeopardizes trust. A decrease in trust in the authorities and in the public sector more broadly, could in turn undermine both the feasibility of the Smart City (as it presupposes trust between citizens and authorities) and, indeed, its purpose. Consider again the emphasis on citizen wellbeing and quality of life and the idea that the Smart City could be a vehicle to forward democracy and citizen engagement and participation. Diminishing trust could potentially weaken the achievability of such goals and thus be counterproductive to the establishment of the Smart City. Seeking consent would be very challenging from a practical perspective, but the authorities could be transparent through communication, such as publicly sharing the intention, practical purpose and method in advance in a manner comprehensible to non-experts, and by also making sure the information is not used outside the domain from which it was collected. This would echo central ethical research principles, as well as Nissenbaum's emphasis on contextual integrity mentioned earlier in the paper.

Our second principle concerns the purpose and the potential for an outcome that benefits the collective. It seems clear, in all three cases, that such a benefit exists. It is very much in the interest of society as a whole to find a strategy that promotes the safety of the general public and can manage a pandemic, and to find ways to stop global warming and other negative climate effects.

Is there perhaps a difference in urgency between these examples? One might intuit that data collection without consent or information is more permissible in the pandemic case. Here the situation was very time-sensitive. It was essential for the authorities to find out as soon as possible if their recommendation had the intended impact or if they needed to scale up. From an urgency perspective (understood as "how bad the potential consequences could be on a collective level"), however, the pandemic and the climate seem to be on a par. There is a difference in the temporal proximity between the measures (the data harvesting and processing) and the expected effect, but combatting both the pandemic and global warming are extremely urgent, and reducing those threats forwards the collective good.

Hence, none of the above examples violates the two principles proposed in this article. Harvesting data in the manner described here does not constitute a significant threat to anyone's privacy as it does not substantially undermine autonomy. In addition, these measures were taken in order to promote the collective good and could plausibly be expected to fulfill this goal.

Compare these three examples with the Rongcheng social credit system mentioned in the introduction. An obvious difference between these cases is the *magnitude* (or *degree*) of the privacy intrusion that limits people's autonomy. In the social credit case, it appears that the state is interested in controlling matters that, at least to some

extent, fall within the sphere where decision-making should be a private matter at the discretion of the individual. Even though making the "wrong" choice does not have any negative effect on other people, or compromise other people's quality of life, it is still recorded in a centralized computer system. But what about other actions like throwing the wrapping of a chocolate bar in the street, spitting on the pavement, or picking the flowers in the park? These are small-scale destructive actions that stand to undermine the collective good. But even so, tracking people through a social credit system (even assuming that it succeeds as a deterrent, which on some level then promotes the collective good) is a disproportionate response and constitutes a major intrusion into people's privacy. Such practices violate the first principle as they *significantly* undermine people's autonomy. Consequently, they are impermissible even though they might forward the common good (i.e. the condition formulated in the second principle). Indeed, social credit systems embody many of the ethical concerns regarding large-scale technology rollout. Evidently, not only privacy is at stake. Consider, for example, how such systems risk undermining trust, transparency, accountability, and the rule of law to mention but a few core features of an open, democratic society. We conclude that there is no plausible definition of the common good that would justify a social credit system of the comprehensive type implemented in Roncheng a few years ago.

Concluding discussion

Planning, developing, and regulating a Smart City requires collecting vast amounts of data, as noted at the beginning of our discussion. Some of this data will be sensitive, which leads to tensions between many central ethical commitments. On the one hand, city officials have to make sure that the city operates efficiently. On the other hand, it is also important to not undercut the privacy of the people living in the city. How can our discussion of privacy in a Smart City help navigate this challenging terrain?

Consider the first claim, according to which only privacy intrusions that do not significantly limit people's autonomy are permissible. This claim rules out scenarios where the sacrificing of a few people's autonomy would generate very high collective utility. At the same time it permits minor, even moderate, privacy intrusions for the collective good. Next consider the second claim, according to which the promotion of the collective good is the only permissible reason for intruding on people's privacy. This supports a commitment to justice, fairness and equality, meaning that it is not permissible to use one group to promote another. Combined, the two claims offer a loose structure for how to think about situations where there is tension between the development of the Smart City and citizens' privacy. The principles need to be supplemented by several ethical commitments on the part of the city, such as transparency, accountability, communication, safety and security, and that data is not used for additional purposes.

In summary, complying with the law is not always enough. It is key for the legitimacy of the Smart City to navigate ethical issues not explicitly regulated in the legal code. For citizens to want to share data there must be transparency, trust and

accountability. While not everyone will want to engage actively in the development of the city it does not seem reasonable to reduce the citizens to mere data providers.

Notes

[1] Exactly what counts as a 'minor' privacy infringement? It seems reasonable to believe that the boundary between minor and non-minor privacy intrusions is vague, but that does not mean that it is not possible to give clear examples from both categories. Some examples of minor privacy intrusions are discussed towards the end of the paper.

[2] Evidently one can imagine data which is open and sensitive at the same time, in particular when combined with other data, but we will not explore that here.

[3] For a good example from the healthcare sector, see Benjamin, R. (2019). Assessing risk, automating racism. *Science*, *366*(6464), 421-422. DOI: 10.1126/science.aaz3873; Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453. https://www.science.org/doi/10.1126/science.aax2342

[4] For a useful discussion of how Kant's theory of rights relates to those of Locke and Nozick (and others), see Beck (2006).

[5] See Friedman & Arold (2011).

[6] Note that this does not commit us to any form of technological determinism. We are merely claiming that the technology used by the government had a clear impact on people's autonomy, not that it fully determined it.

References

Albino, V., Berardi, U., & Dangelico, R. M. (2015). Smart cities: Definitions, dimensions, performance, and initiatives. *Journal of Urban Technology*, 22(1), 3-21. https://doi.org/10.1080/10630732.2014.942092

Beck, G. (2006). Immanuel Kant's theory of rights. *Ratio Juris*, *19*(4), 371-401.

Beech, H., (2021). *Myanmar's Military Deploys Digital Arsenal of Repression in Crackdown. New York Times*, March 1, 2021; retrieved March 15, 2022. https://www.nytimes.com/2021/03/01/world/asia/myanmar-coup-military-surveillance.html

Brandeis, L., & Warren, S. (1890). The right to privacy. *Harvard Law Review*, 4(5), 193-220.

Caragliu, A., Del Bo, C., & Nijkamp, P. (2011). Smart cities in Europe. *Journal of Urban Technology*, 18(2), 65-82. https://doi.org/10.1080/10630732.2011.601117

Csukás, M. S., & Szabó, R. Z. (2021). The many faces of the Smart City: Differing value propositions in the activity portfolios of nine cities. *Cities*, 112, 103116. https://doi.org/10.1016/j.cities.2021.103116

Deren, L., JianJun, C., & Yuan, Y. (2015). Big data in smart cities. *Science China Information Sciences*, 58(10), 1-12.

Frick. A., (2020). Mobildata visar svenskarnas resemönster under pandemin. *CIO, IDG.se* 2020-11-19. Retrieved March 15, 2022. https://cio.idg.se/2.1782/1.742402/mobildata-resmonster-corona

Friedman, L. M., & Arold, N. L. (2011). Cannibal Rights: A Note on the Modern Law of Privacy. *Nw. Interdisc. L. Rev.*, *4*, 235.

Hofmans, T., (2021) Bij politiedatabase Camera In Beeld zijn 280.000 openbare camera's aangesloten, [online] Tweakers. Retrieved November 14, 2021. https://tweakers.net/nieuws/184428/bij-politiedatabase-camera-in-beeld-zijn-280000-openbare-cameras-aangesloten.html

Hollands, R. G. (2008). Will the real Smart City please stand up? Intelligent, progressive, or entrepreneurial? *City*, 12(3), 303-320. https://doi.org/10.1080/13604810802479126

Kant, I. (1914). Die Metaphysik der Sitten. In Kants Gesammelte Schriften. Königlich Preußische Akademie der Wissenschaften, vol. 4. Berlin: Reimer.

Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, 79(1), 1-14. https://www.jstor.org/stable/24432611

Kitchin, R. (2015). Making sense of smart cities: addressing present shortcomings. *Cambridge journal of regions, economy and society*, 8(1), 131-136. https://doi.org/10.1093/cjres/rsu027

Kitchin, R., Cardullo, P., & Di Feliciantonio, C. (2018). Citizenship, justice, and the right to the Smart City. *The Programmable City Working Paper 41*. http://progcity.maynoothuniversity.ie/

Kitchin, R., Coletta, C., Evans, L., Heaphy, L., & Mac Donncha, D. (2017). Smart cities, urban technocrats, epistemic communities and advocacy coalitions: *The Programmable City Working Paper 26*. https://mural.maynoothuniversity.ie/9230/1/RK-Smart-2017.pdf

Laufs, J., Borrion, H., & Bradford, B. (2020). Security and the Smart City: A systematic review. *Sustainable cities and society, 55*. https://doi.org/10.1016/j.scs.2020.102023

Locke, J., (1689/1960). *Two treatises of government: A critical edition with an introduction and apparatus criticus by Peter Laslett*. Cambridge university press.

Nast, C., (2021). *The complicated truth about China's social credit system*. WIRED UK. Retrieved on November 14, 2021. https://www.wired.co.uk/article/china-social-credit-system-explained

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington. Law Review*, 79, 119.

Nozick, R. (1974). Anarchy, State and Utopia. Oxford: Blackwell.

US Supreme Court (1973), *Roe v. Wade*, 410 U.S. 113.

Rawls, J. (1999): The Law of Peoples, Harvard University Press, Cambridge, MA.

Ross, W. D. (1930/2002). *The right and the good*. Oxford University Press.

Ryan, M., & Gregory, A. (2019). Ethics of using Smart City AI and big data: The case of four large European cities. *The ORBIT Journal*, 2(2), 1-36; https://doi.org/10.29297/orbit.v2i2.110

Snijders, D., M. Biesiot, G. Munnichs, R. van Est, with the assistance of Stef van Ool and Ruben Akse (2020). *Citizens and sensors – Eight rules for using sensors to promote security and quality of life*. The Hague: Rathenau Instituut. https://www.rathenau.nl/sites/default/files/2020-02/Citizens%20and%20sensors.pdf

Solove, D. J. (2008). *Understanding privacy*. Harvard University Press.

Sourbati, M., & Behrendt, F. (2021). Smart mobility, age and data justice. *new media & society*, 23(6), 1398-1414. https://doi.org/10.1177/1461444820902682

Stockholms miljöbarometer (2021). *Miljöbarometern*, City of Stockholm, Retrived on March 15, 2022, http://miljobarometern.stockholm.se/klimat/utslapp-av-vaxthusgaser/

Yang, Z. (2023) China just announced a new Social Credit Law. Here's what it means, *MIT Technology Review*. Available at: https://www.technologyreview.com/2022/11/22/1063605/china-announced-a-new-social-credit-law-what-does-it-mean/ (Accessed: 09 April 2024).

Yigitcanlar, T., Kamruzzaman, M., Buys, L., Ioppolo, G., Sabatini-Marques, J., da Costa, E. M., & Yun, J. J. (2018). Understanding 'smart cities': Intertwining development drivers with desired outcomes in a multidimensional framework. *Cities*, 81, 145-160. https://doi.org/10.1016/j.cities.2018.04.003

Ziosi, M., Hewitt, B., Juneja, P., Taddeo, M., & Floridi, L. (2022). Smart Cities: Mapping their Ethical Implications. *SSRN 4001761*.