Mitigating Inference Risks with the NIST Privacy Framework

Christopher B. Landis
clandis@alumni.cmu.edu
Naval Postgraduate School
Monterey, California, USA

Joshua A. Kroll
jkroll@nps.edu
Naval Postgraduate School
Monterey, California, USA

ABSTRACT

The NIST Privacy Framework describes itself as a comprehensive approach to organization-wide privacy program management. However, inferences can yield sensitive information of identities or attributes from nonsensitive information. Privacy governance must protect this information. Although many people and organizations are expanding their privacy definitions to include inferences, our gap analysis reveals that the framework's mapped controls are insufficient for managing inference-driven risk. The framework does not attend organizational focus to privacy inference risk sufficiently to support its stated claim of comprehensive risk management. Applying the framework to past incidents where ostensibly protected information was re-inferred, we analyze how organizations can better mitigate inference-based privacy violations. Finally, we recommend detailed improvements to the framework's controls to account better for inferences. Our recommendations encompass augmenting and mapping additional privacy risk controls to increase implementing organizations' awareness of inference risks, updating controls that depend on protecting specific PII categories, and enhancing organizations' proficiency in translating legal and policy requirements into technical implementations.

KEYWORDS

inferences, privacy framework, re-identification, assessment, NIST, operational security, privacy controls

1 INTRODUCTION

Organizations face substantial challenges in practical privacy risk management. Although privacy enhancing technologies (PETs) can support better privacy outcomes, many organizations struggle to identify the nature and scope of their privacy risk and, consequently, their use of risk controls (including PETs). Risk governance within organizations is often ad-hoc and focused on legal compliance [2, 85]. Beyond cybersecurity, few tools exist to structure risk assessment and mitigation into repeatable practices. The preeminent privacy-focused risk management tool is the U.S. National Institute of Standards and Technology (NIST) Privacy Framework [56]. Through exemplary applications, we show that—while a risk-based approach to governance can improve substantive privacy outcomes—the current NIST Privacy Framework attends mostly to risks borne of unauthorized information flows and cybersecurity failures while under-attending to the problem of managing data *inferences*, a key problem on which the privacy research literature focuses. We believe a risk-based approach to governance can improve substantive privacy outcomes. That is, this gap limits both practical privacy risk management and the extent to which PETs can be brought to bear on practical problems.

Organizations face increasing levels of risk when using sensitive data. Despite the many legally mandated data protection practices,[1] technocratic decisions, like placing a privacy team in an advisory role rather than an embedded role, hinder privacy efforts [85]. They impede organizational risk professionals and leaders from foreseeing compliance problems. The NIST Privacy Framework is meant to aid organizations in addressing this gap. How well does the framework outline a risk management program that identifies and mitigates privacy risks of all sorts, arising both for individuals and organizations?

The problem of bringing privacy risk governance into practice is important because it drives data handling decisions during product design and deployment. Much privacy research has focused on the risk of *re-identification* of anonymized data, explicitly rejecting the popular notion that redacting personally identifiable information (PII) prevents re-identifying an individual in de-identified data [15, 64, 71]. As such, the community has largely refocused its efforts around provable guarantees and quantifiable metrics, either of formal indistinguishability properties like differential privacy [24] or of system-level information-hiding properties [33, 84]. Privacy practitioners face the problem of how to avoid the disclosure of sensitive facts about individuals and organizations.

The NIST Privacy Framework [56] describes itself as a comprehensive approach to organization-wide privacy program management through enterprise risk management. It aspires to enable dialogue among executives, managers, and practitioners so as to organize, assess, plan, and execute a privacy program in any organization, customized to that organization's needs. The framework asserts that cybersecurity can exist without data privacy safeguards but that protecting privacy is not possible without effective cybersecurity [56]. For privacy programs to effect their necessary practicalities, the framework maps categorized risks to controls [55]. Cybersecurity and the framework's mapped controls come short of sufficiently managing the risk of attackers inferring attributes [43], e.g., as in mosaic theory [4].

We analyze how well the NIST Privacy Framework identifies and controls the risks of sensitive data inferences by applying it to four inference incidents. To our knowledge, this is the first published work analyzing the effectiveness of the NIST Privacy Framework against real incidents of unexpected, sensitive inferences. In general,

Proceedings on Privacy Enhancing Technologies 2024(1), 217–231

https://doi.org/10.56553/popets-2024-0013

[1] Including those prescribed by the General Data Protection Regulation (GDPR) in Europe, comparable laws in the U.S. (e.g., the California Consumer Privacy Act (CCPA) and Virginia Consumer Data Protection Act (VCDPA)), and other jurisdictions (e.g., signatories to the Council of Europe's Convention 108).

we find that the framework does not guide organizations' attention to enough of the various inferences risks that they should consider mitigating, depends too much on an obsolete practice (i.e., protecting privacy by protecting PII categories), and has limited guidance for translating policies into technical implementation. We further make recommendations for improving the framework, focusing on its mapped controls. We consider inference risks to be one of two distinct forms and establish a taxonomy to describe them:

Re-identification Inferences Associating anonymized data with the original individuals, regardless of the presence of so-called PII [64].

Operational Inferences Inferring sensitive attributes about individuals and organizations. Note that even with perfect disclosure controls to inhibit re-identification inferences, operational inferences unrelated to PII may still be possible, depending on the threat model [29, 67].

After reviewing related work in §2, we consider how the NIST Privacy Framework provides organizations with a measure of their privacy risk management, especially regarding inferences, in §3. Our primary contribution in §4 is to validate the framework's capacity to identify and mitigate inference risks by applying it to four incidents. In §5, we propose recommendations for how organizations could have disrupted these incidents' structural and systemic causes. Our recommendations to organizations and for improving the NIST Privacy Framework can help organizations better mitigate inference-based privacy violations. Finally, we conclude in §6.

2 RELATED WORK

Privacy has become too complex for an individual to handle [39]. Instead, organizations must develop a comprehensive privacy risk governance program, and the NIST Privacy Framework exists to guide this effort. Definitions of privacy diverge, however, which complicates privacy protections [8, 23, 48]. Because legal requirements often drive organizations' privacy programs, we first review how the law accounts for inferences and individuals' rights related to inferences made about them. We then consider the information industry's competing interests in inferring sensitive information for profit. Complementary to the framework, we examine how PETs and privacy enhancing techniques can help mitigate inferences.

2.1 Accounting for Sensitive Inferences in Law

Unlike some notable U.S. privacy and privacy-related laws that do not mention inferences, we expect that newer privacy laws will incorporate inferences. Researchers have reported this shortfall in the Health Insurance Portability and Accountability Act (HIPAA) [79, 82], Children's Online Privacy Protection Act (COPPA) [42, 80], and Family Educational Rights and Privacy Act (FERPA) [81, 89]. Because the authors of some cybersecurity and privacy controls mapped to the NIST Privacy Framework derived their content from laws [58], we explore how laws account for inferences, including individuals' rights. The impact of inference attacks as a privacy risk are more significant than the way the law currently treats them. Of note, nearly all legal references to inferences that we found relate to re-identification, not operational inferences.

Holder [40] analyzes the GDPR and CCPA—the most progressive privacy laws available—for inferences and found that these laws were too ambiguous or contradictory, largely leaving aspects concerning inferences open to interpretation. He proposes that inferences drawn from seemingly innocuous data can be sensitive and should be protected by law and furthermore, that there is a hierarchy of harms that determines the sensitivity of personal data [40], amplifying that some privacy violation harms are intangible [75]. In other words, private data, regardless of whether inferred, needs protection to minimize the harm that they can cause to people.

Wachter and Mittelstadt [83] make recommendations to improve the GDPR to account for protecting against inferences by establishing an individual's "right to reasonable inferences", i.e., a "right on how to be seen" as complementary to a "right to be forgotten." They examine whether inferences are personal data, which would require the "right to know about, rectify, delete, object to, or port" [83], except that deletion can come with prohibitions on (re)creating new inferences. This points suspiciously to adding another PII category rather than considering individuals' personal privacy, which we address in §5.3.

They also investigate organizations' facilitation of subject access requests, i.e., individuals' rights to their inferred personal data. If allowing individuals to access their inferences would reveal an organization's trade secrets or intellectual property, then Wachter and Mittelstadt recommend that organizations not be obligated to comply [83]. On the other hand, Edwards and Veale [25] argue that the exception to divulging trade secrets or intellectual property should not be needed because a "right to an explanation" may not actually yield the expected result. Instead, they assert that the rights to deletion and porting are far more important.

As privacy laws evolve to incorporate inferences and individuals' rights pertaining thereto, organizations will need to adapt accordingly. As such, we highlight potential definitions and arguments discussed in our community and, as part of our recommendations in §5.4, include steps for enhancing organizations' proficiency in translating legal requirements into policy and technical implementations.

2.2 Inferences for Marketing

Although there are many business models that depend on inferences, including determining creditworthiness, insurance risk, and suitability for matching with an employer, mate, etc., perhaps the most prolific field is marketing [49]. For example, Google [47], Facebook [7], Proctor & Gamble [21], and other companies have yielded huge returns on their targeted advertising investments. While identifying people's potential shopping interests—based on inferences from their metadata or activities—is not typically a malicious act, there is consternation regarding whether society condones this type of inference [73, 85]. After its reputation experienced a temporary dip, Target [38] and other organizations learned to conceal how well they can infer individuals' attributes. Nevertheless, organizations continue to direct their advertising at specific users because it is more effective than generalized advertising [7, 73]. Although targeted advertising often involves proprietary methods and data that are never intended to be released beyond an organization's authorized analysts, organizations often share inferences with their contracted third-party analysts, circumventing sharing restrictions [44]. Our analysis shows how businesses depending on

the NIST Privacy Framework may not be identifying and mitigating inference risks as well as expected.

2.3 Privacy Enhancing Technologies and Techniques

The NIST Privacy Framework's mapped controls recognize various PETs and privacy enhancing techniques to help in the practical work of mitigating inferences. Even so, de-identification is non-trivial, especially for mitigating inferences [50]. Narayanan, Huey, and Felten [51] explore this challenge and present a "precautionary approach" to safeguarding privacy in data sets, positing that future re-identification capabilities may be unknowable. Nevertheless, technologies and techniques that place the burden of provable privacy on the data set owner may help in lasting ways for organizations desiring to protect their constituents' privacy.

2.3.1 Differential Privacy to De-identify Data. Differential privacy is one of the most prevalently discussed PETs today [19]. Garfinkel, Abowd, and Powazek [31] report on the U.S. Census Bureau's work in implementing differential privacy for the 2020 census, highlighting the challenges experienced. Surprisingly, their three recommendations for furthering the implementation of differential privacy—a technology—into the Bureau's work are not technical, but relate to organizational communications, bridging management and technical personnel [31]. This parallels many of the cybersecurity and privacy controls in NIST Special Publication (SP) 800-53, which affect multiple levels of implementing organizations, not just technical aspects [58], like PL-8 Security and Privacy Architectures, described in §4.2.4.

2.3.2 Applying Contextual Integrity to Inferences. Nissenbaum explores the concept of privacy protection in the vein of contextual integrity (CI). By definition, CI "is preserved when information flows generated by an action or practice conform to legitimate contextual informational norms; it is violated when they are breached" [59]. Analysts identify each flow by the five-tuple {Subject, Sender, Recipient, Information Type, Transmission Principle}. Applying the analogy of a food chain to explain the concept of inferences, primitive data are at the bottom of the food chain but are elevated when consumed by analytic processes, becoming inferences. From inferences, analysts can form additional, higher-level inferences or predict instantiation of primitive data. For example, Target inferred that a customer was pregnant by analyzing her shopping history (primitive data up to inference) and predicted which products she might consequently purchase to support her while pregnant (inference down to probable primitive data) [38]. Organizations' privacy programs should consider inferences' context from the subjects' perspectives so as to better protect individuals' privacy.

Martin and Nissenbaum [46] use CI to explore the important question of privacy in public, which relates to our question because it helps conceptualize how people may support or oppose inferences because of their data source or destination. Investigating the preservation and violation of CI, they discover how people can have privacy concerns regarding publicly available records in some contexts (i.e., inappropriate flows) but not others (i.e., appropriate flows). In their study, participants express the greatest concern for privacy violations when the source of the information flows were data brokers. Participants perceive data brokers' as having increased capability to "extract knowledge that is attractive to other stakeholders in various sectors," [46], that is, to infer sensitive information. The idea that sensitive information can change hands without there being a direct, observable flow increases the threat vectors to privacy risk. Knowing this should prompt data brokers to guard better against privacy violations from the inferences that they generate.

3 NIST PRIVACY FRAMEWORK

The ever-increasing frequency of privacy incidents evoked people's awareness of their need for better privacy protections. Cybersecurity compliance alone proving insufficient, NIST published its Privacy Framework [56]. Much of the work in protecting individuals' and organizations' privacy involves cybersecurity and many potential adopters were already using the NIST Cybersecurity Framework [53]. The Privacy Framework's authors acknowledge that they based the structure on the Cybersecurity Framework [56]. Beyond structure, the privacy framework's close relation to the cybersecurity framework is also evident: 57 of 100 privacy subcategories are the same as or have similar intended effect as a cybersecurity subcategory [54] and both frameworks rely on the same list of mapped controls [55, 58]. NIST acknowledges that checklist-like compliance requirements lend themselves readily to assessments that focus on compliance (not the framework's purpose) rather than on "achieving a positive outcome for privacy" [8]. As inferences are still possible, the inability to protect privacy by checklist compliance is evident in the framework.

While the NIST Privacy Framework [56] has achieved substantive goals for some adopters [57], our gap analysis reveals that the framework's mapped controls fall short in identifying and mitigating inference risks. Our recommendations to update the framework's mapped controls to account better for inferences will help organizations mitigate this risk.

3.1 Framework Structure and Implementation

The NIST Privacy Framework' guiding method to organization-wide privacy program management relies heavily on communications throughout the organization. It aspires to enable dialogue among executives, managers, and practitioners so as to organize, assess, plan, and execute a privacy program in any organization, customized to that organization's needs [56]. Because no statute or regulation requires framework adoption and its unique application per organization, using it for privacy assessment and compliance is unsuitable. Beyond compliance, many organizations have still benefited by it [57]. Even so, we find that it is ineffective for helping organizations mitigate inference risks.

Organizations will be most effective in implementing any control (and really, the whole framework) as a whole-of-organization effort with cognizance of interactions between the system and component levels. Protecting privacy is a sociotechnical endeavor. Approaching it strictly from a technical standpoint will lead to failure [45]. For example, SA-8(33) spans policy, personnel training, procedures, and technical components. Thus, a one-dimensional implementation could limit the other value of effecting this control.

The Privacy Framework [56] describes organizations' risk posture in terms of profiles, a core, and implementation tiers. The idea is for organizations to iterate through the framework's *core* subcategories in their *profile*, and select applicable mapped *controls* for each. They can self-assess their competency with the *implementation tiers*.

The framework encourages organizations to generate *profiles* to model their privacy activities and risk tolerance: one representing their current posture and at least one target profile modeling their desired privacy risk management goals. Organizations can use these profiles as a way to self-evaluate their requirements vis-à-vis posture. Each profile maps to the core's functions, categories, and subcategories.

The *core* embodies the highest levels of an organization's privacy activities as functions: identify, govern, control, communicate, and protect. Each function is divided into overarching, programmatic-level desired privacy outcomes called categories and then technical- and management-level desired privacy outcomes as subcategories. Each of the subcategories maps to various controls in NIST SP 800-53 [58] that serve as suggestions for mitigating the risk to that subcategory.

Organizations can assess their own capability to manage privacy risk using the Framework's *implementation tiers*, a non-compulsory progression of privacy program management proficiency. The four tiers are partial, risk informed, repeatable, and adaptive [56].

First published in 2005, NIST updated SP 800-53 [58] in 2020 to its current version, revision 5, to incorporate privacy controls, providing the framework with hundreds of security and privacy controls grouped into 20 "families," such as "Planning (PL)." Each family consists of base controls, each of which may have more specific control enhancements, identified in parenthetical suffixes to the control number. For example, AU-16(3) means "Audit and Accountability (AU) Control #16 (Enhancement #3)." A control enhancement's name is in the format: base control name, the pipe character '|', and the enhancement name. AU-16(3)'s name is "Cross-Organizational Audit Logging | Disassociability" [58]. Control names are necessarily succinct and do not, by themselves, provide sufficient descriptions of their controls. NIST maps most framework subcategories to specific controls and control enhancements [55]. Because mapped controls are suggestions to help achieve the subcategory posture and some overlap in their effects, each adopting organization needs to determine which controls would contribute positively to the organization's privacy goals. Selecting a control enhancement implies selecting its base control as well. Each control may have related controls and references cited to help practitioners choose the best controls. There is a balance between applying every mapped control to avoid deciding against a control—which is simple but potentially wasteful [9]—and not selecting an applicable control. The subjectivity necessary to implement the NIST Privacy Framework well makes it unsuitable as a compliance or certification framework.

3.2 Inferences in the Framework

Guarding against inference incidents is a complex endeavor, despite the framework having only one subcategory that refers to inferences by name or conceptually. Incidentally, this is the framework's only mention of "inference," in the Control (CT) Function, Disassociated Processing (DP) Category, Privacy Subcategory #3 (P3) (CT.DP-P3).

> **CT.DP-P3:** Data are processed to limit the formulation of inferences about individuals' behavior or activities (e.g., data processing is decentralized, distributed architectures). [56]

This definition applies to operational inferences, dealing with attributes rather than re-identification. Because many initially think of re-identification as the inference risk to be mitigated, some may question whether the NIST Privacy Framework should incorporate operational inferences. Consider that page 4 of the framework explains that problematic data actions, which lead to exposure of private data, impact individuals, "singly or in groups (including at a societal level)" [56]. Furthermore, the California Attorney General (AG) opines that inferences about identified individuals are personal data [5]. As such, organizations must influence others' inferred perceptions to mitigate the risk of operational inferences.

Subcategory CT.DP-P3 maps directly to nine controls and control enhancements, listed in Table 1 [55]. Although most of these controls focus more on technical than organizational aspects, three of the mapped controls (PL-8, PM-7, and SA-17) pertain to high-level actions that management can employ to bolster the organization's privacy risk management activities and give practitioners greater authority to execute policy.

The framework addresses re-identification in subcategory CT.DP-P2 but focuses on data processing so as to "limit the identification of individuals" [56]. It does not address inferences directly. Of its eight mapped controls, three also map from CT.DP-P3 (AC-23, SA-8(33), and SI-19) [55]. The other five controls relate to re-identification not necessarily involving inferences.

4 ANALYSIS

Past incidents demonstrate how organizations have suffered harms from inferences and elicit opportunities for learning how to mitigate them. We select four incidents that occurred before NIST published its Privacy Framework [56], knowing that the involved organizations could not have relied on the framework. This forms a baseline from which we evaluate how well the NIST Privacy Framework could have helped in mitigating inference risk. To understand the value of the framework for mitigating inferences we assess how well it could have helped mitigate the incidents described in §4.1. In extending our qualitative analysis, we find in §4.2 that the framework is largely insufficient to help organizations mitigate inference risk in general.

4.1 Inference Incidents

We refer to our four incidents as the EdX, NYC Taxi, Strava Heat Map, and Pizza Index incidents. Our taxonomy categorizes inference incidents as either re-identification or operational incidents. We considered all of the following incidents but chose these in §§4.1.1–4.1.4 to convey the essential incident structure for each inference type, enabling their generalizability to other inference incidents. The limitations of our selection process include the following: the lack of an exhaustive search for inference incidents, including those not documented publicly, and there may be edge cases not covered

Table 1: Controls from NIST SP 800-53 [58] Mapped [55] to Privacy Framework [56] Subcategory CT.DP-P3 and our summaries thereof. Adapted from [58]. Each control number is serialized, for example, AU-16(3) means "Audit and Accountability (AU) Protection #16 (Enhancement #3)" and its name has the format "Base Control Name | Control Enhancement Name."

Control Family	Control Number	Control Name
		Summary
Access Control	AC-23	Data Mining Protection
		Limit database queries and responses. Apply differential privacy or homomorphic encryption and accountability notifications.
Audit and Accountability	AU-16(3)	Cross-Organizational Audit Logging | Disassociability
		Implement privacy-enhancing cryptography to disassociate individuals' identities from audit information.
Identification and Authentication	IA-8(6)	Identification and Authentication (Non-Organizational Users) | Disassociability
		Decrease identity attribute visibility to transmitting parties.
Planning	PL-8	Security and Privacy Architectures
		Develop security and privacy architectures.
Program Management	PM-7	Enterprise Architecture
		Develop and maintain an enterprise architecture.
System and Services Acquisition	SA-8(33)	Security and Privacy Engineering Principles | Minimization
		Minimize collection, processing, and retention of private data.
System and Services Acquisition	SA-17	Developer Security and Privacy Architecture and Design
		Produce a design specification and security and privacy architectures.
System and Communications Protection	SC-2(2)	Separation of System and User Functionality | Disassociability
		Store applications separately from users' state information and interactions.
System and Information Integrity	SI-19	De-Identification
		Remove elements of private data from data sets and evaluate re-identification viability.

by this taxonomy. Nevertheless, our selections exemplify and are representative of privacy inference incidents in general.

The New York City (NYC) Taxi and Limousine Commission (TLC) passenger re-identification incident in §4.1.2 is exemplary of the ability for attackers to combine anonymous data with publicly available information from sources beyond the affected organization's control. Here, stalkers can map-out patterns of life without ever being within proximity of the victim. The concept of combining information follows the inference re-identification pattern in the Netflix prize [52], AOL search history [3], and Massachusetts Group Insurance Commission [70] data set incidents. Our other re-identification incident, EdX in §4.1.1, involves how a future employer could use EdX course performance as a hiring discriminator.

Of our operational inferences, the Strava incident in §4.1.3 demonstrates how an aggregation of decisions about protecting non-sensitive data can enable sensitive data inferences, a problem identified in the access control and privacy literature for decades [64, 66]. Similar inferences from smart meter activity (inferring providers' potentially unjust controlled blackouts) were possible because of an electric company's lack of protecting sensitive data [76]. Similar to the Pizza Index incident in §4.1.4 (in which Domino's inferred attributes of the Department of Defense (DoD) and its members), the Cambridge Analytica scandal involved inferring Facebook users' political affinity (an attribute), resulting in penalties for Facebook [27, 44]. Note that the victim is often not the primary organization involved in operational inferences.

4.1.1 Re-identification Inferences: MIT and Harvard EdX Data. In 2014, Massachusetts Institute of Technology (MIT) and Harvard published a data set of student online course performance on their EdX platform. They adopted an anonymization approach using quasi-identifier (QI)-based de-identification to achieve k-anonymity, facially consistent with the requirements of FERPA. Nevertheless, Cohen [14] was able to re-identify anonymized records by executing downcoding and predicate singling-out (PSO) attacks on the high-dimensional data set. He further showed that individuals without specialty education, experience, or tools, including "a prospective employer, a casual acquaintance, and an EdX classmate," could re-identify students' records [14].

Cohen speculated that the EdX experts de-identified the data set on one attribute after another, followed by selective row deletion. This "disjointed," rather than atomic, de-identification process degraded the initially-established *k*-anonymity [14]. Furthermore, the widespread technique of using QI-based de-identification on highly dimensional data sets can result in the preservation of sufficient information to enable re-identification whereas differential privacy techniques can overcome this [15]. If FERPA's authors had incorporated re-identification inference threats or if the EdX data set de-identification had used an atomic rather than a disjointed process, re-identification would have been much more difficult.

4.1.2 Re-identification Inferences: NYC Taxi Data. In response to a Freedom of Information Law request, the NYC TLC released its taxi data from 2013 [20]. The TLC anonymized the data by hashing the PII fields, including driver license and medallion numbers. Of course, a dictionary attack quickly resolved the finite range of these fields. From an inferences perspective, though, analysts can re-identify passengers and infer patterns of life, whether favorable, benign, or unbecoming, by combining the data with information that is publicly available or otherwise known. Published examples of privacy inference re-identification include stalking celebrities and re-identifying a man who visited multiple strip clubs [77].

Beyond the shortfall in de-identifying the data—exhibiting a disparity between the law and technical compliance [28, 60]—the opportunity for inferences poses a greater risk to individuals and liability to the TLC. Tockar [77] demonstrated how differential privacy techniques could have aided in preventing inference-based re-identification. However, Douriez et al. [20] demonstrated that differential privacy is insufficient after transforming the taxi data set into a moving object database of trajectories. Similar re-identification opportunities exist today [10, 18, 30].

4.1.3 Operational Inferences: Strava Heat Map. The Strava fitness app collects shared geo-location tracking data from its users and aggregates them to generate a global heat map of workout locations, the brightness directly proportionate to the popularity [37], enabling operational inferences. Only days passed between Strava publishing its heat map in 2018 before people noticed stark workout paths in the Afghanistan wilderness. An easy inference identifies them as coalition forward operating bases, inhabited by deployed military personnel. If the bases are well-known locations, this inference reveals no new information. On the other hand, the heat map also helped reveal a covert Central Intelligence Agency (CIA) site near Camp Lemonnier in Djibouti, an inference which could be fatal to the lives of those agents.

Strava de-identified the data for its heat map by aggregation. An analyst looking only at the heat map would be hard pressed to re-identify any individual. Nevertheless, many considered some of the inferences that observers made regarding the heat map to be violations of privacy [37]. Strava's position on the issue was that it only included tracked workouts in its heat map that users shared with them, corresponding with each user's privacy settings. Strava also noted that users could opt-out of sharing [37], thus blaming its users for revealing covert military and paramilitary bases. In other words, individuals' behaviors induced harm to an organization unexpectedly and that organization transferred that harm back to the individuals' inadvertent disregard for their organizations' operational security (OPSEC) policies, which might not have been specific enough to this risk. Still, perhaps the best place to limit the spread of sensitive information is at its source. Many communities use OPSEC principles but, incidentally, the U.S. Government (USG), especially the DoD, has the most explicit program policy publicly available [67].

4.1.4 Operational Inferences: Pentagon Pizza Orders. In 1998, the Washington Post interviewed the owner of 59 Domino's franchises in the DC area. He had established the "Washington Pizza Index," the number of pizza orders directly correlated with the level of USG activity [69], enabling operational inferences. Furthermore, based on the area of the franchise, he could conjecture whether the pizza was headed for the White House, congressional offices, or the Pentagon. Although this incident is relatively old, we select it because of its straightforward analysis and availability of relevant public information. As of the date of the interview, the record for the greatest number of pizzas ordered was by people working in the Pentagon during the Persian Gulf War [69].

During a war, one might expect the DoD to be working longer hours so an unusually high volume of pizza orders may be normal. However, extra hours devoted to planning in the days leading up to a major operation may reveal too much, potentially harming the success of the operation by signaling its imminence to adversaries. Likewise, a corporate build-up working toward the launch of a new product could give rise to similar inference opportunities. Also, the signals are not limited to food orders as parking lot fullness, employee shuttle activity levels, etc. can likewise provide opportunities for unwanted inferences. These operational inferences are a matter of OPSEC. As with the Strava incident, it may be best to limit the spread of sensitive information is at its source.

4.1.5 Lessons Observed. The themes that would have helped in better mitigating re-identification inferences were largely technical and those that would have helped in better mitigating operational inferences were largely sociotechnical. To identify and mitigate inference risks, one must understand them. In aggregate, these observations indicate the unlikelihood that the framework's authors considered such inference risks. The following contributed to the inference re-identification of people in de-identified data:

- retaining too many QIs (EdX, Netflix, MA Health)
- a lack of ability to translate legal requirements unambiguously into technical implementations (EdX, NYC)
- disjointed, rather than atomic, de-identification (EdX)
- lack of understanding de-identification methods (NYC, AOL)

Sociotechnical aspects, including human behavior and configuration decisions were the primary contributors to enabling operational inferences. Note that the harmed organizations (coalition forces, CIA, and DoD) in these operational inference incidents were involuntarily harmed by the behavior of their members (deployed and headquarters forces) and third-party organizations (Strava, Domino's, and the Washington Post). Both were unexpected OPSEC failures. The root of these observations is a lack of thought towards the potential consequences of actions, including:

- risks of sensitive inferences from personal data were not identified (all operational inferences)

- interaction of design choices with end users' defaults concealed risks (all operational inferences)
- irregular behavior in aggregate spurred actions to reveal sensitive information (Strava, Pizza, Smart Meter)

4.2 Framework Effectiveness via Controls Mapped for Inferences

This section constitutes our assessment of the effectiveness of the NIST Privacy Framework subcategory CT.DP-P3 and SP 800-53 controls to which it maps for mitigating the risks involved in the inferences described in §4.1. Table 1 lists these controls and their enhancements. For each control, we assess whether it or its enhancements apply and would have helped perceive and mitigate risk. By "applicable," we mean that if the control had been implemented properly, it would have meaningfully contributed to mitigating risk in that incident. If the control is applicable, we estimate its degree or likelihood of implementation. If we find that it was partially or fully implemented or likely implemented, we analyze why it fell short in mitigating the inference. We provide a summary of our findings in Table 2.

Table 2: A Summary of Our Analysis of the Effectiveness of NIST SP 800-53 Controls [58] Mapped [55] to Privacy Framework [56] Subcategory CT.DP-P3 Against the Incidents Presented in §4.1. Each cell's symbol in the middle four columns represents the degree to which the control was implemented. The last column indicates the lesson observed.

Control Number	EdX	NYC Taxi	Strava	Pizza Index	Lesson(s) Observed
AC-23	N	N	N/A	N/A	N/A
AU-16(3)	N	N	N/A	N/A	N/A
IA-8(6)	P	P	N/A	N/A	PII/QI
PL-8	L	U	F	F	TPTI
PM-7	L	L	P	P	TPTI
SA-8(33)	N	P	F	N/A	CPAI
SA-17	N/A	N/A	N/A	N/A	N/A
SC-2(2)	P	N	N/A	N	PII/QI
SI-19	P	P	P	P	CPAI, TPTI, & PII/QI

Legend: N/A - not applicable, **U** - unknown and unlikely, **L** - unknown and likely, **N** - not implemented, **P** - partially implemented, **F** - fully implemented, **CPAI** - the ability to combine anonymous data with publicly available information, **PII/QI** - the insufficiency of using solely PII categories and/or quasi-identifiers (QIs) to protect individuals from re-identification, **TPTI** - difficulty translating policy into technical implementation

4.2.1 AC-23: Data Mining Protection. As indicated in Table 2, none of the organizations implemented this control. Applying differential privacy well would have worked to safeguard both the EdX and taxi re-identification inference incidents; Cohen and Tockar both claimed as much [14, 77]. Limiting database queries or enabling accountability notifications for atypical database queries or accesses would not have helped because the data sets were made public. Applying homomorphic encryption [88], a PET suggested by this control, would not have helped either because the plaintext data were public. Publishing encrypted data would have had no value.

For generating its heat map, because Strava executed an aggregate query on its database, none of the PETs suggested by this control would have helped prevent the operational inferences. Arguably, however, aggregation clearly abated data mining opportunities in the heat map, but operational inferences were still possible. Similarly, for the Pizza Index, management deliberately disclosed the inference publicly during a newspaper interview [69]. This control is therefore not applicable for operational inferences.

AC-23 has no control enhancements.

4.2.2 AU-16(3): Cross-Organizational Audit Logging | Disassociability. AU-16 focuses primarily on the protection of users' identities when coordinating auditing with external organizations, like linking tables through blinded keys to disassociate logs from their data. Although the principles of this control's "disassociability" enhancement could apply, despite no organization implementing it, auditing does not specifically pertain to the incidents in §4.1.

Similar to AC-23, privacy-enhancing cryptography—any of several schemes that enables data processing without revealing the data [1, 6]—may have helped mitigate re-identification inference incidents. Specifically, private information retrieval [72] techniques may have been able to help in both the EdX and taxi incidents, if, for example, there was a query interface for researchers in lieu of making the whole "anonymized" data sets public. This control would not have helped mitigate the operational inferences so it is not applicable to the operational inferences.

4.2.3 IA-8(6): Identification and Authentication (Non-Organizational Users) | Disassociability. IA-8 focuses primarily on organizations' external users' interactions with an organization's information systems. Because each of the incidents involved publicly released data, non-organizational users do not authenticate to access the data. This base control does not apply; though, the principles of IA-8(6) could still apply.

Interestingly, each of the organizations in the inference incidents partially implemented a principle of this control, specifically, that they sought to "make identity attributes less visible to transmitting parties" [58], but their de-identification attempts were ineffective at or not applicable to mitigating the inferences. This is more evidence demonstrating the ineffectiveness of protecting privacy only by PII categories. In the EdX data, the QIs remaining after de-identification contributed to analysts' re-identification inference capability. For the taxi data set, the ability to combine publicly available information with released de-identified data enabled passenger re-identification. Strava's and the (perhaps inadvertent) Pizza Index's disassociation of individuals' identities through aggregation—like striving for k-anonymity with sufficient k—may be an application of a principle of this control, but was not applicable to and could not have mitigated the subsequent operational inferences.

4.2.4 PL-8: Security and Privacy Architectures. The NIST SP 800-53 describes security and privacy architectures as system-level manifests comprised of three components: requirements for protecting organizational information and PII, a description of how they support the enterprise architecture (explained in control PM-7 in §4.2.5), and assumptions regarding the system and its dependencies. Here, the presence of an architecture proved insufficient to mitigate the inferences, indicating organizational difficulty with translating intent into policies and policies into technical implementations.

For re-identification inferences, the data controller's architecture protects the data. While a private organization's system-level security and privacy architecture may not be publicly available, we found that EdX had a privacy policy dated February 6, 2012 that was still in effect when the 2013 academic year began [26], the year of data comprising the EdX data set. The policy comprehensively covered the entire student experience, including that private data may be made available publicly, "to the extent permitted by FERPA" [26]. As such, EdX likely had some type of system-level security and privacy architecture from which they derived this privacy policy, indicating apparent compliance, yet re-identification inferences were still possible. Similarly, NYC law governed the TLC's handling of PII. Specifically, chapter 5 of title 10 of the administrative code [12] defines PII and lists requirements for publicly disclosing that a security breach occurred; however, in 2013, there was no statutory requirement for agencies of the city to have privacy policies or architectures nor is there public evidence that the TLC had them. Therefore, the TLC likely had some level of awareness of its need to protect individuals' privacy but likely did not have a security and privacy architecture. In this analysis, it is clear that there was a gap in translation. Without additional internal details from these organizations, we speculate that the breakdown occurred in translating intent into policies, which would indicate the polices were insufficient, or translating policies into technical implementation.

For operational inferences, PL-8 differs from the other controls described thus far because the victim organization's architecture protects the data, not so much the data controller's. In the cases of both Strava and the Pizza Index, note that the U.S. DoD—generalized here from the "Coalition Forces" that would have been in Afghanistan in 2018—was the affected party, not Strava, the data controller, or any individual. The DoD definitely has a system-level security architecture [74] and privacy program [63], but neither has a role in preventing the inference of sensitive information, which the OPSEC program [78] governs. Apparently, the DoD OPSEC program was insufficient in translating these types of operational inference risks for DoD employees to take into account.

Although subcategory CT.DP-P3 maps to the PL-8 base control, framework adopters may choose to implement any of PL-8's two control enhancements:

(1) Defense in Depth. This enhancement focuses on applying a defense-in-depth posture [34] while developing and administering security and privacy architectures. Because the EdX and NYC Taxi data sets were made public, i.e., outside an organization's control, layering additional defenses would not have mitigated these inferences. Defense-in-depth could have contributed to informing an OPSEC program, however, in the Strava and Pizza Index incidents via additional safeguards to limit the signals coming from these DoD sites.

(2) Supplier Diversity. This enhancement focuses primarily on addressing the potential issues of a monoculture [86] and would not have applied in these incidents because a monoculture did not contribute to the incidents.

4.2.5 PM-7: Enterprise Architecture. In contrast with PL-8, PM-7 is a higher-level control that encompasses security and privacy—among a host of other organizational concerns—integrating systems that might each have their own security and privacy architecture. In these incidents, the presence of an architecture proved insufficient to mitigate inference risk, indicating organizational difficulty with translating policies into technical implementations.

Similar to a system-level security and privacy architecture, the existence of an organization's enterprise architecture may not be public knowledge. EdX did not indicate on its website that it had an enterprise architecture, but EdX probably had some type of overarching business model that would meet the spirit of this control. NYC's website included a 2016 job posting that enabled us to infer that the city has an enterprise architecture [13], though being a big and disparate organization, this may be a miscategorization. Even so, without access to these architectures, we are unable to determine how well they account for privacy risks. In contrast, the DoD—the generalized victim of the operational inference incidents—has a publicly available enterprise architecture [11]; however, it is not actually a specific architecture but a framework for developing architectures within the DoD [17]. Nevertheless, it mentions risk management and guiding security and information assurance requirements but not privacy [11] so we gauge this as a partial implementation. Without additional details from these organizations, we speculate whether the breakdown occurred in translating intent into policies or translating policies into technical implementation.

"Offloading" is PM-7's sole control enhancement. It recommends that organizations move all non-essential supporting functions to separate the functions from critical systems and data, applying a "least authority" principle [68]. There is also an implied assumption that the component or contracted organizations performing these functions would be experts therein and thus have better quality and more efficient security and privacy safeguards specific to those functions, thus decreasing the likelihood of privacy violation. While the degree to which EdX offloads non-essential supporting functions is unclear, governments typically delegate or contract many functions so it is likely that both NYC and DoD offloaded functions to some degree. For all of the above, there is no indication that offloading contributed to the incidents.

4.2.6 SA-8(33): Security and Privacy Engineering Principles | Minimization. SA-8 and its enhancements serve as an avenue through which organizations can incorporate security and privacy engineering principles. It also applies a lens by which the framework reminds its adopters to operationalize these principles at all stages of a system's life cycle. SA-8(33) focuses on minimizing PII, a concept we discuss further in §5.3.

The EdX case demonstrates how protecting privacy via confidentiality fails. It collected students' level of education, gender, and year of birth, each as optional fields and inferred a country of

residence from the user's public Internet Protocol (IP) address [14]. According to its privacy policy [26], EdX uses these data for at least nine different purposes, of which nearly all relate to some form of data analysis. Being optional, none of these PII fields matter to enrolling in and completing a course, the core function of the EdX platform. As such, EdX did not apply this security control. These extra PII fields enabled re-identification inferences by allowing the combination of QIs [14].

The NYC TLC partially implemented SA-8(33), its data containing unnecessary PII fields for drivers but not passengers. These fields enabled analysts to learn the annual income of re-identified taxi drivers [20]; however, minimizing passenger PII proved ineffectual at preventing re-identification because of the ability to combine publicly available information with precise time and location data in the data set. Interestingly, unlike normal taxi service, mobile app based ride share services (e.g., Lyft and Uber) require passenger identification to support payment and "other purposes."

For Strava, minimization via aggregation did not prevent operational inferences [37]. Even without PII present in or contributing to the heat map, operational inferences were still possible because of the information publicly available related to current events [36]. On the other hand, there is no indication that Domino's stored any of their customers' PII beyond transaction completion; such PII would not have been applicable to minimize to mitigate the operational inferences anyway. But for the other three incidents, the ability to combine publicly available information with the de-identified data sets enabled these inferences.

4.2.7 SA-17: Developer Security and Privacy Architecture and Design. This control guides adopting organizations in setting requirements for external system developers. Since we have no indication that any of the incidents in §4.1 employed or contracted external developers, this control in not applicable for these cases. Instead, control PL-8, analyzed in §4.2.4, applies to internal developers.

4.2.8 SC-2(2): Separation of System and User Functionality | Disassociability. SC-2 focuses primarily on separating user-level and system- or privileged-level functionality. Given the public nature of the data sets and inferences in the incidents, this base control would not have applied. The principles of SC-2(2) could still apply, though. For the one organization that appeared to have implemented this control enhancement (i.e., EdX), remaining quasi-identifiers (QIs) were problematic for preventing inferences.

For re-identification inferences, EdX's privacy policy [26] clearly communicated that it tracks users' interactions with the website; however, course enrollments and the number of forum posts per course—both QIs—were the only interaction data included in the data set [14]. In other words, we suspect that EdX did not normally separate its users' interaction data (so as to not hinder internal data analysis) but did so partially to release a de-identified data set. The NYC taxi data controllers did not implement this control. To arrive at this conclusion, we considered that drivers' system interaction state information to be the association of the specific driver with each route completed.

For operational inferences, Strava's data aggregation to generate the heat map removed all users' state information from the data, but it had no effect on revealing clandestine locations. Thus, SC-2(2) is not applicable to the Strava operational inference. If the DoD OPSEC Program alerted Pentagon employees to the risk, they could have omitted information from their pizza orders (i.e., delivery address) by taking-out or dining-in, especially if at a further location than the closest Domino's to the Pentagon. Following this practice, Domino's franchise owner would have had to incur a greater burden to associate the orders with the Pentagon.

4.2.9 SI-19: De-Identification. EdX, NYC, and Strava each partially implemented SI-19 but, as previously explained, de-identified data insufficiently to prevent inferences. On the other hand, the Domino's franchises' owner de-identified his customers but not their association with the USG, which, in isolation, is not necessarily a sensitive correlation. Arguably, though, exposing this operational inference was whole purpose of his interview with the Washington Post [69].

Although subcategory CT.DP-P3 maps to the SI-19 base control, framework adopters may choose to implement any of SI-19's eight control enhancements:

(1) Collection. De-identify the data a priori by limiting the collection only to the necessary fields. We analyze this concept of minimization under control SA-8(33) in §4.2.6.

(2) Archiving. To protect data stored long-term, this control enhancement urges de-identification of data before archiving them such that private data—whose intended utility was temporal—do not need to and will not be archived. Data archiving did not play a role in the presented incidents because the goal was data release.

(3) Release. De-identifying data before releasing it outside of the organization is the fundamental goal; however, the presented incidents demonstrated that it is a challenging endeavor. See also our discussion of disassociability and control IA-8(6) in §4.2.3.

(4) Removal, Masking, Encryption, Hashing, or Replacement of Direct Identifiers. In the EdX data set, fields that could serve as direct identifiers were either removed or replaced prior to its public release [14] but enough information remained to comprise QIs. The NYC TLC hashed direct identifiers, but did not employ a key or salt to the hash as the control recommends, which resulted in the re-identification of taxi drivers [77]. Strava removed direct identifiers through aggregation. We analyze the Pizza Index de-identification from the Domino's perspective with this base control and from the Pentagon perspective with control SC-2(2) in §4.2.8.

(5) Statistical Disclosure Control. This control enhancement applies to situations in which multiple versions of a data set enable analysts to infer specific attributes because of changes from one version to the next. For example, an analyst could capture images of the Strava heat map at various times to create a longitudinal data set and make inferences. For the other incidents, this control is not applicable because only one version of each data set was released.

(6) Differential Privacy. This is one of the most promising de-identification PETs to mitigate against re-identification inferences. We briefly present some of the foremost challenges to implementing differential privacy in §2.3.1 and analyze the potential success for differential privacy with control AC-23 in §4.2.1.

(7) Validated Algorithms and Software. There are many privacy preserving technologies available that help guard against inferences in specific use cases. For example, First, Hasan and Fritz [35] devise

a method that protects students from having their genders inferred from their online coursework interactions. Second, Zhang et al. [90] develop a method to anonymize data generated by human-wearable devices, inhibiting analysts' ability to infer age, gender, height, and weight while preserving the data well enough to maintain the ability to categorize activities, e.g., walking and running. Of the presented incidents, none employed a validated de-identification algorithm like either of these examples.

(8) Motivated Intruder. Analogous to a penetration test for an information system, which is covered by control CA-8 Penetration Testing, this control enhancement involves attempting to re-identify de-identified data. There is no indication that any party in the presented incidents applied this control.

5 RECOMMENDATIONS

Based on our analysis of NIST's Privacy Framework helping organizations mitigate privacy inferences, summarized in Table 2, we recommend that privacy programs incorporate the following inference-related components. We also propose how to incorporate these components into the framework [56]. To better mitigate an attacker's ability to combine anonymized data with other information, we seek to increase the coverage of inferences in mapped controls in §§5.1–5.2. The aim is to perceive risk and engage controls to offset it by increasing organizations' awareness of inference risks and ability to foresee potential exposures. First, we examine existing inference-related controls that are not mapped to subcategory CT.DP-P3, described in §3.2. Then, we recommend augmenting specific controls with inference-relevant verbiage and references and then mapping those controls to CT.DP-P3. In §5.3, we propose updating controls pertaining to PII to account for the obsolescence of solely safeguarding PII categories to protect privacy. Finally, we discuss challenges associated with translating legal requirements and policy into implementable technical solutions in §5.4. Although we recommend only small changes to the framework's controls and their mappings to subcategory CT.DP-P3, making inference risk much more visible will improve organizations' risk awareness and ability to mitigate problems.[2]

5.1 Inference-Related Controls not Mapped

In the course of our analysis, we found additional existing inference-related controls in the NIST SP 800-53 [58] that were not mapped to CT.DP-P3 [55]. Specifically, we found two controls containing words derived from "infer" and four other inference-related controls, none of them mapped to the inferences subcategory.

Of the two controls containing derivatives of "infer," we analyzed SI-19(5) above in §4.2.9. The other control—PL-4(1) "Rules of Behavior | Social Media and External Site/Application Usage Restrictions"—would have helped mitigate the Strava operational inference. It relates to social media users' interactions pertaining to the organization's information. Had deployed coalition forces considered their location to be organizational information—why else would they be in a foreign combat zone?—then they could not, by policy, have shared their tracked workouts without violating this control [37]. An OPSEC risk analysis could determine whether this information was okay to release. PL-4(1) is not mapped to any NIST Privacy Framework subcategory [55] whatsoever but should be mapped to CT.DP-P3.

Although numerous controls in the NIST SP 800-53 [58] can help mitigate violations of privacy in general, the following four additional controls would help organizations mitigate inferences.

- *AC-4(9) Information Flow Enforcement | Human Reviews.* Sometimes humans can predict and mitigate the potential for inferences, as was the U.S. Census Bureau's standard practice prior to adopting differential privacy [31]. In the NYC taxi incident, humans with the requisite expertise possibly could have foreseen the re-identification of taxi drivers and the ability to infer passengers' identities.
- *AC-21 Information Sharing.* AC-21(1) "Automated Decision Support" can help with AC-4(9) by employing PETs. AC-21(2), "Information Search and Retrieval," relates to the query interface suggested with control AU-16(3) in §4.2.2 and differential privacy, discussed in §2.3.1 and §4.2.1.
- *SC-38 Operations Security.* Nearly all organizations have sensitive information to conceal from others. Although both the Strava and Pizza Index incidents are related to military organizations, private organizations can also protect their sensitive information via OPSEC concepts. National Security Decision Directive (NSDD) 298 [67] provides an overview the OPSEC process and expounds on the importance of employee training to increase awareness of the organization's critical information and how to protect it against potential threats so as to better assure the organization's mission. Organizations may also direct their members to implement individual OPSEC controls, like not sharing fitness tracking data when at work.
- *SR-7 Supply Chain Operations Security.* SR-7 applies the principles of SC-38 to organizations' supply chains. Specifically, SR-7 advises…

  > …determining indicators that potential adversaries might obtain that could be interpreted or pieced together to derive information in sufficient time to cause harm to organizations …and considering how aggregated information may expose users or specific uses of the supply chain. [58]

  With Pentagon employees using Domino's as part of their [unofficial] supply chain, SR-7 might have helped mitigate this inference.

We recommend mapping each of the above four controls to NIST Privacy Framework subcategory CT.DP-P3.

5.2 Other Controls for Addressing Inferences

In addition to inference-related controls that were not mapped to CT.DP-P3, we found controls facially unrelated to inferences but that actually have an inference relevance. We recommend augmenting these controls—pertaining to organizational literacy of inferences, privacy violation disclosures, and organizations' privacy policies—and creating one new control to address inferences.

[2]We have discussed our findings and recommendations with the privacy engineering group at NIST, which is responsible for the framework.

To mitigate an inference-based compromise, one must first recognize and understand the threat. While there are many proprietary methods to develop threat models, Linking, Identifying, Non-repudiation, Detecting, Data Disclosure, Unawareness, Non-compliance (LINDDUN) specifically applies to privacy and, more recently, accounts for inferences [87]. To achieve this in the framework, we recommend bolstering control PM-16 "Threat Awareness Program" to incorporate inference-related threats. From the discussion section of PM-16's description, we recommend this modification in which we add the italicized portion of this quote:

> Because of the constantly changing and increasing sophistication of adversaries, especially the advanced persistent threat (APT), it may be more likely that adversaries can successfully breach or compromise organizational systems *or infer private organizational information.* [58]

Then, PM-16 will directly support mapping CT.DP-P3 to a new control enhancement to base control AT-2 "Literacy Training and Awareness," which currently has six control enhancements [58]. This seventh control enhancement, drafted in Figure 1, would be all about inferences, including how they can occur and have occurred and how to mitigate their potential to damage the organization.

(7) LITERACY TRAINING AND AWARENESS | INFERENCES

Provide literacy training on recognizing adversarial opportunities for conducting re-identification and operational inference of sensitive organizational information.

Discussion: Numerous inference-based incidents pepper recent history and have harmed organizations and individuals. Minimizing PII is insufficient for mitigating privacy inferences as motivated adversaries use all capabilities within their means to re-identify individuals in de-identified data and infer organizations' sensitive operational information.

Related Controls: AC-4, AC-22, AC-23, AC-24, PM-16, SC-38, SI-19

Figure 1: A New Control Enhancement for Control AT-2 in the NIST SP 800-53 [58]

Increasing organizations' literacy on inferences naturally leads to a more effective execution of other controls, including AC-1 "Policy and Procedures," PM-28 "Risk Framing," and RA-8 "Privacy Impact Assessments," to which multiple privacy framework subcategories map. In addition, these organizations are likely to develop and maintain more effective architectural controls, PL-8 and PM-7, to which multiple subcategories also map, including CT.DP-P3 [55]. Inference training is especially crucial for the personnel involved in executing controls AC-22 "Publicly Accessible Content" and AC-24 "Access Control Decisions" because of their role in publicizing data from which threat actors could make damaging inferences.

In addition to increasing awareness of inferences in general, organizations' disclosure procedures, which relate to incident response, for sensitive information also need modification. Therefore, we also recommend expanding the following disclosure-related controls with our italicized additions to incorporate inferences.

- *AU-13 Monitoring for Information Disclosure.* "Examples of organizational information include personally identifiable information retained by the organization or proprietary information generated by *or inferred about* the organization."
- *PM-21 Accounting of Disclosures.* "Develop and maintain an accurate accounting of disclosures of personally identifiable information *and incidents of re-identified individuals in purportedly anonymized organizational data*, including…"
- *RA-8 Privacy Impact Assessments.* "A privacy impact assessment is an analysis of how personally identifiable information *and sensitive organizational information ~~is~~ are* handled to ensure that handling conforms to applicable privacy requirements, determine the privacy risks associated with an information system or activity, and evaluate ways to mitigate privacy risks." [58]

Other subcategories of the NIST Privacy Framework map to these controls [55], which contribute to the organization-wide privacy preserving effort advocated by the framework itself [56].

Organizations should also incorporate inference concepts into their privacy program documentation. For example, an organization can increase its privacy transparency by publishing a privacy policy that clearly addresses the sources used in making inferences and communicates, as suggested by Wachter and Mittelstadt [83]:

- Its intent to infer sensitive data from non-sensitive data.
- Its intended purpose for the inference and how the source data are relevant, which should align ethically with their use, e.g., not inferring a predilection to gambling for targeted advertising of gambling opportunities.
- The assumed and proven accuracy and reliability of its inference methods.

Control PM-20(1) "Dissemination of Privacy Program Information | Privacy Policies on Websites, Applications, and Digital Services" addresses this well, stipulating that organizations' privacy policies are relatively easy to understand and enable informed consent [58]. "Notice and consent" is the prevalent legal standard [85] but opponents posit that transparency is actually more important. Transparency rather than compelling consent to privacy policies would lead to "inaccessible [privacy-exposing application programming interfaces (APIs)], coding to prevent scraping, automatic deletion of data, and blocking of cookies" [85].

5.3 PII Protection Requirements Are Obsolete

Over a decade ago, researchers established that protecting privacy by protecting only information in PII categories is ineffective [64, 71]. Cohen further demonstrates the insufficiency of k-anonymity and other QI-based anonymization techniques [15] on the EdX data set [14]. QI-based de-identification relies on "unstated assumptions on the data distribution" [14] and likewise, is insufficient to mitigate re-identification, even when employed in FERPA-compliant anonymization. The gap between legal requirements and technical implementations enables workarounds—both unintentional and intentional—to the de-identification process [15, 28, 60]. Incorporating the concept of inference generation from lingering QIs into controls like SI-19(6) "Differential Privacy" would prompt organizations to mitigate this inferences vulnerability.

Incorporates Inferences	A Type of Sensitive Information	PM-26	Safeguarding PII	PT-2(1)	Needs to Incorporate Inferences	IR-8(1)	PT-4(2)	SI-18(2)
AC-23	AC-3(11)	RA-5(5)	AC-2	PT-2(2)	AC-16	PL-2	PT-4(3)	SI-18(4)
PL-4(1)	AC-4(1)	SA-3(2)	AC-3(14)	PT-3	AT-2	PL-8	PT-5	SI-18(5)
PT-7(2)	AC-4(17)	SA-11(5)	AU-3(3)	PT-3(1)	AT-3	PM-5(1)	PT-5(1)	SI-19
SI-19(7)	AC-21	SI-7	AU-9	PT-3(2)	AT-3(3)	PM-9	RA-3	SI-19(2)
	AU-2	SI-11	AU-12(4)	PT-7	AT-3(5)	PM-11	RA-8	SI-19(3)
	AU-3	SI-12(3)	AU-14	PT-7(1)	CM-13	PM-20(1)	SA-8(33)	SI-19(6)
	AU-13	SI-20	PE-8(3)	SI-12(1)	IR-2(3)	PM-21	SA-15(12)	
	MP-6		PL-8(2)	SI-18	IR-4	PM-25	SC-7(24)	
	PM-23		PM-22	SI-18(1)	IR-8	PT-1	SC-16	
			PT-2	SI-18(3)		PT-4	SC-42(4)	
				SI-19(1)		PT-4(1)	SI-12(2)	

Figure 2: Controls Mentioning "Personally Identifiable Information" in the NIST SP 800-53 [58]

Protecting PII categories is still important; however, we found that multiple controls' dependence on the concept of using PII and/or QIs as sole safeguard(s) against privacy violations limited their effectiveness. As such, we recommend updating all controls in the NIST SP 800-53 [58] that depend on identifying PII categories in determining what data to safeguard. These updates are in the same spirit as Holder's recommendations to update the GDPR and CCPA language for inferences [40, pp. 1352–1353]. These would also improve framework subcategory CT.DP-P2, described in §3.2.

The NIST SP 800-53 lists 83 controls in that mention PII, listed here in Figure 2 [58]. Propitiously, we found 4 (5%) that already have wording sufficient to incorporate inferences. In addition, 17 (20%) controls mention PII simply as an example of sensitive information to organizations and associated individuals. We do not recommend any change to these controls to account for inferences because of their indication that PII categories are not the sole privacy concern. Similarly, we found that 22 (27%) controls are adequate as they are because they describe safeguarding PII—which is still important—without depending directly on identifying and protecting its categorized data as the sole method of privacy protection. Finally, 40 (48%) controls need modification to incorporate inferences. In these controls, appending "and other sensitive information that could lead to re-identification of individuals" to instances of "personally identifiable information" would be minimally sufficient.

5.4 Translating Legal Requirements into Technical Implementations

As a society, we generally value privacy; however, there are many concepts of privacy, including what it is and entails, and diverse opinions regarding sharing, trust, obfuscation, decisions to control one's privacy, and protection from those that might intrude into or exploit one's private matters [62, 85]. In an effort to protect people's privacy, legislatures make laws requiring organizations protect it. Legislatures often define principles and standards in laws rather than defining clear rules; whereas the former is often not computable in logic standards, computers apply rules very well [28, 61]. This distinction matters because organizations need to translate legal requirements into policies for the laws to have their intended effect. In turn, organizations must translate those policies into technical implementations. Strategic ambiguity and delegation of detail in laws can also be good, though, because of the afforded flexibility. Another issue with legal and policy requirements that can arise with NIST standards and guides is *checklist compliance*.

Many of the controls in the NIST SP 800-53 [58] mention the term "law" or "legal" but none explicitly describe translating law to the organizational policy or technical levels. Most uses of "law" refer to complying with "applicable laws" or referencing "law enforcement." Most instances of "legal" recommend seeking legal counsel for matters addressed by a control [58]. In consideration of these usages, we do not recommend any modification to the NIST Privacy Framework or controls to improve organizations' ability to translate law into policy and technical implementations.

Instead, legislatures and organizations can take actions to help overcome translation challenges, which could reduce litigation in the judiciary. Legislatures can work with scientists to establish a formal, mathematically provable definition of privacy or just implementable requirements [22, 60, 61]. People with the International Association of Privacy Professionals (IAPP)'s Certified Information Privacy Professional (CIPP) certification—which exists for the purpose of "putting privacy law and policy to work" [41]—may benefit both legislatures and organizations. Although Pasquale [65] envisions a future in which artificial intelligence (AI) would replace much of the tedious labor in the legal profession, we are not there yet. For the present, improvisational narrative and discourse to address cases not previously examined still need human involvement because there is no precedent in the training data [65]. As such, systems at all levels (legislative, policy, technical, and judicial) that implement legal logic standards must be adaptable to new laws or else risk obsolescence in a "technological–legal lock-in" that stifles legal evolution [16] and, to a lesser extent, industry.

Checklist compliance—completing a checklist for the sake of compliance and believing that the subject is now secure—can arise easily within organizations striving to meet legal or policy compliance requirements. In analyzing controls in §§4.2.2, 4.2.3, 4.2.5, 4.2.6, and 4.2.8, we described how applying the principle of a control is most effective in contrast to its verbatim implementation, which is like checklist compliance. Checklists have their place and, when developed and implemented well (a challenging endeavor), can be highly useful for guiding professionals through most any

complex process [32]. However, even NIST recognizes that compliance requirements lend themselves readily to assessments that focus on compliance rather than on "achieving a positive outcome for privacy" [8]. For example, the EdX data were treated per FERPA de-identification requirements but re-identification was still possible [14]. An appropriate checklist for the NIST Privacy Framework might include overarching goals, like communicate within the organization, identify risk, mitigate risk, and repeat continuously. Such a checklist could serve as a risk communication tool. Privacy impact assessments (PIAs) provide another checklist compliance opportunity as organizations strive to comply with applicable laws and government regulations; they help protect organizations by their mere existence as a paper trail but, by themselves, do not provide security or privacy [85]. In summary, we caution organizations to avoid checklist compliance for privacy.

6 CONCLUSION

As organizations face growing complexity in protecting privacy and risk in using sensitive data, resources to help organizations identify the nature and scope of their privacy risk also need to evolve. The practical need for complying with privacy laws and meeting people's expectations for organizations to protect their privacy led us to ask how we could improve the NIST Privacy Framework [56], a prominent guide to implementing organizational privacy programs. We focus on mitigating inference-based privacy violations, taxonomically defining inferences as re-identification or operational inferences. To determine how to improve organizations' defenses against privacy inferences and the NIST Privacy Framework, we apply the framework to past incidents of re-identification and operational inferences. This analysis revealed shortcomings in the framework's capacity to identify inference risk and recommend offsetting mitigations. Our recommendations include increasing organizations' awareness of inferences by expanding the mapping of inference-related controls and augmenting selected other controls to account for inferences, updating controls that depend on protecting specific PII categories or quasi-identifiers (QIs) as sufficient for protecting privacy, and improving the ability for organizations to translate legal requirements into policy and policy into technical implementations. Further analyses of NIST Privacy Framework effectiveness would contribute to this field of research, especially if conducted by framework-implementing organizations on privacy incidents or near-incidents involving inferences.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive feedback. We appreciate the time and feedback that the privacy engineering group at NIST gave to us, especially Naomi Lefkovitz and Dylan Gilbert. We also specifically thank Nitin Kohli and Chris Hoofnagle for their feedback. Although the authors conducted this research as an official function of their U.S. Navy employment, this research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Views expressed are solely those of the authors and do not represent official positions of the U.S. Navy, Department of Defense, or U.S. Government.

REFERENCES

[1] Almudena Alcaide, Esther Palomar, Israel Barroso-Perez, and Ali E. Abdallah. 2011. Privacy-Enhancing Cryptography-Based Materials. In *Proceedings of the International Conference on Security and Cryptography*. 379–382. https://ieeexplore.ieee.org/document/6732417

[2] Kenneth A Bamberger and Deirdre K Mulligan. 2015. *Privacy on the Ground: Driving Corporate Behavior in the United States and Europe.* MIT Press, Cambridge, MA.

[3] Michael Barbaro and Tom Zeller. 2006. A Face Is Exposed for AOL Searcher No. 4417749. The New York Times. https://www.nytimes.com/2006/08/09/technology/09aol.html

[4] Steven M. Bellovin, Renée M. Hutchins, Tony Jebara, and Sebastian Zimmeck. 2014. When Enough is Enough: Location Tracking, Mosaic Theory, and Machine Learning. *New York University Journal of Law & Liberty* 8 (2014), 555–628. https://digitalcommons.law.umaryland.edu/fac_pubs/1375/

[5] Rob Bonta and Susan Duncan Lee. 2022. Office of the Attorney General, State of California, Opinion. No. 20-303. https://oag.ca.gov/system/files/opinions/pdfs/20-303.pdf

[6] Luís T. A. N. Brandão and René Peralta. 2021. *Privacy-Enhancing Cryptography to Complement Differential Privacy.* Technical Report. PEC Project, Cryptographic Technology Group, NIST. https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=932789

[7] Malte Brettel, Jens-Christian Reich, Jose M. Gavilanes, and Tessa C. Flatten. 2015. What Drives Advertising Success on Facebook? An Advertising-Effectiveness Model Measuring the Effects on Sales Of "Likes" and Other Social-Network Stimuli. *Journal of Advertising Research* 55, 2 (2015), 162–175. https://doi.org/10.2501/JAR-55-2-162-175

[8] Sean Brooks, Michael Garcia, Naomi Lefkovitz, Suzanne Lightman, and Ellen Nadeau. 2017. *An Introduction to Privacy Engineering and Risk Management in Federal Systems.* National Institute of Standards and Technology, U.S. Department of Commerce, Gaithersburg, MD. https://doi.org/10.6028/NIST.IR.8062

[9] Marc B. Bucks and Daphne Williams. 2021. *Maintaining a Competitive Advantage: An Analysis of the Efficiency and Effectiveness of the Marine Corps' Assessment and Authorization Process.* Master's thesis. Naval Postgraduate School, Monterey, CA. https://hdl.handle.net/10945/68302

[10] Jen Caltrider, Misha Rykov, and Zoë MacDonald. 2023. It's Official: Cars Are the Worst Product Category We Have Ever Reviewed for Privacy. Mozilla. https://foundation.mozilla.org/en/privacynotincluded/articles/its-official-cars-are-the-worst-product-category-we-have-ever-reviewed-for-privacy/

[11] Chief Information Officer, U.S. Department of Defense. 2010. DoDAF Architecture Framework Version 2.02. https://dodcio.defense.gov/Library/DoD-Architecture-Framework/

[12] City of New York. 2005. Disclosure of Security Breach. Administrative Code, Title 10, Chapter 5. https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=441345&GUID=8ECF99C5-52FA-48A2-B776-D5817E8D463A&Options=ID|Text|

[13] City of New York. 2016. Computer Systems Manager. Department of Information Technology and Telecommunications, Job Posting Notice. https://www1.nyc.gov/assets/doitt/downloads/jobs/276549_dir_infrastructure_arch.pdf

[14] Aloni Cohen. 2022. Attacks on Deidentification's Defenses. https://doi.org/10.48550/arxiv.2202.13470 Unpublished.

[15] Aloni Cohen and Kobbi Nissim. 2020. Towards Formalizing the GDPR's Notion of Singling Out. *Proceedings of the National Academy of Sciences* 117, 15 (2020), 8344–8352. https://doi.org/10.1073/pnas.1914598117

[16] Rebecca Crootof. 2019. "Cyborg Justice" and the Risk of Technological–Legal Lock-In. *Columbia Law Review* 119, 7 (2019), 233–251. https://www.jstor.org/stable/26960742

[17] Steven H. Dam. 2006. *DoD Architecture Framework: A Guide to Applying System Engineering to Develop Integrated, Executable Architectures.* SPEC, Marshall, VA.

[18] Yves-Alexandre de Montjoye, Cesar A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. 2013. Unique in the Crowd: The Privacy Bounds of Human Mobility. *Scientific Reports* 3, 1 (2013), 1376–1376. https://doi.org/10.1038/srep01376

[19] Damien Desfontaines and Balázs Pejó. 2020. SoK: Differential Privacies. In *Proceedings on Privacy Enhancing Technologies*. 288–313. https://doi.org/10.2478/popets-2020-0028

[20] Marie Douriez, Harish Doraiswamy, Juliana Freire, and Cláudio T. Silva. 2016. Anonymizing NYC Taxi Data: Does It Matter?. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA).* 140–148. https://doi.org/10.1109/DSAA.2016.21

[21] Charles Duhigg. 2012. How Companies Learn Your Secrets. The New York Times Magazine. https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?pagewanted=1&_r=1&hp

[22] Cynthia Dwork, Nitin Kohli, and Deirdre Mulligan. 2019. Differential Privacy in Practice: Expose your Epsilons! *Journal of Privacy and Confidentiality* 9, 2 (2019), 1–22. https://doi.org/10.29012/jpc.689

[23] Cynthia Dwork and Deirdre K. Mulligan. 2013. It's Not Privacy, and It's Not Fair. *Stanford Law Review Online* 66, 35 (September 2013),

35–40. https://stanfordlawreview.org/wp-content/uploads/sites/3/2016/08/DworkMullliganSLR.pdf
[24] Cynthia Dwork and Aaron Roth. 2014. *The Algorithmic Foundations of Differential Privacy.* Now Publishers, Inc., Boston, MA. https://doi.org/10.1561/0400000042
[25] Lilian Edwards and Michael Veale. 2017. Slave to the Algorithm? Why a 'Right to an Explanation' is Probably Not the Remedy You Are Looking For. *Duke Law and Technology Review* 16, 1 (2017), 18–84. https://discovery.ucl.ac.uk/id/eprint/1574817/1/Veale_slavetothealgorithm_published.pdf
[26] EdX. 2012. Privacy Policy. Wayback Machine. https://web.archive.org/web/20120920014400/https://www.edx.org/privacy Accessed May 16, 2022.
[27] Federal Trade Commission. 2019. FTC Sues Cambridge Analytica, Settles with Former CEO and App Developer. https://www.ftc.gov/news-events/news/press-releases/2019/07/ftc-sues-cambridge-analytica-settles-former-ceo-app-developer
[28] Joan Feigenbaum and Daniel J. Weitzner. 2018. On the Incommensurability of Laws and Technical Mechanisms: Or, What Cryptography Can't Do. In *Security Protocols XXVI.* 266–279. https://doi.org/10.1007/978-3-030-03251-7_31
[29] Md Sadek Ferdous, Soumyadeb Chowdhury, and Joemon M. Jose. 2016. Privacy Threat Model in Lifelogging. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct.* 576–581. https://doi.org/10.1145/2968219.2968324
[30] Geoffrey A. Fowler. 2019. What Does Your Car Know about You? We Hacked a Chevy to Find Out. The Washington Post. https://www.washingtonpost.com/technology/2019/12/17/what-does-your-car-know-about-you-we-hacked-chevy-find-out/
[31] Simson L. Garfinkel, John M. Abowd, and Sarah Powazek. 2018. Issues Encountered Deploying Differential Privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society.* 133–137. https://doi.org/10.1145/3267323.3268949
[32] Atul Gawande. 2010. *The Checklist Manifesto: How to Get Things Right.* Metropolitan Books, New York.
[33] Ian Goldberg. 2007. Privacy-Enhancing Technologies for the Internet III: Ten Years Later. In *Digital Privacy: Theory, Technologies, and Practices,* Alessandro Acquisti, Stefanos Gritzalis, Costos Lambrinoudakis, and Sabrina di Vimercati (Eds.). Auerbach Publications, New York, Chapter 1, 3–18. https://doi.org/10.1201/9781420052183-7
[34] Stephen Groat, Joseph Tront, and Randy Marchany. 2012. Advancing the defense in depth model. In *2012 7th International Conference on System of Systems Engineering (SoSE).* 285–290. https://doi.org/10.1109/SYSoSE.2012.6384127
[35] Rakibul Hasan and Mario Fritz. 2022. Understanding Utility and Privacy of Demographic Data in Education Technology by Causal Analysis and Adversarial-Censoring. *Proceedings on Privacy Enhancing Technologies* 2022, 2 (2022), 245–262. https://doi.org/10.2478/popets-2022-0044
[36] Alex Hern. 2018. Fitness Tracking App Strava Gives Away Location of Secret US Army Bases. The Guardian. https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases
[37] Alex Hern. 2018. Strava Suggests Military Users 'Opt Out' of Heatmap as Row Deepens. The Guardian. https://www.theguardian.com/technology/2018/jan/29/strava-secret-army-base-locations-heatmap-public-users-military-ban
[38] Kashmir Hill. 2012. How Target Figured Out a Teen Girl Was Pregnant Before Her Father Did. Forbes. https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did
[39] Christian Pieter Hoffmann, Christoph Lutz, and Giulia Ranzini. 2016. Privacy Cynicism: A New Approach to the Privacy Paradox. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace* 10, 4 (December 6, 2016), article 7. https://doi.org/10.2139/ssrn.3319830
[40] Allan E. Holder. 2020. What We Don't Know They Know: What To Do About Inferences in European and California Data Protection Law. *Berkeley Technology Law Journal* 35, 4 (2020), 1331–1364. https://doi.org/10.15779/Z38MP4VP1V
[41] International Association of Privacy Professionals. 2022. CIPP. https://iapp.org/certify/cipp/ Accessed May 2, 2022.
[42] Suzanne Kaufman. 2022. The Invisible, Yet Omnipresent Ear: The Insufficiencies of the Children's Online Privacy Protection Act. *NYU Annual Survey of American Law* 78 (2022), 101–136. https://annualsurveyofamericanlaw.org/wp-content/uploads/2022/11/Kaufman-78.1.pdf
[43] Joshua A. Kroll, Nitin Kohli, and Paul Laskowski. 2019. Privacy and Policy in Polystores: A Data Management Research Agenda. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*, Vijay Gadepally, Timothy Mattson, Michael Stonebraker, Fusheng Wang, Gang Luo, Yanhui Laing, and Alevtina Dubovitskaya (Eds.). LNCS, Vol. 11721. Springer, Cham, 68–81. https://doi.org/10.1007/978-3-030-33752-0_5
[44] Issie Lapowsky. 2019. How Cambridge Analytica Sparked the Great Privacy Awakening. Wired. https://www.wired.com/story/cambridge-analytica-facebook-privacy-awakening/
[45] Nancy G. Leveson. 2016. *Engineering a Safer World.* The MIT Press, Cambridge, MA. https://doi.org/10.7551/mitpress/8179.001.0001
[46] Kirsten Martin and Helen Nissenbaum. 2017. Privacy Interests in Public Records: An Empirical Investigation. *Harvard Journal of Law & Technology* 31, 1 (2017), 111–143. http://jolt.law.harvard.edu/articles/pdf/v31/31HarvJLTech111.pdf
[47] Patrick McGee. 2021. Google Advertising Boom Lifts Profits to Record. https://www.proquest.com/docview/2532469468
[48] Deirdre K. Mulligan, Colin Koopman, and Nick Doty. 2016. Privacy is an Essentially Contested Concept: A Multi-Dimensional Analytic for Mapping Privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374, 2083 (2016), 20160118. https://doi.org/10.1098/rsta.2016.0118
[49] Soichi Nakajima. 2019. Digital Advertising: The Google/Facebook Duopoly Continues. In *Digiworld Yearbook.* Institut de l'Audiovisuel et de Telecommunications en Europe (IDATE), Montpellier, 80–81. https://www.proquest.com/docview/2445995003/
[50] Arvind Narayanan and Edward W. Felten. 2014. No Silver Bullet: De-identification Still Doesn't Work. (July 9, 2014). https://www.cs.princeton.edu/~arvindn/publications/no-silver-bullet-de-identification.pdf Unpublished.
[51] Arvind Narayanan, Joanna Huey, and Edward W. Felten. 2016. A Precautionary Approach to Big Data Privacy. In *Data Protection on the Move: Current Developments in ICT and Privacy/Data Protection*, Serge Gutwirth, Ronald Leenes, and Paul De Hert (Eds.). Springer, 357–385. https://doi.org/10.1007/978-94-017-7376-8_13
[52] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust De-anonymization of Large Sparse Datasets. In *2008 IEEE Symposium on Security and Privacy (S&P).* 111–125. https://doi.org/10.1109/SP.2008.33
[53] National Institute of Standards and Technology. 2018. *Framework for Improving Critical Infrastructure Cybersecurity, version 1.1.* U.S. Department of Commerce, Gaithersburg, MD. https://doi.org/10.6028/NIST.CSWP.04162018
[54] National Institute of Standards and Technology. 2020. *Crosswalk from the NIST Privacy Framework Core to the Framework for Improving Critical Infrastructure Cybersecurity V1.1.* U.S. Department of Commerce, Gaithersburg, MD. https://www.nist.gov/privacy-framework/resource-repository/browse/crosswalks/cybersecurity-framework-crosswalk
[55] National Institute of Standards and Technology. 2020. *Cybersecurity Framework/Privacy Framework to NIST Special Publication 800-53, Revision 5 Mapping.* U.S. Department of Commerce, Gaithersburg, MD. https://www.nist.gov/privacy-framework/nist-privacy-framework-and-cybersecurity-framework-nist-special-publication-800-53
[56] National Institute of Standards and Technology. 2020. *NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management, version 1.0.* U.S. Department of Commerce, Gaithersburg, MD. https://doi.org/10.6028/NIST.CSWP.01162020
[57] National Institute of Standards and Technology. 2022. Privacy Framework Perspectives and Success Stories. https://www.nist.gov/privacy-framework/getting-started-0/perspectives-and-success-stories
[58] National Institute of Standards and Technology Joint Task Force. 2020. *Security and Privacy Controls for Information Systems and Organizations.* U.S. Department of Commerce, Gaithersburg, MD. https://doi.org/10.6028/NIST.SP.800-53r5
[59] Helen Nissenbaum. 2019. Contextual Integrity Up and Down the Data Food Chain. *Theoretical Inquiries in Law* 20, 1 (2019), 221–256. https://doi.org/10.1515/til-2019-0008
[60] Kobbi Nissim, Aaron Bembenek, Alexandra Wood, Mark Bun, Marco Gaboardi, Urs Gasser, David R O'Brien, Thomas Steinke, and Salil Vadhan. 2018. Bridging the Gap between Computer Science and Legal Approaches to Privacy. *Harvard Journal of Law & Technology* 31, 2 (2018), 687–780. https://jolt.law.harvard.edu/assets/articlePDFs/v31/02.-Article-Wood-7.21.pdf
[61] Kobbi Nissim and Alexandra Wood. 2018. Is Privacy *Privacy*? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, 2128 (2018), 20170358. https://doi.org/10.1098/rsta.2017.0358
[62] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. 2018. Turtles, Locks, and Bathrooms: Understanding Mental Models of Privacy Through Illustration. *Proceedings on Privacy Enhancing Technologies* 2018, 4 (2018), 5–32. https://doi.org/10.1515/popets-2018-0029
[63] Office of the Director, Administration and Management, DoD 2007. *Department of Defense Privacy Program.* Office of the Director, Administration and Management, DoD, Washington, DC. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodm/540011r.pdf
[64] Paul Ohm. 2010. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. *UCLA Law Review* 57, 6 (2010), 1701–1777. https://www.uclalawreview.org/pdf/57-6-3.pdf
[65] Frank Pasquale. 2019. A Rule of Persons, Not Machines: The Limits of Legal Automation. *George Washington Law Review* 87, 1 (January 2019), 1–55. http://www.gwlr.org/wp-content/uploads/2019/01/87-Geo.-Wash.-L.-Rev.-1.pdf
[66] Jovan Powar and Alastair R. Beresford. 2023. SoK: Managing Risks of Linkage Attacks on Data Privacy. *Proceedings on Privacy Enhancing Technologies* 2023, 2 (2023), 97–116. https://doi.org/10.56553/popets-2023-0043
[67] Ronald Reagan. 1988. *National Operations Security Program.* The White House. https://catalog.archives.gov/id/6879871
[68] Jerome H. Saltzer and Michael D. Schroeder. 1975. The Protection of Information in Computer Systems. *Proc. IEEE* 63, 9 (1975), 1278–1308. https://doi.org/10.1109/PROC.1975.9939

[69] Sarah Schafer. 1998. With Capital in Panic, Pizza Deliveries Soar. *Washington Post* (December 19, 1998), D1. https://www.washingtonpost.com/wp-srv/politics/special/clinton/stories/pizza121998.htm

[70] Seth Schoen. 2009. What Information is "Personally Identifiable"? Electronic Frontier Foundation (EFF). https://www.eff.org/deeplinks/2009/09/what-information-personally-identifiable

[71] Paul M. Schwartz and Daniel J. Solove. 2011. The PII Problem: Privacy and a New Concept of Personally Identifiable Information. *New York University Law Review (1950)* 86, 6 (2011), 1814–1894. https://www.law.berkeley.edu/files/bclt_Schwartz-Solove_NYU_Final_Print.pdf

[72] Elaine Shi, Waqar Aqeel, Balakrishnan Chandrasekaran, and Bruce Maggs. 2021. Puncturable Pseudorandom Sets and Private Information Retrieval with Near-Optimal Online Bandwidth and Time. In *Advances in Cryptology – CRYPTO 2021*. 641–669. https://doi.org/10.1007/978-3-030-84259-8_22

[73] Ido Sivan-Sevilla, Wenyi Chu, Xiaoyu Liang, and Helen Nissenbaum. 2020. Unaccounted Privacy Violation. In *FTC PrivacyCon 2020*. 1–25. https://www.ftc.gov/system/files/documents/public_events/1548288/privacycon-2020-ido_sivan-sevilla.pdf

[74] Greg Slabodkin. 2013. Defending DOD Networks with a Single Security Architecture. Defense Systems. https://defensesystems.com/it-infrastructure/2013/07/defending-dod-networks-with-a-single-security-architecture/190685/

[75] Daniel J. Solove and Danielle Keats Citron. 2018. Risk and Anxiety: A Theory of Data-Breach Harms. *Texas Law Review* 96 (2018), 737–786. https://doi.org/10.2139/ssrn.2885638

[76] Mikael Thalen. 2021. Hacker Reveals Smart Meters are Spilling Secrets about the Texas Snowstorm. Daily Dot. https://www.dailydot.com/debug/hacker-smart-meter-texas-snowstorm/

[77] Anthony Tockar. 2014. Riding with the Stars: Passenger Privacy in the NYC Taxicab Dataset. Wordpress. https://agkn.wordpress.com/2014/09/15/riding-with-the-stars-passenger-privacy-in-the-nyc-taxicab-dataset/

[78] Under Secretary of Defense for Intelligence, DoD 2020. *DoD Operations Security (OPSEC) Program Manual*. Under Secretary of Defense for Intelligence, DoD, Washington, DC. https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodm/520502m.PDF

[79] United States Congress 1996. *Health Insurance Portability and Accountability Act of 1996*. United States Congress, Washington, DC. https://www.govinfo.gov/link/plaw/104/public/191?link-type=pdf&.pdf

[80] United States Congress 1998. *Children's Online Privacy Protection Act of 1998*. United States Congress, Washington, DC. https://uscode.house.gov/view.xhtml?path=/prelim@title15/chapter91&edition=prelim

[81] United States Congress 2013. *Family Educational Rights and Privacy Act of 1974*. United States Congress, Washington, DC. https://www.law.cornell.edu/uscode/text/20/1232g

[82] U.S. Department of Health and Human Services. 2017. *Health Information Privacy Beyond HIPAA: A 2018 Environmental Scan of Major Trends and Challenges*. Technical Report. National Committee on Vital and Health Statistics (NCVHS) and its Privacy, Security, and Confidentiality Subcommittee. https://www.ncvhs.hhs.gov/wp-content/uploads/2018/02/NCVHS-Beyond-HIPAA_Report-Final-02-08-18.pdf

[83] Sandra Wachter and Brent Mittelstadt. 2019. A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI. *Columbia Business Law Review* 2019, 2 (2019), 1–130. https://ssrn.com/abstract=3248829

[84] Isabel Wagner and David Eckhoff. 2018. Technical Privacy Metrics: A Systematic Survey. *ACM Computing Surveys (CSUR)* 51, 3 (June 2018), 1–38. https://doi.org/10.1145/3168389

[85] Ari Ezra Waldman. 2021. *Industry Unbound: The Inside Story of Privacy, Data, and Corporate Power*. Cambridge University Press, Cambridge. https://doi.org/10.1017/9781108591386

[86] J.A. Whittaker. 2003. No Clear Answers on Monoculture Issues. *IEEE Security Privacy* 1, 6 (2003), 18–19. https://doi.org/10.1109/MSECP.2003.1266365

[87] Kim Wuyts, Laurens Sion, and Wouter Joosen. 2020. LINDDUN GO: A Lightweight Approach to Privacy Threat Modeling. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. 302–309. https://doi.org/10.1109/EuroSPW51379.2020.00047

[88] Xun Yi, Russell Paulet, and Elisa Bertino. 2014. Homomorphic Encryption. In *Homomorphic Encryption and Applications*. Springer, Cham, 27–46. https://doi.org/10.1007/978-3-319-12229-8_2

[89] Elise Young. 2015. Educational Privacy in the Online Classroom: FERPA, MOOCs, and the Big Data Conundrum. *Harvard Journal of Law & Technology* 28, 2 (2015), 549–592. https://jolt.law.harvard.edu/articles/pdf/v28/28HarvJLTech549.pdf

[90] Dalin Zhang, Lina Yao, Kaixuan Chen, Guodong Long, and Sen Wang. 2019. Collective Protection: Preventing Sensitive Inferences via Integrative Transformation. In *2019 IEEE International Conference on Data Mining (ICDM)*. 1498–1503. https://doi.org/10.1109/ICDM.2019.00197