

A Risk Assessment Framework for Digital Identification Systems

Allison Woodruff
Google
woodruff@acm.org

Will Drewry
Google
drewry@google.com

Dirk Balfanz
Google
balfanz@google.com

Mariana Raykova
Google
marianar@google.com

Abstract

We introduce a risk assessment framework for digital identification systems, as well as recommended best practices to enhance privacy, security, and other desirable properties in these systems. To generate these resources, we created a casebook of a wide range of digital identification systems, and we then applied expert analysis and critique to identify patterns. We piloted the framework on several reviews within our organization over a period of approximately one year, and found it to be robust and helpful for those reviews. This work is intended to inform product review and development, product policy, and standards efforts, and to help guide a consistent responsible approach to digital identification across the broader digital identification ecosystem.

Keywords

digital identification, privacy, risk assessment, security

1 Introduction

In recent years, many digital identification systems have been launched or proposed. For example, the press has reported on refugees scanning their irises to authenticate themselves to receive food and cash in certain refugee camps [44, 46, 76]. The company deploying the technology has promoted it as a secure way for refugees to efficiently prove their identities, e.g., to quickly check out of a grocery store at a refugee camp [46]. Proponents have also pointed out that iris scans are an effective way to minimize fraud from people who “dual register” to get extra benefits [46, 76]. At the same time, refugee advocates have raised substantial privacy and human rights concerns such as creating a database of stable identifiers of vulnerable populations [44, 46, 76]. Beyond distribution of goods and services [66, 83], many other use cases for digital identification systems have emerged as well, e.g., mobile driver’s licenses [16], transportation cards [41, 92], health certificates [2, 7, 52], human verification [8, 40], social media verification [62], driver authentication [39, 85, 86], identification of homeless individuals [19], and more.

Such digital identification systems offer many benefits such as convenience and authentication. Further, digital identification offers substantial advantages over physical identity cards, for example,

allowing people to present the minimum information necessary rather than all information on their government identity card, e.g., digital identification may allow a user to show their age without revealing additional personal information such as their address. At the same time, digital identification may create substantial risks such as impersonation, fraud, surveillance, increased government access to data, or broad data breaches.

Demand for digital versions of government identification, health certificates, and other documentation is growing rapidly [66], accompanied by public policy calls for responsible development of such technology [27, 28]. However, many aspects of digital identification are as yet unresolved, and best practices for digital identification have not yet been fully established. Even when best practices have been articulated, they have often not been achieved in practice, for example, because of lack of awareness or because improvements rely on cooperation across a large ecosystem of issuers,¹ verifiers,² and technology providers.

We introduce a framework of risk analysis and best practices for digital identification, based on expert experience and case analysis. The work is intended to inform product review and development, product policy, and standards efforts, and to help guide a consistent responsible approach to digital identification across the wider digital identification ecosystem.

Digital identification is a broad space. Our exploration centers on issues such as provisioning, managing, and presenting ID/credentials (e.g., driver’s licenses, health cards) issued by governments/authorities or by non-government entities—this may also include, for example, mechanisms to verify or impute government ID (e.g., the use of facial recognition to infer a person’s legal identity), or the use of non-government identification systems to provide government/social services. While authentication is not our main focus, we have considered it in our analysis and provide select illustrative examples related to authentication. Issues such as the following are beyond the scope of our exploration: labeling identity attributes (e.g., assigning gender [69]), or personal presentation [33] in digital spaces (e.g., social identity management).

Further, in this work, we focus largely on identifying and reducing potential risks to users and society, many of which revolve around privacy and security. Even within this scope, we do not consider this work comprehensive, and beyond that we recognize that other considerations must also be taken into account when assessing digital identification systems. Accordingly, we view the

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies 2026(2), 242–258

© 2026 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2026-0046>



¹We define an **issuer** as an organization that generates a digital card and provides it to an individual, often through an app a user has installed on their phone or by email. For example, a healthcare provider or pharmacy may issue a vaccine card.

²We define a **verifier** as any organization that reads the information in the credential. For example, a concert venue may check health certificates as a condition of entry.

framework to be just one tool among many to consider the safer deployment of technology.

The remainder of the paper is structured as follows. We begin by reviewing relevant background. Next we describe our methodology and introduce the framework. We then provide additional discussion and conclude.

2 Background

In this section, we review relevant digital identification systems, products, and standards, and we then turn to risk and harms analysis frameworks.

2.1 Digital Identification Systems

Digital identification is a broad and still developing space. While digital identities have existed for decades [88], and efforts such as the W3C’s work on decentralized identifiers (DIDs) [73] have advanced standardization, much work remains to increase awareness, availability, and consistent application of best practices.

2.1.1 Legal Digital Identity. The digitization of legal identity is at various levels of deployment in countries globally. Electronic passports may be considered the first widely deployed digital identity, but their use is primarily for international travel (see [77] for a discussion of digital support for “use cases that normally require a physical passport,” and [48, 70] for discussion of considerations and concerns with the use of biometrics in border control and e-passport applications). In the European Union, extensive standardization work has been underway for some time to ratify eIDAS (electronic Identification, Authentication, and Trust Services) 2.0 [26]. In the US, the deployment of mobile driver’s licenses (mDL) has been tested in several jurisdictions [16].

Legal digital identity deployments often involve many parties: government agencies, government-approved technology vendors, and consumer technology vendors. For example, in the US, government digital identification is the responsibility of each state’s department of motor vehicles (DMV) [42] which need to meet emerging federal guidance [78]. These government entities then outsource the technical integration work to approved vendors which then integrate with consumer technology providers through standardized or non-standardized means. Each of these steps provides opportunities for risk assessment, consideration of best practices, and application of standards. However, given the many layers of integration, even identity systems which have been undergoing standardization for years, such as mDL, will have gaps in standards coverage—such as a lack of standards specifying provisioning protocols.

2.1.2 Consumer Digital Identity. Consumer digital identity is simply the mapping of a person to their associated data with an online software service. In the past, usernames and passwords, which acted as a public identifier and a secret knowledge factor, were the ways that consumers managed that association. Many consumer devices now support both secure cryptographic key management and local measurement of biometrics, which enable newer approaches, such as FIDO’s Passkeys [29].

Consumer digital identities are usually purpose-driven, from accessing email services to providing loyalty discounts at a chain of stores. While there are many different standards, there is no

obligation to rely on any given standard. Additionally, the purpose-oriented nature may change the requirements that a service provider holds a consumer to, such as providing a valid credit card or photo of a legal identity document.

2.1.3 Physical Identity. Underpinning both legal and consumer identity are the associated person’s actual physical identity. Historically, biometrics were collected mainly in support of legal identity, binding paperwork to a physical person. With facial recognition technologies, it has become increasingly common to map from biometrics to legal digital identity. Further, with technological advancements in mobile devices, biometrics now underlay not only legal identity but consumer digital identities as well.

2.2 Risk and Harms Analysis Frameworks

Disciplines such as security and privacy have a valuable tradition of sharing guidelines that outline best practices, identify risk, and minimize harms [10, 21, 35, 38, 50]. Such frameworks allow experienced practitioners and researchers to synthesize and share their insights, making them more accessible to developers, regulators, advocates, and others. In this way, risk and harms frameworks can be used to evaluate and improve technical systems. For example, they may be used during formal processes such as launch reviews or adversarial testing, or more informally during product discussions or design sprints. These frameworks can also drive standards and policy agendas. Such frameworks are particularly valuable for emerging technologies and topics, where complex real-world practical challenges may not yet have been widely articulated or discussed.

Frameworks emphasizing risk and harms have become more prominent, particularly with the emergence of responsible innovation and responsible AI approaches [74, 87]. Shelby et al. developed a taxonomy of sociotechnical harms of algorithmic systems, based on a review of research on this topic [71]. Relatedly, Scheuerman et al. employed a grounded theory approach to develop a framework of harmful content online [68]. Neither of these frameworks are specific to digital identification, and we view them as complementary to our approach in that they articulate a large range of potential harms which may be considered as part of risk assessment during the application of our framework. Warford et al. introduce a framework that synthesizes research on risk factors for at-risk users [84]. While this work is not specific to digital identification, it describes real-world circumstances and practices of at-risk users, and is therefore complementary to our approach in that it may be used to help assess the impact of certain implementation decisions on at-risk users.

NIST 800-63 presents Digital Identity Guidelines which cover interaction with government information systems, encompassing identity proofing, authentication, and federated identity [78]. By contrast, our work gives broad consideration to a wide range of developers and products, including government and non-government actors and systems. The NIST Digital Identity Guidelines articulate a general impact assessment process consistent with the NIST Guide for Conducting Risk Assessments [64], comprised of listing groups that may be impacted, listing potential harms, assessing severity, and determining the appropriate assurance level to be

followed [78]. We view our work as complementary in that we articulate specific risks, common pitfalls, and concrete best practices in the development and implementation of digital identification systems, providing a short, practical, targeted list of considerations and capturing issues not highlighted in other frameworks. This complementary approach is supported by the NIST Guide for Conducting Risk Assessments, which encourages the simultaneous application of multiple risk assessment frameworks and advocates for tailored approaches to risk management [64]. The NIST Guide for Conducting Risk Assessments further calls for organizations to share their risk-related learning and expertise whenever possible [64], which is our goal in sharing this framework.

3 Methodology

To frame our exploration of the space, we constructed a casebook of 37 digital identification products representing a broad range of issues. We constructed the casebook from October 2021 through November 2022, and reviewed and analyzed it from June 2022 through January 2023; we overlapped the casebook construction, review, and analysis so that we could iteratively select informative cases and also incorporate newly emerging issues. We chose the casebook approach because digital identity issues are highly contextual, so specific cases elucidate important risks and mitigations. To construct the casebook, we consulted internal stakeholders and reviewed external literature to identify a wide variety of real and fictional systems/products (e.g., [15, 18, 19, 22, 39, 43, 46, 47, 63, 80, 85]) and then wrote a précis of each. We completed case collection when we reached saturation of risk factors [32, 67]. For more information about the cases, see Appendix A, Table 6.

The cases were reviewed and analyzed by four co-authors who are experts in digital identification systems; amongst these reviewers there is also expertise in security, privacy, AI ethics, health ethics, human-computer interaction, human rights, and related issues. We reviewed and independently scored each case according to (1) risk and (2) whether a technology company should build the product. We also optionally wrote open-ended comments on our analysis of each product.

Our evaluation strategy was within the context of a global human rights perspective, implemented by cross-disciplinary experts working in a multinational context. In that context, cultural or regulatory norms may be acceptable to product creators, such as sharing user-identifiable biometric embedding data with a centralized government database, but still introduce vectors of human rights abuse by the product creator or others who access the data through legitimate or illegitimate means. Further, security and privacy mechanisms often generate broader concerns, necessitating holistic review. Accordingly, we considered both first- and second-order effects, e.g., the construction and potential abuse of centralized databases, as well as foundational techniques that can minimize the risk of harmful second-order effects. Our primary focus was on risks and harms that may impact identity holders and society; while we considered holistic perspectives and practical constraints for all actors in the ecosystem, we primarily rely on operational mechanisms beyond our framework to advocate for the needs of issuers and verifiers.

We then held a series of reconciliation meetings, supplemented by occasional written communication, to understand and/or resolve discrepancies. Most discrepancies were due to different priorities and/or missed risks; in the reconciliation discussions a reviewer often adjusted an outlier score to a higher risk level, which brought scores closer together. Discrepancies were typically not due to different assumptions about the facts of the cases. We took detailed notes throughout the reconciliation, for example, noting best practices that in our experience are often overlooked.

We used an inductive approach to explore emerging themes and common patterns in the cases [79]. Drawing on reflexive thematic analysis approaches [11, 12], we iteratively developed and refined a list of risk factors that were most salient for each of the digital identification cases. In some instances these risk factors stemmed from our own prior experiences reviewing digital identification products. In other instances, we built on and/or extended existing concepts from security, privacy, and related literatures to develop risk factors, e.g., trusted versus non-trusted parties, selective disclosure [56, 82], unlinkability [35, 60], revocability [77], on-device processing [1, 55], shared use [51], notice and consent [4, 23, 58, 65], and inclusion and equity [5, 83, 91]. We do not claim the list of risk factors is exhaustive; part of our contribution here is to prioritize and select from a larger set of possible considerations.³ Our final list included 22 risk assessment factors, and we then organized those risk factors into five main categories. As we progressed, we also took note of other guidance such as common pitfalls and mitigations for individual risk factors. At the end of the process, we had a casebook with a précis and a risk and mitigation analysis for each case, and a broader framework with 22 risk factors, recommended best practices, common pitfalls, mitigations, open questions, and other guidance for digital identification systems.

As an early pilot of the framework, we worked through two live product reviews in detail; gathered feedback from nine experts in security, privacy, at-risk users, law, and public policy; and iterated on the content of the framework. In one live product review, we engaged directly with the product team as well as multiple cross-functional stakeholders over an extended multi-month period during the product development cycle and review process. In the other, we primarily served the role of supporting launch reviewers performing a complex risk analysis. In both of these reviews, we found the framework covered key issues and was helpful in adding weight to the importance of best practices.

We then piloted the framework in eight additional product reviews, engaging with approximately 20 cross-functional stakeholders (representatives from product teams, AI ethics, human rights, legal, privacy, product policy, public policy, and security) within our organization over a period of approximately one year. The bulk of the reviewing took place throughout 2023 with minor activity continuing into 2024. Our typical approach was to begin our analysis as we normally would, by learning about the product and the current thoughts of other reviewers and stakeholders. Once we had formed

³For example, usability is a consideration for digital identification systems [45] but we do not include an overarching risk factor for it, as in our experience that is often well-covered in other aspects of product development and review. At the same time, we do include risk factors for specific aspects of usability which are highly salient and often problematic in digital identification systems, such as **shared use** and **notice and consent**.

our opinions, we revisited the framework to compare the risk factors with issues that had been identified thus far. Broadly, we found that issues that were not within clear scope of established review functions were sometimes overlooked or de-prioritized by others, but were highlighted in the framework, illustrating the importance of integrating review across organizational boundaries. In other words, the framework was useful for drawing attention to issues (often, second-order effects or societal concerns) that fall in organizational gaps. We also saw that the framework identified issues missed by less experienced reviewers, and there was inconsistency in which factors might be overlooked, illustrating the utility of the framework for product teams and reviewers less versed in these issues. We found that experts with a high level of practical, multi-disciplinary experience assessing digital identification systems or similar products did not typically miss factors in the framework, but it was a useful checklist and also a useful tool for advocating for the prioritization of issues that might otherwise receive less attention. We end by observing that, as noted in Section 1, a wide range of considerations must be taken into account, and our work was used as one tool among many to consider the safer deployment of technology.

4 Framework

We begin by presenting our framework, comprised of 22 risk assessment factors, shown in Tables 1 through 5. Each factor is presented with a more desirable “lower risk” choice on the lefthand side and a less desirable “higher risk” (or unnecessary risk) choice on the righthand side. For example, in Table 1 a **trusted parties** approach (lefthand side) is typically lower risk than a **non-trusted parties** approach (righthand side). In some cases these choices are binary while in others they may represent endpoints in a spectrum. The risk factors are grouped into the following five main categories of risk assessment considerations for digital identification systems which emerged from our analysis:

- **High-Level Approach** (Table 1)
- **Desirable Properties** (Table 2)
- **Recommended Techniques** (Table 3)
- **Undesirable Technical Consequences** (Table 4)
- **Societal Impact** (Table 5)

After presenting our framework, we discuss how to use the framework, give an example application of the framework, and highlight nuances of specific risk factors.

4.1 How to Use the Framework

In a typical approach to using the framework, we envision it as a list of considerations that reviewers can work through when evaluating a given system, product or proposal. This evaluation may occur at various stages in the product development life cycle, from ideation to launch or even post-launch review. Consistent with privacy by design and related approaches that advocate early reviewer engagement and proactive consideration of best practices [17], ideally we recommend that evaluation begin during the design stage (a high leverage moment to identify issues and recommend implementation paths) and continue through the life of the product (perhaps with lightweight engagement if issues were addressed early). However,

the framework can be useful at any stage, and in practice, it may often be invoked at launch review.

The reviewers may be the developers or designers of a system, perhaps without expertise in digital identification systems. At the same time, given that digital identification is in many ways a fairly nascent space, it is often beneficial to include expert reviewers; these may include cross-functional experts from a range of domains, such as human rights, law, privacy, security, user experience, and more, working together to share ideas and facilitate a holistic review.

Reviewers can proceed by working through the framework, considering each risk factor in turn.⁴ For each risk factor, reviewers can assess whether the product design follows the “lower risk” practice, falls on the spectrum in between the “lower risk” practice and the “higher risk” practice,⁵ or follows the “higher risk” practice; or alternatively determine the factor is not applicable. If the factor is applicable and the system does not follow the “lower risk” practice, the reviewers can debate whether it would be possible to adjust the product to follow the more desirable practice, or move closer towards it, and if so, the pros and cons of making that change.

In some cases it is not viable to follow or more closely approximate the lower risk practice. For example, certain best practices for privacy and security are more easily followed in some infrastructures than others. Consider on-device processing, which often offers protections that server-side processing does not.⁶ The architecture of some offerings makes it difficult or infeasible to limit processing to on-device.

Overall, **if a system does not follow the desirable “lower risk” practice, reviewers should robustly consider potential harms and alternative product choices.** This analysis should naturally take into account a wide range of considerations not only from this framework but from many other tools and stakeholders as well, leading to a recommendation whether to move forward as is, or to modify or even cancel the product.

The description above notwithstanding, we hope our readers may use the framework however it best suits their needs, and we imagine it could productively be used as a type of checklist, or modified for use as a scoring system. Naturally we also encourage organizations to tailor the framework as appropriate, consistent with [78]. Based on our own experience with internal testing, the following may also be particularly useful in deployment in other organizations: apply the framework in the design stage and beyond; encourage use of the framework by non-experts; and attend carefully to issues in the framework that have no clear owner.

⁴We have found it reasonable to work through the risk factors in the order presented, but it is certainly sensible for others to follow whatever order seems most useful for their own purposes.

⁵As mentioned above, some factors may be binary and some may fall on a spectrum, depending on the risk factor and its interaction with specific system details.

⁶We do however recognize the limitations of on-device approaches. For example, while on-device approaches protect biometric templates being accessed or stored, they do not protect against physical acquisition of devices.

Table 1: Risk assessment factors for **High-Level Approach** (major design choices). Risk assessment factors in Tables 1–5 may be binary or may represent endpoints in a spectrum. Unless otherwise noted, examples throughout the paper are intended to be abstract and/or fictional rather than representing real-world products.

High-Level Approach	
Lower Risk	Higher Risk
<p>Trusted Parties</p> <p><i>Example: A service to present legal ID will display credentials only to government verifiers such as the TSA</i></p>	<p>Non-Trusted Parties: Sharing credentials with non-trusted parties can lead to data proliferation</p> <p><i>Example: Many parties (often non-trusted verifiers) request users provide health credentials to access physical spaces such as restaurants, bars, gyms, and more</i></p>
<p>Effective: The proposed mechanism is sufficiently accurate and feasible to achieve the desired results</p> <p><i>Example: Face Unlock stops some, although not all, authentication attacks</i></p>	<p>Ineffective: The proposed mechanism is not sufficiently accurate or feasible to achieve the desired results</p> <p><i>Example: Relying on voice identification for authentication in situations where high accuracy is needed</i></p>
<p>Biometrics Uniquely Advantageous^a: If biometrics are used, biometrics appear inherently useful and well-suited to the use case; it may be possible to demonstrate that non-biometrics can not achieve the product goal</p> <p><i>Example: Use of facial recognition in combination with other credentials for access to high security areas in a workplace</i></p>	<p>Biometrics Used When Not Necessary: Biometrics are used but a non-biometric mechanism is available that would be less risky, and equally or more effective</p> <p><i>Example: Use of palm prints at point of purchase; while palm prints appear preferable to face matching, it may further reduce risk to rely on a proxy device or card [47, 89]</i></p>
<p>Government Actor Performing Government Function</p> <p><i>Example: Government builds an app and verification process to issue and present digital version of legal ID</i></p>	<p>Non-Government Actor Performing (Pseudo) Government Function</p> <p><i>Example: Third party becomes purveyor of legal ID, or third party verifies user owns a given legal ID</i></p>

^a Here, when we refer to biometrics, we are considering the case where the product/service provider will verify the biometrics and perhaps store them server side, not just use existing mechanisms like Touch ID. This guidance is not meant to discourage developers from using Touch ID.

Table 2: Risk assessment factors for Desirable Properties.

Desirable Properties	
Lower Risk	Higher Risk
<p>Reveals Minimum Information Necessary to Achieve the Goal: Selective disclosure^a</p> <p><i>Example: A health app presents vaccination status (e.g., up-to-date, not up-to-date) and no further details^b</i></p>	<p>Reveals More Information Than Necessary to Achieve the Goal</p> <p><i>Example: A health app displays all vaccination information such as the person’s name and date/location of vaccinations</i></p>
<p>Unlinkability: Cannot link user identity across different uses of the same credential</p> <p><i>Example: An age verification app generates one-time use credentials so that a user’s sessions can not be linked</i></p>	<p>Links An Identity to Other Data: Links one identity, such as a legal identity, to other data, such as online activity, OR Creates a Record of Transactions Linked to An Identity</p> <p><i>Example: Subway pass and associated transit activity are linked to user account including name, payment information, and other data [92]</i></p>
<p>Revocability: The issuer can easily revoke and reissue the credential</p> <p><i>Example: An app to display a government-issued driver’s license is architected in such a way that the government can revoke the license and/or issue an updated version of the license, e.g., with a new expiration date</i></p>	<p>Not Easily Revocable: The issuer can not easily revoke and reissue the credential^c</p> <p><i>Example: An app to display a government-issued driver’s license is architected in such a way that the government can not revoke and/or issue an updated version of the license</i></p>

^a Selective disclosure very often arises in the context of electronic presentation of government-issued credentials, such as government IDs or health certificates.

^b Consider also the application of selective disclosure for online age verification, where it has been proposed as an improvement over less privacy-conscious online mechanisms [9, 49, 75]; at the same time, some argue that in this domain online processes substantially increase risk over in-person presentation even with protections such as selective disclosure in place [24, 34].

^c We emphasize here potential disadvantages to users who may have difficulty getting up-to-date credentials, or societal risks of out-of-date credentials which enable unauthorized access or permissions.

Table 3: Risk assessment factors for **Recommended Techniques**.

Recommended Techniques	
Lower Risk	Higher Risk
<p>Ephemeral: Risk is lower if data is deleted (almost) immediately, even though data deletion is hard to guarantee</p> <p><i>Example: Biometric data is gathered when user is first issued a certificate, and then it is deleted; after issuance non-biometric credentials are used</i></p>	<p>Permanent Record: Sharing and holding data carries risk; stable identifiers are more risky than non-stable identifiers</p> <p><i>Example: Biometric data is gathered for authentication and then stored permanently</i></p>
<p>On Device</p> <p><i>Example: The use of facial recognition for unlocking a user’s phone can often be performed on device, greatly reducing potential access to biometric data</i></p>	<p>Server Side</p> <p><i>Example: Users can use their palm prints at point of purchase, enabled by a server-side database of palm prints [47, 89]</i></p>
<p>Initial Launch Incorporates All Privacy and Security Best Practices</p> <p><i>Example: Initial launch is delayed to implement a stronger security mechanism</i></p>	<p>Initial Launch Does Not Incorporate All Privacy and Security Best Practices</p> <p><i>Example: Initial launch includes a reasonable security protection mechanism, but not the strongest standard option available</i></p>
<p>Effective Protection for Shared Use: System thoughtfully supports multiple users/accounts per device; supports ID/credentials for minors</p> <p><i>Example: An app to present government ID provides separate profiles for different family members, and allows parents to manage ID for their children</i></p>	<p>Does Not Have Effective Protection for Shared Use: System does not sufficiently account for more than one user/account per device; does not support ID/credentials for minors</p> <p><i>Example: An app to present government ID allows ID from multiple legal names per account but does not offer security protections for the different identities</i></p>
<p>Notice/Consent Provided: Subjects are notified, provided meaningful consent, and have a viable opt-out^a</p> <p><i>Example: Meeting attendees are offered the opportunity to opt-in to a system that identifies speakers and creates a transcript with names</i></p>	<p>Notice/Consent Not Provided: Subjects are not notified, are not provided meaningful consent, and/or do not have a viable opt-out</p> <p><i>Example: People are observed in a public place for urban planning purposes, but are not made aware of ongoing identification and/or are not offered an opt-out^b</i></p>
<p>A Plan is In Place for Outages and Failures: Users will be well-supported in the event of outages and failures</p> <p><i>Example: Reasonable service guarantees are established</i></p>	<p>A Plan is Not In Place for Outages and Failures: Users will need to rely on physical documents or expend undue effort in the event of outages and failures</p> <p><i>Example: Users of a digital ID presentation app are told to carry their physical ID with them at all times^c</i></p>

^a We note that as discussed in the literature, while notice and consent can be remediating factors, they may not fully ameliorate risks [4, 58, 65].

^b While details are not specified in the article, notice/consent issues may pertain in a situation such as the one posed in [15].

^c Conversely, consider that in some situations it may be advisable to use physical ID rather than digital ID [81].

Table 4: Risk assessment factors for Undesirable Technical Consequences.

Undesirable Technical Consequences	
Lower Risk	Higher Risk
<p>Does Not Significantly Increase Risk of Fraud/Abuse: System is resistant to fraud/abuse, does not create new threats, and/or reduces pre-existing threats</p> <p><i>Example: Rolling out smart cards for credit cards reduced existing fraud (detection of which previously relied on risk models)</i></p>	<p>Facilitates Fraud/Abuse: System is easily circumvented and/or creates new threats^a</p> <p><i>Example: Third party gathers biometric data which if breached could in theory be synthesized or replayed at other sites in order to gain access or privileges</i></p>
<p>Data Set Already Exists: Risk is lower if the data set already exists, particularly if it is already public</p> <p><i>Example: Use of public images of celebrities to identify/prevent celebrity impersonation on social media sites</i></p>	<p>Constructs a Biometric Data Set: System creates a new biometric data set; concerning even if held privately by trusted party, and even more concerning if publicly available or held by non-trusted party</p> <p><i>Example: Creating a new biometric data set to facilitate provision of services</i></p>
<p>Verifier Already Has Subject’s Information: Verifier already knows the information being gathered^b</p> <p><i>Example: Travel authority already knows that a user with a given legal name has purchased a ticket from one location to another, and is simply verifying that the person holding the ticket has that legal identity</i></p>	<p>Verifier Acquires Subject’s Information</p> <p><i>Example: A bartender checking a person’s government ID to confirm their age does not previously have access to their legal name, address, or other information shown on their ID</i></p>

^a In considering fraud and abuse, we emphasize here potential disadvantages to users including not only the risk of falling victim to fraud and abuse but also potential burdens placed on identity holders, such as increased effort or knowledge. Accordingly, we advocate that costs should lie as much as possible with the issuer, with some costs likely falling to verifiers if they have additional privileges or capabilities over credential holders, and costs should be minimized for users to the extent possible.

^b In many workplace scenarios, or scenarios with high physical security, the verifier already has the subject’s information.

Table 5: Risk assessment factors for Societal Impact.

Societal Impact	
Lower Risk	Higher Risk
<p>Misidentification Does Not Cause Substantial Harm^a</p> <p><i>Example: Identify individuals for the purpose of computing aggregate metrics about space usage in the workplace and then discard the identifications</i></p>	<p>Misidentification May Cause Substantial Harm: False positives/negatives may result in substantial harm</p> <p><i>Example: A facial recognition system that identifies homeless individuals [19] might mistakenly label an individual as “assaultive” resulting in dangerous confrontation with law enforcement</i></p>
<p>Surveillance of Specific Individuals for Specific Reasons</p> <p><i>Example: In a system to monitor financial transactions, identification is made only for transactions above a large mandatory reporting amount^b</i></p>	<p>Surveillance of Broad Population: Broad surveillance without specific evidence against all individuals being surveilled</p> <p><i>Example: Automatically connecting to a government resource (e.g., each time a user wants to present their ID to a verifier) could leak device information, IP address, or other data from all users^c</i></p>
<p>Follows Best Practices for Inclusion and Equity: Product experience is inclusive and/or expected impacts are equitable</p> <p><i>Example: Identification is roughly equally accurate for members of different groups</i></p>	<p>Is Not Inclusive or Equitable: Product experience is not inclusive and/or expected impacts are not equitable [53, 66]</p> <p><i>Example: Has a negative impact on populations without access to high-end devices</i></p>
<p>Follows Widely Accepted Principles of Human and Civil Rights^d</p> <p><i>Example: Vaccine card product undergoes a human rights assessment prior to launch in order to identify and mitigate potential issues</i></p>	<p>Product Experience/Impacts Do Not Follow Widely Accepted Principles of Human and Civil Rights</p> <p><i>Example: A health app makes it easier to identify certain at-risk groups^e</i></p>
<p>Similar Jurisdictions Are Offered Similar Products and Protections</p> <p><i>Example: A driver’s license product provides the same privacy and security protections to every state in the US</i></p>	<p>Similar Jurisdictions Are Offered Different Products and Protections</p> <p><i>Example: A driver’s license product provides different privacy and security protections to different states in the US</i></p>
<p>Compelling Benefit: Use case is compelling, so a full risk/benefit analysis may be warranted</p> <p><i>Example: Provide identification for people who are incapacitated</i></p>	<p>Limited Benefit: Use case is not compelling, so even minor risk seems unwarranted</p> <p><i>Example: Use biometric identification to provide minor customer service experience improvements, such as greeting guests^f</i></p>

^a Here we generally consider misidentifications in which the system makes an inadvertent mistake, although some of these harms could also arise from intentional fraud or abuse.

^b Consider [57].

^c Consider also use cases discussed in [31].

^d For discussion of human rights considerations in digital identification systems, see [7].

^e As another example, consider concerns raised in [53].

^f Consider [59].

4.2 Example Application of the Framework

Here we return to an example raised in the introduction, the use of iris scans for authentication for food and cash distribution in refugee camps [44, 46, 72, 76]. In practice, one would work through all risk factors in the framework. Here we focus on several of the more salient risk factors to consider for that case:

- **Biometrics Used When Not Necessary.** While iris scans offer advantages such as convenience and may open up access to people without devices and/or reduce fraud [46, 76], lower tech or non-biometric techniques such as a code may be sufficiently effective and incur less risk [46].
- **Constructs a Biometric Dataset.** It is likely the case that a database of iris scans does not previously exist, and constructing a data set of stable, non-alterable identifiers of a sensitive population introduces risk of misuse or security attacks [46, 76].
- **Product Experience and/or Impacts May Not Follow Widely Accepted Principles of Human and Civil Rights.** Given arguments raised by refugee advocates regarding surveillance, consent, and personal safety threats [44, 46, 76], a robust human rights impact assessment should likely be performed to assess possible harms [76].

4.3 Risk Factor Details

While we do not discuss each risk factor individually for the sake of brevity, many of the risk factors are largely self-explanatory and do not require additional detail. However, some risk factors benefit from additional comment. In this section, we elaborate on more complex risk factors and highlight common pitfalls. Readers may find it useful to begin by reading Tables 1–5, consulting the details in this section as needed.

4.3.1 Trusted Parties ... Non-Trusted Parties (Table 1). Policies and terms of service can specify how identification systems will be used, and can limit access to trusted parties. For example, access to sensitive technologies such as facial recognition can be controlled via allow lists (see e.g., [3]). While limiting access to trusted parties is ameliorative, and is generally lower risk, one should not assume this is a complete remediation. For example, it is a pitfall to mistakenly assume good actors can keep data secure, as public sector actors (e.g., public schools) and others may have not have strong protective measures in place. Relatedly, sometimes an argument is made that it is okay for a trusted actor to execute a potentially problematic practice because the trusted actor will execute it responsibly; however, this is a problematic argument since practices and tools (e.g., APIs) that are established and normalized by “good” actors can empower “bad” actors. We recommend carefully evaluating such assumptions and arguments regarding trusted parties.

4.3.2 Effective ... Ineffective (Table 1). One common pitfall in digital identification systems is to choose an approach that is not feasible for the population/situation, e.g., asking rideshare drivers to authenticate themselves while driving, thereby creating hazardous conditions [85]; or incorrectly assuming a given user population has capacity to manage a particular type of secret, account, or device. An advisable mitigation is to engage with community partners and/or conduct user research to confirm that a given approach is

appropriate (see Section 5.3 for more discussion of stakeholder engagement).

Another common pitfall is to employ voice and/or affect recognition as a primary factor for identification, a task for which they typically are not sufficiently accurate/secure. Voice has fewer features than facial recognition and is often not performed in a protected environment on device. Further, deep fake technology can make it easy to circumvent voice (see [20] for an example of an AI-generated synthetic voice successfully authenticating a user at a financial institution); voice clips from high profile users may be readily available, or people who have been close to the user can provide data, to relatively easily train a generative model that is sufficiently good to pass voice authentication models. Rather than relying on voice and/or affect recognition as a primary factor, it is advisable to use a different mechanism entirely [20]; and/or to use voice, affect recognition, and other lower accuracy identifiers as secondary factors (i.e., use them as a component in multi-factor identification).

4.3.3 Biometrics Uniquely Advantageous ... Biometrics Used When Not Necessary (Table 1). We recommend assessing whether biometrics offer a genuine advantage for the use case, versus being used unnecessarily. Biometrics are often proposed when less invasive measures can achieve comparable effects. It is perhaps natural for developers to gravitate towards biometric approaches as, for example, they are a convenient way to provide usable security without asking humans to remember a password or other secret. However, biometrics introduce a number of substantial risks, e.g., they are often (near) stable identifiers over the lifetime of an individual, so once the information has been gathered, it can enable long-term tracking [25, 36]; or they may be gathered non-volitionally.⁷ Due to the sensitivity of biometrics, it is often advisable to use alternative mechanisms that provide a level of indirection. It is often wise to scope a system to achieve exactly what is intended and no more, for example, to create a temporary identity that can be used only in a specific situation and not more broadly.

Here is one common pitfall: unnecessarily using biometrics when “device as proxy” meets the use case. As the sophistication of consumer devices has improved, it is now possible for those devices to make strong guarantees about access control to cryptographic key material and associated operations. This capability is often leveraged where the device acts as a proxy for the user’s identity by registering key material the device holds on behalf of the user and uses only after authenticating them.

Although “device as proxy” is already the de-facto standard for many use cases (e.g., sign-in, autofill), too often products across the industry fail to adopt this best practice. We recommend guarding against this, and using “device as proxy” when practicable.⁸ Ever since Touch ID on the iPhone, and perhaps before, it has been a well-known, well-accepted security best practice to gather and keep biometrics on device and then use the device as a proxy to submit something other than biometrics to a server. A typical model is

⁷For example, law enforcement may compel people to unlock their devices with biometrics, even if they can not compel them to provide passwords ([13, 81]). As another example, face biometrics may be gathered easily in public spaces.

⁸In assessing practicality, we recommend considering the bottom tier of devices that can provide such proxy functionality, and assessing whether it meets the use case.

that the user inputs biometrics such as a fingerprint to authenticate themselves on a device in their control, and then the device becomes the interface to other less trusted services. In this way, a user's device can become a trusted proxy to translate between on-device authentication and external services, bridging between easy-to-use authentication for the user and authentication with desirable properties for online use (e.g., cryptographic security and revocability). Some apps use QR codes for this purpose, and for specific use cases (e.g., payments, presentation of mobile driver licenses) NFC-based protocols are available.

4.3.4 Unlinkability ... Links An Identity to Other Data (Table 2). A system may be designed for unlinkability (see e.g., the approach outlined in [41]), such that user identity can not be linked across different uses of the same credential. This affords a number of privacy advantages, and is desirable for many everyday use cases [36]. In fact, Goodell and Aste argue that “support for multiple, unlinkable identities is an essential right” [36]. Unlinkability can be supported by generating one-time use versions of the credential, and/or by using credential formats and methods that support unlinkability, e.g., generating credentials that do not re-use key material, time stamps (credential expiration, etc.), or other correlatable data across credential presentations.

However, unlinkability may be at odds with certain types of personalization or other potential benefits, as it is in opposition to recognizing the same user at a repeat visit. Some systems may therefore support linkability, e.g., creating a record of transactions linked to an identity. Using an anonymized identity does not provide full protection in this case, as the record of transactions may reveal deanonymizing or otherwise private information.

4.3.5 Revocability ... Not Easily Revocable (Table 2). An often overlooked complexity in issuing digital identity and credentials is that they may expire (e.g., consider expiration dates on driver's licenses) or be revoked (e.g., consider a driver's license revoked by a state agency). Presentation of physical identity documents often relies on the notion that a human will examine the document in person and notice if it is out-of-date. However, in many use cases digital identity documents will likely not be examined by a human (for example, they may be checked automatically) and/or the expiration date may not be visible on a digital identity document due to selective disclosure. Accordingly, digital identification systems should consider mechanisms to manage the lifecycle of identity and credentials, including renewal and revocation.

In developing a plan to allow issuers to reissue and revoke credentials, it is important to explore considerations such as the following: whether the approach is resilient to malicious users and/or attackers; and whether the approach will leak data about the user. For example, there may be a strong correlation between criminal activity such as a DUI and revoked credential status, so it is advisable to consider privacy-protective mechanisms for revoking credentials.

4.3.6 Does Not Significantly Increase Risk of Fraud/Abuse ... Facilitates Fraud/Abuse (Table 4). A given digital identification system may substantially impact the risk of fraud/abuse. For example, it may be easy to circumvent a given system, or it may be easy to use deep fake technology or otherwise mislead or bypass it. Further,

for some high profile individuals, much or all of the necessary information to obtain a digital identity credential may be publicly accessible, so additional protections may be appropriate for such individuals.

As another example, a given identification system may introduce a new exploit or line of attack, creating significant new risk of fraud or abuse beyond the product. For example, it may facilitate impersonation in other services, e.g., by constructing a high quality voice model that could be used for voice generation. Alternatively, it may allow a fraudulent purveyor to gather identity metadata from many customers and reuse or sell that data, or the system may be designed such that it can be appropriated to intentionally inflict individual harm.

Another pitfall is that recovery mechanisms can enable identity theft. Recovery mechanisms are a difficult problem and, currently, they are often (one of) the most insecure part(s) of a system; unfortunately, account recovery often depends on less secure information than the primary authentication mechanism. Therefore, recovery mechanisms are frequently targeted for fraud and abuse. One advisable rule of thumb is that the recovery mechanism should be at least as strong as the actual mechanism.

4.3.7 Surveillance of Specific Individuals for Specific Reasons ... Surveillance of Broad Population (Table 5). Many commenters have observed that digital identification systems pose significant surveillance risks. One common pitfall is choosing an approach that is unnecessarily expansive, e.g., gathers highly sensitive data when less sensitive data would work as well. In this situation it is recommended to choose instead a targeted mechanism for the precise use case.

Linkability is also an important consideration for surveillance, and there are often opportunities to limit it. For example, one can implement static rather than dynamic issuance of government ID for presentation (in the dynamic case, a system could pass each request from the verifier to the issuer, posing privacy challenges). As another example, if long-term surveillance is in place (for example, for urban planning) and biometrics are used, such biometrics can be used exclusively to trade (bootstrap) for a code for regular use to avoid long-term use and surveillance with biometrics.

Another often overlooked consideration is the potential for a service from a “good” actor to be used in a different way. Think of dual-use of a given technology, an issue often raised in responsible AI (see e.g., [61]), and assess how quickly and easily a provider might switch from providing a service to conducting surveillance. Further, consider whether the product or feature introduces additional risk of search or data requests. For example, a user stopped by law enforcement might be required to use their phone to transmit their mobile driver's license data. Consider whether there is any risk the phone might then be subject to additional search.

5 Discussion

In this section, we discuss some of the overarching considerations regarding privacy, security, and safety that arise across multiple risk factors in the framework.

5.1 Standards That Enhance Privacy and Security

Digital identification requires coordination across a large ecosystem of issuers, verifiers, technology providers, and others. It is difficult for a single player to implement a best practice for digital identity in isolation, although in some cases it is possible for a single technology provider to take an existing standard and layer a best practice on top of it. Standards are a good mechanism for coordination but take a long time (typically years) to put in place, and often lag current technology and best practice. To ensure the ecosystem moves towards best practices as efficiently as possible, standards for digital identity are an important area of investment, especially to support desirable properties called out in the framework such as **selective disclosure, unlinkability, and revocability**.

As an illustrative example, consider one of the most valuable affordances of digital identification, selective disclosure, i.e., the ability to present only the minimal information required in a particular circumstance. For example, when presenting a physical driver's license to prove their age in order to purchase alcohol at a bar, a person typically shows the front of their driver's license which includes not only their age but also personal information such as their home address. By contrast, with digital identification, that same person could show only their age along with a photo.⁹ However, in practice, selective disclosure is often not used. For example, electronic presentation of proof of COVID-19 vaccination has typically involved display of full information about vaccine administration (type, number, date) from which sensitive information such as age or medical status could be inferred. A more security-conscious implementation would simply confirm whether or not the person is up-to-date on their vaccinations [2].

Individual actors can not easily establish or enforce selective disclosure, as it must be implemented across issuers, verifiers, and presentation apps; rather this technique is more effectively enforced through cooperative agreements and standards. Standards bodies should more aggressively pursue this important strategy.

5.2 Launch Velocity

Companies may experience significant pressure to bring digital identification systems or products to market quickly. For example, they may feel time pressure to bring a partially developed product to market to "hold their place" in a rapidly changing landscape, even if current functionality offers only limited user benefit. Further, they may favor the popular launch-and-iterate innovation strategy that emphasizes agility and continuous improvement based on releasing products with reasonable (but perhaps not ideal) functionality to get real-world experience and feedback.

These pressures notwithstanding, in this domain it is generally inadvisable to prioritize launch velocity over best practices in privacy, security, fairness, human rights, and other values. The launch-and-iterate innovation strategy founders when actions with serious consequences can not be undone. In high risk or sensitive domains such as digital identity, real-world experimentation can be

⁹We note that, while highly protective, selective disclosure may not offer complete privacy protection if credentials with selective disclosure have opaque but persistent, unique value. It is advisable to implement and advocate for standards for single-use credentials with opaque attributes that appear different across different instances of the user's credential.

risky. If something harmful is launched, it may not be possible to roll it back. Key data may already have been exposed, or harmful precedent may have been set for non-ideal practices that other companies then emulate.

In the digital identification space, it is advisable to move thoughtfully and carefully, minimizing unnecessary exposure by drawing as much as possible on well-established practices such as adversarial testing. We also note that prioritizing launch velocity often creates false dilemmas that can be resolved by thoughtful design and/or additional engineering work. At the same time, "launching and iterating" and "launching safely" are not entirely at odds. Carefully designed product roadmaps that lay out a path for the next several launches, accompanied by principled protections for each launch, can go a long way towards mitigating potential hazards of an incremental and agile launch process.

5.3 Stakeholder Engagement

Digital identification systems raise many complex sociotechnical issues. For example, equity is a key consideration for identity, particularly where government-issued credentials are concerned. Accessibility and fairness metrics (for example, measuring performance for different groups) are now increasingly common during product development and launch review. Analyses should also go beyond these issues to consider broadly who is vulnerable and how the product might affect them, for example, children, elders, survivors of intimate partner abuse, refugees, and others. For example, are recovery mechanisms equally accessible to different groups, e.g., do in-person recovery mechanisms support urban and rural users equally well? Consider also non-users who may be affected by the product, for example, because it normalizes use of digital identity tools to which they do not have ready access and adds friction to the use of physical documents. Is deployment of the technology likely to create friction for the use of physical documents, disadvantaging those without access to digital identification?

Such questions are best addressed in partnership with potentially affected communities. Community members, advocates, and regional experts can provide valuable insight regarding real-world impacts of digital identification, including weighing in on speculated harms and proposed mitigations. For example, intimate partner violence scenarios are nuanced and it can be difficult to predict how a given digital identification feature may be (mis)appropriated, especially given complexities with shared devices. Advocates can shed light on harmful scenarios and help prioritize mitigations.

As a result, stakeholder engagement increasingly features in regulatory perspectives, and robust community partnerships are highly advisable to meet these expectations and thoughtfully address sociotechnical issues.

6 Conclusions

We introduced a risk assessment framework for digital identification systems, as well as best practices for them. It is intended as a step towards systematizing knowledge in this emergent and complex domain, drawing on expert experience and assessment [90] to prioritize among a large set of possible considerations and to highlight certain issues that in our experience are often overlooked. In presenting these resources, we hope to provide utility across a

spectrum of users, from developers and regulators to those building agentic systems for process automation to ensure digital identification products are both effective and beneficial.

Beyond the work we have presented here, this space is rife with opportunity for new innovations, standards, and best practices. For example, the technology industry has not yet produced standard protection mechanisms for shared devices or accounts, an issue which interacts heavily with digital identification systems. While this is a challenging problem, it is well worth making an effort to advance the state of the art in this area.

Acknowledgments

We thank Elie Bursztein, Emily Everett, Jen Gennai, Vi Vien Hoang, Patrick Gage Kelley, Kurt Thomas, Amanda Walker, Lawrence You, and our other excellent colleagues at Google for their valuable collaboration, insights, and feedback on this work.

References

- [1] Abdulmalik Alwarafy, Khaled A. Al-Thelaya, Mohamed Abdallah, Jens Schneider, and Mounir Hamdi. 2021. A Survey on Security and Privacy Issues in Edge-Computing-Assisted Internet of Things. *IEEE Internet of Things Journal* 8, 6 (2021), 4004–4022. <https://doi.org/10.1109/JIOT.2020.3015432>
- [2] Anonymous. 2021. Privacy-Preserving Medical Credentials for Access Authorization. Technical Disclosure Commons. https://www.tdcommons.org/dpubs_series/4334/.
- [3] Parker Barnes and Andrew Schwartz. 2019. Celebrity Recognition now available to approved media & entertainment customers. Google Cloud Blog. <https://cloud.google.com/blog/products/ai-machine-learning/celebrity-recognition-now-available-to-approved-media-entertainment-customers>.
- [4] Solon Barocas and Helen Nissenbaum. 2009. On Notice: The Trouble with Notice and Consent. In *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information*. <https://ssrn.com/abstract=2567409>.
- [5] Solon Barocas and Andrew D. Selbst. 2016. Big Data’s Disparate Impact. *California Law Review* 104, 3 (June 2016), 671–732. <https://www.jstor.org/stable/24758720>.
- [6] BBC. 2021. Schools warned over facial recognition systems. *BBC* (19 Oct. 2021). <https://www.bbc.com/news/technology-58969136>.
- [7] Ana Beduschi. 2021. Rethinking digital identity for post-COVID-19 societies: Data privacy and human rights considerations. *Data & Policy* 3 (2021), e15. <https://doi.org/10.1017/dap.2021.15>
- [8] Angus Berwick. 2024. Sam Altman’s Worldcoin Is Battling With Governments Over Your Eyes. *The Wall Street Journal* (18 Aug. 2024). <https://www.wsj.com/tech/sam-altman-openai-humanness-iris-scanning-4d0e1dab>.
- [9] Olivier Blazy. 2024. Online Age Verification and Privacy Protection: An Impossible Equation? Stanford Cyber Policy Center Seminar. <https://cyber.fsi.stanford.edu/events/may-7-online-age-verification-and-privacy-protection-impossible-equation>.
- [10] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2012. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *2012 IEEE Symposium on Security and Privacy*. 553–567. <https://doi.org/10.1109/SP.2012.44>
- [11] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. <https://doi.org/10.1080/2159676X.2019.1628806>
- [12] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology* 18, 3 (2021), 328–352. <https://doi.org/10.1080/14780887.2020.1769238>
- [13] Thomas Brewster. 2021. FBI: You Don’t Have To Tell Us Which Body Part Unlocks Your Smartphone—But We Will Guess Anyway. *Forbes* (29 Nov. 2021). <https://www.forbes.com/sites/thomasbrewster/2021/11/29/fbi-no-consent-required-for-forced-phone-unlocks-via-finger-or-face/>.
- [14] Ryan Browne. 2022. Mastercard launches tech that lets you pay with your face or hand in stores. *CNBC Tech* (17 May 2022). <https://www.cnbc.com/2022/05/17/mastercard-launches-tech-that-lets-you-pay-with-your-face-or-hand.html>.
- [15] Emma Bubola. 2021. Venice, Overwhelmed by Tourists, Tries Tracking Them. *The New York Times* (4 Oct. 2021). <https://www.nytimes.com/2021/10/04/world/europe/venice-tourism-surveillance.html>.
- [16] California DMV (Department of Motor Vehicles). [n.d.]. CA DMV Wallet & mDL Pilot. <https://www.dmv.ca.gov/portal/ca-dmv-wallet/>.
- [17] Ann Cavoukian. 2009. Privacy by Design: The 7 Foundational Principles. https://student.cs.uwaterloo.ca/~cs492/papers/7foundationalprinciples_longer.pdf.
- [18] Chris Aston Chen. 2021. Introducing Age Verification. <https://blog.roblox.com/2021/09/introducing-age-verification/>.
- [19] Joseph Cox. 2022. Tech Firm Offers Cops Facial Recognition to ID Homeless People. *Vice* (8 Feb. 2022). <https://www.vice.com/en/article/wxdp7x/tech-firm-facial-recognition-homeless-people-odin>.
- [20] Joseph Cox. 2023. How I Broke Into a Bank Account With an AI-Generated Voice. *Vice* (23 Feb. 2023). <https://www.vice.com/en/article/how-i-broke-into-a-bank-account-with-an-ai-generated-voice/>.
- [21] Lorrie Faith Cranor. 2008. A framework for reasoning about the human in the loop. In *Proceedings of the 1st Conference on Usability, Psychology, and Security* (San Francisco, California) (*UPSEC ’08*). USENIX Association, Article 1, 15 pages. https://www.usenix.org/legacy/events/upsec08/tech/full_papers/cranor/cranor.pdf.
- [22] A.J. Dellinger. 2019. US prisons are reportedly creating ‘voice print’ databases. *engadget* (31 Jan. 2019). <https://www.engadget.com/2019-01-30-us-prison-voice-print-database.html>.
- [23] Rachna Dhamija and Lisa Dusseault. 2008. The Seven Flaws of Identity Management: Usability and Security Challenges. *IEEE Security & Privacy* 6, 2 (2008), 24–29. <https://doi.org/10.1109/MSP.2008.49>
- [24] Electronic Frontier Foundation (EFF). [n.d.]. Age Verification Harms Users of All Ages. <https://www.eff.org/document/age-verification-harms-users-all-ages>.
- [25] Electronic Frontier Foundation (EFF). [n.d.]. Biometrics. <https://www.eff.org/issues/biometrics>.
- [26] European Commission. [n.d.]. Discover eIDAS. <https://digital-strategy.ec.europa.eu/en/policies/discover-eidas>.
- [27] European Commission. [n.d.]. European Digital Identity. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en.
- [28] Federal Trade Commission. 2023. FTC Warns About Misuses of Biometric Information and Harm to Consumers. <https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-warns-about-misuses-biometric-information-harm-consumers>.
- [29] FIDO Alliance. [n.d.]. Passkeys: Accelerating the Availability of Simpler, Stronger Passwordless Sign-Ins. <https://fidoalliance.org/passkeys/>.
- [30] Rainhard D. Findling and Rene Mayrhofer. 2012. Towards face unlock: on the difficulty of reliably detecting faces on mobile phones. In *Proceedings of the 10th International Conference on Advances in Mobile Computing & Multimedia* (Bali, Indonesia) (*MoMM ’12*). Association for Computing Machinery, New York, NY, USA, 275–280. <https://doi.org/10.1145/2428955.2429008>
- [31] Eve Gaumond. 2021. Artificial Intelligence Act: What Is the European Approach for AI? *Lawfare*. <https://www.lawfaremedia.org/article/artificial-intelligence-act-what-european-approach-ai>.
- [32] Barney G. Glaser and Anselm L. Strauss. 1967. *Discovery of Grounded Theory: Strategies for Qualitative Research*. Aldine.
- [33] Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. Bantam Doubleday Dell Publishing Group.
- [34] Eric Goldman. 2025. The “Segregate-and-Suppress” Approach to Regulating Child Safety Online. *Stanford Technology Law Review* 28, 2 (2025), 173–232. <https://law.stanford.edu/publications/the-segregate-and-suppress-approach-to-regulating-child-safety-online/>.
- [35] Marta Gomez-Barrero, Javier Galbally, Christian Rathgeb, and Christoph Busch. 2018. General Framework to Evaluate Unlinkability in Biometric Template Protection Systems. *IEEE Transactions on Information Forensics and Security* 13, 6 (June 2018), 1406–1420. <https://doi.org/10.1109/TIFS.2017.2788000>
- [36] Geoff Goodell and Tomaso Aste. 2019. A Decentralized Digital Identity Architecture. *Frontiers in Blockchain* 2 (2019), 17. <https://doi.org/10.3389/fbloc.2019.00017>
- [37] Google Nest Help. 2025. Learn about familiar face detection. Google Nest Help Center. <https://support.google.com/googlenest/answer/9268625>.
- [38] Peter Leo Gorski, Luigi Lo Iacono, and Matthew Smith. 2023. Eight Lightweight Usable Security Principles for Developers. *IEEE Security & Privacy* 21, 1 (Jan. 2023), 20–26. <https://doi.org/10.1109/MSEC.2022.3205484>
- [39] Sandeep Gupta, Attaullah Buriro, and Bruno Crispo. 2019. DriverAuth: A risk-based multi-modal biometric-based driver authentication scheme for ride-sharing platforms. *Computers & Security* 83 (June 2019), 122–139. <https://doi.org/10.1016/j.cose.2019.01.007>
- [40] Alex Heath. 2025. Sam Altman’s eye-scanning project launches cryptocurrency in the US. *The Verge* (30 April 2025). <https://www.theverge.com/cryptocurrency/659011/worldcoin-us-launch-orb-crypto-sam-altman>.
- [41] Thomas S. Heydt-Benjamin, Hee-Jin Chae, Benessa Defend, and Kevin Fu. 2006. Privacy for Public Transportation. In *International Workshop on Privacy Enhancing Technologies (PET 2006)*. Springer, 1–19. https://doi.org/10.1007/11957454_1
- [42] Ash Johnson. 2024. The Path to Digital Identity in the United States. <https://itif.org/publications/2024/09/23/path-to-digital-identity-in-the-united-states>.
- [43] George Joseph and Debbie Nathan. 2019. Prisons across the US are quietly building databases of incarcerated people’s voice prints. *The Intercept* (30 Jan. 2019). <https://theintercept.com/2019/01/30/prison-voice-prints-databases-securus/>.
- [44] Russ Juskalian. 2018. Inside the Jordan refugee camp that runs on blockchain. *MIT Technology Review* (12 April 2018). <https://www.technologyreview.com/2018/04/12/143410/inside-the-jordan-refugee-camp-that-runs-on-blockchain/>.

- [45] Maina Korir, Simon Parkin, and Paul Dunphy. 2022. An Empirical Study of a Decentralized Identity Wallet: Usability, Security, and Perspectives on User Control. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*. 195–211. <https://usenix.org/conference/soups2022/presentation/korir>.
- [46] Anna Kramer. 2021. Refugees are buying groceries with iris scans. What could go wrong? *protocol* (16 Sept. 2021).
- [47] Dilip Kumar. 2020. Introducing Amazon One—a new innovation to make everyday activities effortless. *Innovation at Amazon* (29 Sept. 2020). <https://www.aboutamazon.com/news/innovation-at-amazon/introducing-amazon-one-a-new-innovation-to-make-everyday-activities-effortless>.
- [48] Ruggero Donida Labati, Angelo Genovese, Enrique Muñoz, Vincenzo Piuri, Fabio Scotti, and Gianluca Sforza. 2016. Biometric Recognition in Automated Border Control: A Survey. *Comput. Surveys* 49, 2, Article 24 (June 2016), 39 pages. <https://doi.org/10.1145/2933241>
- [49] Lauren Leffer. 2024. Online Age Verification Laws Could Do More Harm Than Good. *Scientific American* (April 2024). <https://www.scientificamerican.com/article/online-age-verification-laws-privacy/>.
- [50] Tara Matthews, Elie Bursztein, Patrick Gage Kelley, Lea Kissner, Andreas Kramm, Andrew Oplinger, Andreas Schou, Manya Sleeper, Stephan Somogyi, Dalila Szostak, Kurt Thomas, Anna Turner, Jill Palzkill Woelfer, Lawrence L. You, Izzie Zahorian, and Sunny Consolvo. 2025. Supporting the Digital Safety of At-Risk Users: Lessons Learned from 9+ Years of Research & Training. *ACM Transactions on Computer-Human Interaction* (Feb. 2025). <https://doi.org/10.1145/3716382>
- [51] Tara Matthews, Kerwell Liao, Anna Turner, Marianne Berkovich, Robert Reeder, and Sunny Consolvo. 2016. “She’ll just grab any device that’s closer”: A Study of Everyday Device & Account Sharing in Households. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI ’16). Association for Computing Machinery, New York, NY, USA, 5921–5932. <https://doi.org/10.1145/2858036.2858051>
- [52] Elliot Mbunge, Tafadzwa Dzinamarira, Stephen G. Fashoto, and John Batani. 2021. Emerging technologies and COVID-19 digital vaccination certificates and passports. *Public Health in Practice* (May 2021). <https://doi.org/10.1016/j.puhip.2021.100136>
- [53] Zane McNeill. 2024. Democratic Senators Ask DOJ If Facial Recognition Tech Violates Civil Rights Act. *Truthout* (23 Jan. 2024). <https://truthout.org/articles/democratic-senators-ask-doj-if-facial-recognition-tech-violates-civil-rights-act/>.
- [54] Microsoft Teams. 2021. Prepare offices for hybrid work with new experiences on Microsoft Teams Rooms and Teams devices. <https://techcommunity.microsoft.com/blog/microsoftteamsblog/prepare-offices-for-hybrid-work-with-new-experiences-on-microsoft-teams-rooms-an/2709599>
- [55] Muhammad Baqer Mollah, Md. Abul Kalam Azad, and Athanasios Vasilakos. 2017. Security and privacy challenges in mobile cloud computing: Survey and way ahead. *Journal of Network and Computer Applications* 84 (April 2017), 38–54. <https://doi.org/10.1016/j.jnca.2017.02.001>
- [56] Wojciech Mostowski and Pim Vullers. 2012. Efficient U-Prove Implementation for Anonymous Credentials on Smart Cards. In *Security and Privacy in Communication Networks* (London) (*SecureComm 2011*), Muttukrishnan Rajarajan, Fred Piper, Haining Wang, and George Kesidis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 243–260. https://doi.org/10.1007/978-3-642-31909-9_14
- [57] Karen Nershi. 2022. How Strong Are International Standards in Practice? Evidence from Cryptocurrency Transactions. Stanford Cyber Policy Center Fall Seminar Series. <https://www.youtube.com/watch?v=XHD-3lUaVUE>.
- [58] Helen Nissenbaum. 2011. A Contextual Approach to Privacy Online. *Daedalus* 140, 4 (Oct. 2011), 32–48. https://doi.org/10.1162/DAED_a_00113
- [59] Lewis Page. 2009. Facial-recognition tech now used to greet hotel guests. *The Register* (13 Feb. 2009). https://www.theregister.com/2009/02/13/face_ware_hotel_guests.
- [60] Andreas Pfitzmann and Marit Hansen. 2010. A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management. http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf.
- [61] Sundar Pichai. 2018. AI at Google: our principles. <https://blog.google/technology/ai/ai-principles/>
- [62] Carson Powers, Nickolas Gravel, Christopher Pellegrini, Micah Sherr, Michelle L. Mazurek, and Daniel Votipka. 2024. “I can say I’m John Travolta...but I’m not John Travolta”: Investigating the Impact of Changes to Social Media Verification Policies on User Perceptions of Verified Accounts. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*. USENIX Association, Philadelphia, PA, 353–372. <https://www.usenix.org/conference/soups2024/presentation/powers>
- [63] Alan Rappoport and Kashmir Hill. 2022. I.R.S. to End Use of Facial Recognition for Identity Verification. *The New York Times* (7 Feb. 2022). <https://www.nytimes.com/2022/02/07/us/politics/irs-idme-facial-recognition.html>
- [64] Rebecca M. Blank and Patrick D. Gallagher. 2012. Guide for Conducting Risk Assessments. *NIST Special Publication 800-30, Revision 1* (Sept. 2012). <https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-30r1.pdf>.
- [65] Neil Richards and Woodrow Hartzog. 2019. The Pathologies of Digital Consent. *Washington University Law Review* 96 (2019), 1461–1503. <https://papers.ssrn.com/abstract=3370433>.
- [66] Mia Sato. 2021. The pandemic is testing the limits of face recognition. *MIT Technology Review* (28 Sept. 2021). <https://www.technologyreview.com/2021/09/28/1036279/pandemic-unemployment-government-face-recognition/>.
- [67] Benjamin Saunders, Julius Sim, Tom Kingstone, Shula Baker, Jackie Waterfield, Bernadette Bartlam, Heather Burroughs, and Clare Jinks. 2018. Saturation in qualitative research: exploring its conceptualization and operationalization. *Quality & Quantity* 52 (2018), 1893–1907. <https://doi.org/10.1007/s11135-017-0574-8>
- [68] Morgan Klaus Scheuerman, Jialun Aaron Jiang, Casey Fiesler, and Jed R. Brubaker. 2021. A Framework of Severity for Harmful Content Online. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2, Article 368 (Oct. 2021), 33 pages. <https://doi.org/10.1145/3479512>
- [69] Morgan Klaus Scheuerman, Jacob M. Paul, and Jed R. Brubaker. 2019. How Computers See Gender: An Evaluation of Gender Classification in Commercial Facial Analysis Services. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 144 (Nov. 2019), 33 pages. <https://doi.org/10.1145/3359246>
- [70] Ben Schouten and Bart Jacobs. 2009. Biometrics and their use in e-passports. *Image and Vision Computing: Special Issue on Multimodal Biometrics* 27, 3 (Feb. 2009), 305–312. <https://doi.org/10.1016/j.imavis.2008.05.008>
- [71] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N’Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society (AI/ES 2023)*. 723–741. <https://doi.org/10.1145/3600211.3604673>
- [72] Emrys Shoemaker, Gudrun Svava Kristinsdottir, Tanuj Ahuja, Dina Baslan, Bryan Pon, Paul Curriion, Pius Gumisizira, and Nicola Dell. 2019. Identity at the margins: examining refugee experiences with digital identity systems in Lebanon, Jordan, and Uganda. In *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS ’19)*. Association for Computing Machinery, New York, NY, USA, 206–217. <https://doi.org/10.1145/3314344.3332486>
- [73] Manu Sporny, Dave Longley, Markus Sabadello, Drummond Reed, Ori Steele, and Christopher Allen. 2022. Decentralized Identifiers (DIDs) v1.0: Core architecture, data model, and representations. <https://www.w3.org/TR/did-core/>.
- [74] Bernd Carsten Stahl and David Wright. 2018. Ethics and Privacy in AI and Big Data: Implementing Responsible Research and Innovation. *IEEE Security & Privacy* 16, 3 (2018), 26–33. <https://doi.org/10.1109/MSP.2018.2701164>
- [75] Alan Stapelberg. 2025. It’s now easier to prove age and identity with Google Wallet. <https://blog.google/products/google-pay/google-wallet-age-identity-verifications/>
- [76] Bethan Staton. 2016. Eye spy: biometric aid system trials in Jordan. *The New Humanitarian* (18 May 2016). <https://www.thenewhumanitarian.org/analysis/2016/05/18/eye-spy-biometric-aid-system-trials-jordan>.
- [77] Quinten Stokkink, Georgy Ishmaev, Dick Epema, and Johan Pouwelse. 2021. A Truly Self-Sovereign Identity System. In *2021 IEEE 46th Conference on Local Computer Networks (LCN)*. 1–8. <https://doi.org/10.1109/LCN52139.2021.9525011>
- [78] David Temoshok, Diana Proud-Madruga, Yee-Yin Choong, Ryan Galluzzo, Sarbari Gupta, Connie LaSalle, Naomi Lefkowitz, and Andrew Regenscheid. 2025. Digital Identity Guidelines. *NIST Special Publication 800-63-4* (July 2025). <https://doi.org/10.6028/NIST.SP.800-63-4>
- [79] David R. Thomas. 2006. A General Inductive Approach for Analyzing Qualitative Evaluation Data. *American Journal of Evaluation* 27, 2 (2006), 237–246. <https://doi.org/10.1177/1098214005283748>
- [80] Stephen Totilo. 2021. Youth gaming platform Roblox rolls out new age verification system. *Axios* (21 Sept. 2021). <https://www.axios.com/2021/09/21/roblox-age-youth-identity-verification-gaming>.
- [81] Gaby Del Valle. 2024. Don’t ever hand your phone to the cops. *The Verge* (24 Sept. 2024). <https://www.theverge.com/2024/9/24/24252235/police-unlock-phone-password-face-id-apple-wallet-id>.
- [82] Pim Vullers and Gergely Alpár. 2013. Efficient Selective Disclosure on Smart Cards Using Idemix. In *Policies and Research in Identity Management*, Simone Fischer-Hübner, Elisabeth de Leeuw, and Chris Mitchell (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 53–67. https://doi.org/10.1007/978-3-642-37282-7_5
- [83] Fennie Wang and Primavera De Filippi. 2020. Self-sovereign identity in a globalized world: Credentials-based identity systems as a driver for economic inclusion. *Frontiers in Blockchain* 2 (Jan. 2020). <https://doi.org/10.3389/fbloc.2019.00028>
- [84] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L. Mazurek, Manya Sleeper, and Kurt Thomas. 2022. SoK: A Framework for Unifying At-Risk User Research. In *43rd IEEE Symposium on Security and Privacy (Oakland ’22)*. IEEE, 2344–2360. <https://doi.org/10.1109/SP46214.2022.9833643>
- [85] Elizabeth Anne Watkins. 2021. “In the Pre-Evening Light it Sees So Little”: A Human-Centered, Qualitative Investigation into Computer Vision as Account Verification in Gig Work. Princeton CITP Seminar. <https://www.cs.princeton.edu/events/26112>.
- [86] Elizabeth Anne Watkins. 2023. Face Work: A Human-Centered Investigation into Facial Verification in Gig Work. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1, Article 52 (April 2023), 24 pages. <https://doi.org/10.1145/>

3579485

- [87] Laura Weidinger, Maribeth Rauh, Nahema Marchal, Arianna Manzini, Lisa Anne Hendricks, Juan Mateos-Garcia, Stevie Bergman, Jackie Kay, Conor Griffin, Ben Bariach, Iason Gabriel, Verena Rieser, and William Isaac. 2023. Sociotechnical Safety Evaluation of Generative AI Systems. arXiv:2310.11986 [cs.AI]
- [88] Edgar A. Whitley, Uri Gal, and Annemette Kjaergaard. 2014. Who do you think you are? A review of the complex interplay between information systems, identification and identity. *European Journal of Information Systems* 23 (2014), 17–35. <https://doi.org/10.1057/ejis.2013.34>
- [89] Zack Whittaker. 2021. Amazon will pay you \$10 in credit for your palm print biometrics. *TechCrunch* (2 Aug. 2021). <https://techcrunch.com/2021/08/02/amazon-credit-palm-biometrics/>.
- [90] Allison Woodruff. 2019. 10 things you should know about algorithmic fairness. *Interactions* 26, 4 (June 2019), 47–51. <https://doi.org/10.1145/3328489>
- [91] World Bank Group. 2016. Identification for Development: Strategic Framework. <https://thedocs.worldbank.org/en/doc/21571460567481655-0190022016/render/April2016ID4DStrategicRoadmapID4D.pdf>.
- [92] Aram Zucker-Scharff. 2022. The MTA’s switch to OMNY machines is a privacy nightmare. *Fast Company* (11 Sept. 2022). <https://www.fastcompany.com/90788367/the-mtas-switch-to-omny-machines-is-a-privacy-nightmare>.

A Casebook

Table 6: Types of cases in the casebook, ordered alphabetically. Types were created post-hoc for the purpose of presentation, and are not meant to be mutually exclusive or comprehensive of all digital identification systems.

Type	Description	Example
Age Verification	User provides proof of age (e.g., a copy of their government ID or a privacy-preserving certificate [9, 49]) to access online content or services with a minimum age requirement (e.g., adult entertainment or gambling)	<i>In order to access a gaming platform’s features such as voice chat, the user uploads a copy of their driver’s license, passport, or photo ID, and then uses selfie verification to prove they are the owner of that document [18, 80]</i>
Criminal Investigation	Law enforcement uses digital means to identify, investigate, and locate known criminals or people of interest	<i>Prison authorities gather incarcerated people’s voice prints to create a database of voice prints; this database is used to identify inmates in phone calls, and in some cases, the system analyzes the voices of call recipients outside prisons to identify regular contacts [22, 43]</i>
Device and Account Authentication	User provides identifying information, often biometric, to access a device or account [20]; this information may be used for primary authentication or for two-factor authentication (2FA)	<i>Mobile device uses facial recognition technology to authenticate user access to phone [30]</i>
Enterprise Collaboration Tools	System identifies employees, for example, meeting participants, in order to facilitate collaboration	<i>Virtual meeting environment recognizes people in a video conference in order to add their profile information and name to the meeting [54]</i>
Fraud Detection and Prevention	System actively or passively gathers information to confirm that identity is authentic and has not been misappropriated	<i>Uber app conducts periodic facial recognition checks to verify that a driver’s face matches the ID associated with the account, in order to minimize theft of legitimate drivers’ accounts or unauthorized account sharing [39, 85, 86]</i>
Financial Services	System monitors financial transactions and gathers identity information in order to meet compliance and legal requirements, and to detect issues such as large transactions that may indicate money laundering	<i>System gathers digital identity information to meet mandatory reporting requirements for legal identity associated with financial accounts and transactions [57]</i>
Government-Issued ID or Credentials	Government or non-government entity provides digital versions of government-issued ID or credentials, which user can present in-person and/or online; the provisioning process may include verifying the user’s legal identity	<i>The user verifies their identity and presents their government-issued driver’s license on a mobile device [16, 75]</i>
Government Services	System is used to access government services, such as benefits, based on technologically facilitated in-person or online identification	<i>An identity verification platform automatically verifies digital identity for access to unemployment applications [66] or tax services [63]</i>
Government Surveillance	Government uses digital means, e.g., user check-ins, to surveil a population, e.g., to ensure compliance with specified behaviors or location restrictions	<i>Person provides a selfie to prove they are in a given location [66]</i>

Table 6 (continued from previous page): Types of cases in the casebook.

Type	Description	Example
Health Certificates and Public Health	User presents health certificates to meet public health requirements or recommendations	<i>User can show digital proof of vaccination status [2, 7, 52]</i>
Homeless Identification	System infers legal identity when documentation is not available	<i>System uses facial recognition to enable law enforcement, fire departments, and other first responders to identify homeless people, and then provides information such as date of birth, behavioral labels, or criminal history [19]</i>
Photo Labeling in Consumer Tools	System can match and label faces detected in user’s photos	<i>Photo tool on mobile device recognizes faces to facilitate labeling</i>
Physical Security	System identifies people to support physical security in the workplace and at home	<i>System attempts to detect familiar faces at the door and notifies home owner if it detects people it does not recognize [37]</i>
Point-of-Purchase	User presents digital identification such as biometric information at point-of-purchase to acquire goods or services	<i>Palm print scanners are used for store checkout [14, 47, 89], or facial recognition is used for school lunch payment [6]</i>
Refugee Services	System identifies people in order to provide refugee services [72]	<i>Refugees pay for groceries and receive cash with an iris scan [44, 46, 76]</i>
Urban Planning	System monitors and tracks people’s movement in public spaces to facilitate urban planning	<i>City deploys surveillance cameras and purchases cellphone data to monitor tourists and prevent overcrowding [15]</i>