Anthroencryption

Strategies for Protecting the World's Most Vulnerable People

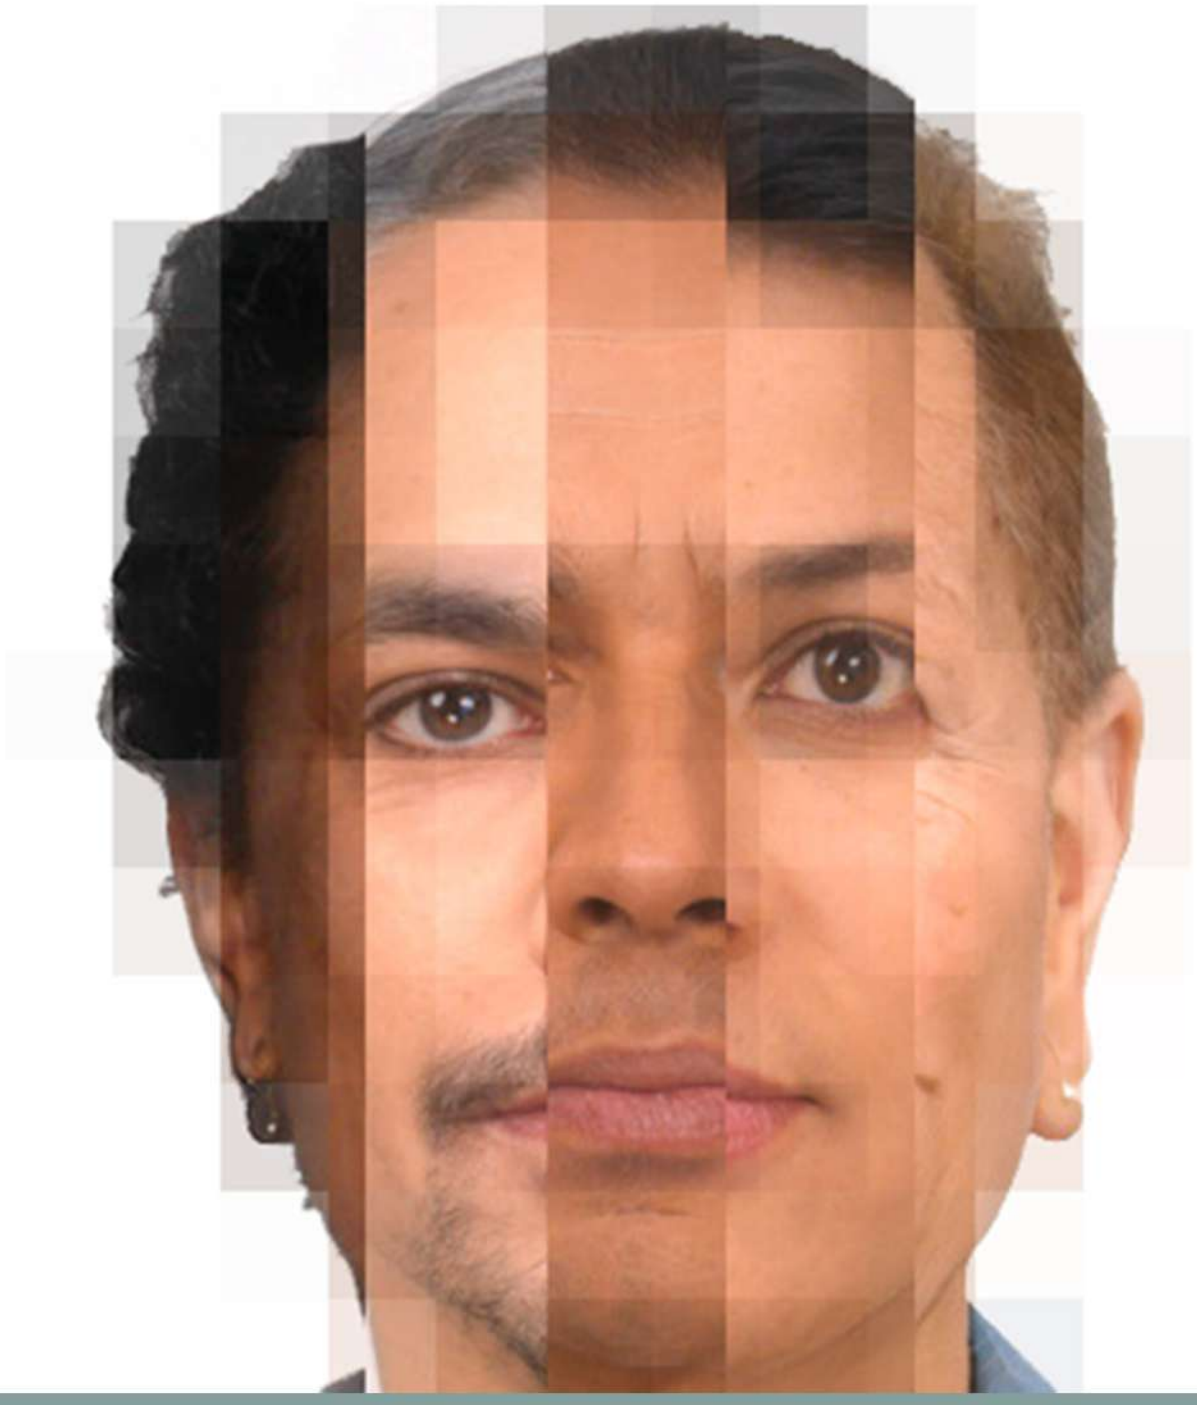

Timothy Quinn

DARKDATA project

Anthroencryption

Strategies for Protecting the World's Most Vulnerable People

Timothy Quinn
Dark Data Project
darkdataproject.org

v 2022.09

Introduction

Governments, humanitarian organizations and private contractors are capturing, storing and sharing an ever increasing volume of identity data, much of it pertaining to "third party" or "end recipient" individuals such as refugees and other vulnerable individuals who may never interact directly with the databases where their data is stored. While many countries have enacted legislation to establish baseline safety protocols for protecting personal data, and online privacy and security standards like GDPR and OWASP are becoming more ubiquitous outside of the technology economy, protecting data about people remains a challenging and underappreciated area of risk for most organizations, and securely sharing that data a significant impediment.

For individuals who don't possess an email address, driver's license, credit card or passport, or who engage with government or NGO services through offline activities, these problems are particularly impactful, whether imposing unnecessary and unwanted authentication like biometric disclosure, or inhibiting the safe sharing of data among humanitarian partners, support services and judicial systems.

A better understanding of the tiered reality of identity protection is critical if we want to establish appropriate strategies suitable to the needs of vulnerable people in a wide variety of contexts. At present, this subset of data security is loosely referred to as "anonymization," a term easily criticized for a blasé "one size fits all" approach to identity protection. A more accurate and nuanced term is anthroencryption: the practice of protecting and safely sharing human identity.

an·thro·en·cryp·tion

n. the practice of protecting and safely sharing human identity

Shareability is a key component of anthroencryption, and informs the type and degree of identity protection an organization should choose for a specific use case. If there was no need to share identity data, organizations could in principle treat personally identifiable data no differently than any other critically sensitive data and protect it using the strongest possible industrial encryption technologies, or secure it on physical media or on air-gapped storage devices. This remains a common approach in evidentiary use cases, as when governments, law

enforcement or transnational justice organizations like the International Criminal Court need to protect the identities of whistleblowers, confidential sources and witnesses of mass atrocities and war crimes. The cost of unshared identity data, however, is also high: refugees are compelled to recount traumatic events multiple times to multiple humanitarian organizations, recipients of emergency aid are required to reregister with multiple support services, and information which should be in the public domain can wind up locked away for decades out of an excess of caution. Anthroencryption provides a better framework for deciding what data should be collected and how it should be shared.

An important feature of anthroencryption is the principle of conservative implementation: data caretakers should capture, store and share the least amount of identity data necessary for a specific use case. Fingerprints should not be required to receive emergency aid, and the lack of a telephone number or postal address should not be a barrier to civic or economic participation. At the same time, anthroencryption is intended to be forward-looking, encouraging a fulsome collection of identity data if there is additional benefit to end recipients when safely shared between humanitarian organizations. An example of such a "one-for-many" benefit would be an immigration process in which legal aides capture sensitive data which is then used to provision multiple services.

Although an understanding of the advantages and disadvantages of each type of anthroencryption primarily benefits vulnerable populations unable to advocate for themselves, the framework is broadly applicable to all identity data, even that which is captured by private entities for commercial purposes. A more nuanced perspective on identity protection ultimately benefits anyone whose endemic, immutable or private attributes will be stored in databases they neither see nor control.

Types of Anthroencryption

1 - Unidentification

The most basic form of anthroencryption is also the easiest to implement: if there's marginal need or benefit to capturing personally identifiable data, don't collect it. A remarkably broad range of humanitarian use cases fall into this category.

Community drop-in centers, food banks, healthcare clinics, homeless shelters and other brief transactional social services generally have a need for incremental data (how many people are

using a service over a given unit of time) or statistical data (what services are being used most, and by which socioeconomic groups), but not identity data. Many such organizations nonetheless collect personally identifiable data, sometimes to prevent duplication of services, sometimes to identify hoarding, profiteering and corruption, and sometimes to satisfy the requests of funders and other partners attempting to make data-driven decisions about efficacy.

Afghan elder shows his inked finger during elections in Barge Matal, Afghanistan, 2009. Photo courtesy Christopher Allison, United States Army.

This sort of data trawling places individuals at unnecessary risk by forcing them to self-police systems that can usually be monitored in other ways. Concerns have been raised in recent years, for example, about the collection and storage of biometric data by emergency aid organizations as a strategy to ensure equal access to resources while inhibiting crisis profiteering, even though other less invasive strategies, such as inked fingertips, have been enormously effective at preventing election fraud with relatively little abuse. For end recipients, it's not hard to imagine the potential "worst case" risks of unnecessary identification. Landlords might be disinclined to rent to those who have experienced previous periods of homelessness or joblessness. Insurers would certainly increase premiums if made aware of specific prior medical conditions (or even just tests for certain medical conditions).

In situations where the risk to end recipients far outweighs any benefit, a better strategy is to simply count. Many humanitarian organizations rely on passive counting, which is the counting of consumed or unconsumed services such as used syringes or remaining MREs, since this can be accomplished without the real-time involvement of beneficiaries while prioritizing ease of access and dignity over precision. An alternative approach is active counting, which is the real-time monitoring of consumption of services, and is accomplished by positioning a counter at the front of a queue or handing out incrementing tickets which are then redeemed for services.

When deciding whether to pursue a strategy of unidentification, it's important to honestly assess risks and benefits. It's easy to overestimate the impact of abuse resulting from too little

data. While inked fingertips can be washed clean, and aid recipients can blend in with a crowd to be served more than once, these abuses have relatively minor impact compared with the myriad disadvantages of over-policing an already marginalized population. Hoarding and profiteering certainly shouldn't be trivialized, since they can have severe social and economic impacts on the people humanitarian organizations are meant to be helping, but in many cases there are better mitigation strategies which don't involve placing the burden of proof on victims.

2 - Aggregation

Many non-individualized datasets nonetheless incorporate identification at the point of capture because accuracy and non-duplication are critical requirements of the data. Elections and censuses in many countries (specifically, countries where there are government ID cards, chains of custody for physical ballots, and other requisite elements of statistical infrastructure) fall into this category: in these cases, it's critical to ensure the integrity of data capture, even though only aggregated data is ever shared.

Governments and humanitarian organizations embarking on aggregated anthroencryption have a choice between attributable aggregation and destructive aggregation. Destructive aggregation requires that identity data be destroyed following aggregation, and is preferable from a privacy perspective because it prevents future unintended misuse, such as that resulting from a security breach. This risk shouldn't be underestimated or equated with the risk of other long-term encryption strategies; long-term storage of unused data poses a particular risk because it's often forgotten in security audits and transfers of responsibility, and is less likely to benefit from organization-wide security upgrades, particularly if the media on which it's stored becomes deprecated (for example, data stored on CDR).

Most aggregation use cases require attributable aggregation, because retaining identity from the point of capture facilitates auditability: in other words, the data needs to be available if the aggregation is ever questioned, and that requirement may persist well beyond any initial data sharing period. Election results are an example of aggregated data which requires both short-term and long-term auditability.

Regardless of whether identity data is retained or destroyed, there are risks with sharing even

eth·no·en·cryp·tion

aggregated data. If a dataset is too small, it can be easy to discern unique indicators of individuals. Polls and surveys are particularly vulnerable to this sort of accidental disclosure; if a dataset is based on a small number of respondents, and that data can be further winnowed down by attributes like location or demographics, it's not impossible to wind up with a dataset representing a single individual.

n. the practice of obfuscating ethnicity to protect populations from discrimination or violence

Many organizations attempt to guard against this by setting a minimum sample size for disclosure and not sharing further reductive attributes. Obfuscating attributes such as ethnicity (ethnoencryption) reduces the risk that individuals can be identified as a result of targeting a specific population, such as occurred during the Rwandan genocide of 1994 when government forces both encouraged, and participated in, a pogrom to systematically identify and kill ethnic Tutsis (a process which was greatly facilitated by the presence of ethnicity data on government-issued ID cards).

Another risk with anthroencryption, and aggregation in particular, is correlation inference, which is more popularly known as "de-anonymization," "re-identification" or the "Mosaic Effect." When a dataset is only partially obfuscated, the remaining data can be cross-indexed with other large datasets (such as voter registration or credit assessment databases) to fill in missing areas, deobfuscating the data. Characteristics such as birthdays, gender and ethnic self-identification have historically been vulnerable to this sort of data reconstruction because they represent endemic qualities of individuals which are unlikely to change over time or across datasets. This innate vulnerability is also the reason why mandating immutable characteristics such as "mother's maiden name" as a form of identity verification is now considered a terrible security practice.

Unobfuscated, or even partially obfuscated, geolocation data is also a prime candidate for correlation inference. Where you are says a lot about who you are; when traveling between home and work with geolocation enabled by default on your phone, you unknowingly risk exposing your identity through correlation with public residency and employment records. When your travels take you away from home and work, the uniqueness of your geolocation data can also represent a risk. Notoriously, US military personnel stationed in Africa in recent years

have been identified from the public data generated by GPS-enabled fitness trackers. It's not hard to imagine similar scenarios where the locations of humanitarian workers or journalists are identified based on the unique footprint of technology that isn't ubiquitous in the region.

Even approximate location data can pose a risk of correlation inference. Postal codes and telephone area codes are often used as proxies for location, but because of variances in population density can make individuals in rural and remote areas more vulnerable than individuals in dense urban areas. Similarly, IP addresses and IMEI numbers are often used to identify unique network-connected devices without explicitly identifying the owners of those devices, but, like any representational identifiers, can be correlated with other datasets (such as those maintained by telecommunications companies) to infer ownership identity. IP addresses are potentially more vulnerable than IMEI addresses because of how frequently they're captured and stored to augment user authentication and inform usage analytics.

The best way to mitigate the risks of data aggregation is to ensure large, non-reductive sample sizes, avoiding intersective data points which can be correlated with other datasets. Obviously, if these reductive datapoints are the essential purpose of gathering the data, such as demographic data resulting from a census or ethnicity data used to mitigate conflict, removing them won't be an option, and the only recourse will be to increase the size of the underlying source data.

3 - Irreversible Encryption

An often overlooked form of anthroencryption is irreversible one-to-one encryption, whereby personally identifiable data is transformed into some sort of unique representation of the source data. The notion of one-way encryption will be familiar to security and database administrators as "hashing," a technique commonly employed to encrypt a password in a database: the password is run through an algorithm that transforms it into a seemingly random string of characters and then saved in the database so that the next time the user enters that password, the two hashes can be compared without the original password ever having been retained. Data like this which is deliberately encrypted prior to storage, whether reversibly or irreversibly, is referred to as being encrypted "at rest" because the data can be maintained indefinitely in an encrypted state.

The purpose of irreversible identity encryption is to create a reasonably unique "fingerprint" of the individual so that it can be distinguished from other identities, even while it remains encrypted and forever unreadable. (Note that this is a metaphorical fingerprint, not a literal fingerprint, which would obviously defeat the purpose of encryption.) This notional fingerprint will be only reasonably unique because the encryption process will likely be "lossy": it will reduce information density in the interest of manageability. That said, a suitable algorithm working with enough source data can create a fingerprint that is unique to such a degree than any practical risk of duplication is non-existent.

To understand how irreversible anthroencryption would work, imagine taking the personally identifiable portion of an identity (name, date of birth, place of birth, driver's license number, passport number, etc.) and concatenating everything into a single string of characters. This string of characters is then "salted" (meaning that a domain-specific identifier is appended to ensure that any hashes resulting from the current system don't exactly resemble hashes in other systems) and encrypted using a cryptographically secure algorithm. The resulting hash is unique to the encrypted identity and can be used as an identifier to label any non-personally identifiable information. Because the hash has been created using one-way encryption and the source data is then destroyed, there is no chance of future disclosure from a data breach.

(It should be noted that while there exist today cryptographically secure algorithms that are suitable for this task, it's impossible to guarantee that today's unbreakable encryption won't become vulnerable in the future as a result of innovations in computer processing power. Advances in quantum computing, for example, may eventually accomplish brute force decryption which is impossible today.)

The advantages of irreversible encryption are that it's extremely secure and it preserves more source data than aggregation. The primary disadvantage is that while identity data can be encrypted into a unique fingerprint, a lengthy string of characters in lieu of a name may be unwieldy for governments and humanitarian organizations, and therefore subject to typographic error (a risk which can be somewhat mitigated by using a check digit as the final character, as with an ISBN or UPC code).

Founded in 2009 and a recent partner of the Dark Data Project (2022), the Rural Development Network (ruraldevelopment.ca) is a Canadian NGO which facilitates sustainability and investment in rural communities by addressing social issues that are unique to, or exacerbated by, reduced capacity and the reduced visibility of remote commun5. A core focus of RDN is the challenge of estimating levels of rural homelessness, for which they have created a methodology that involves broad-based data collection over a time-delimited period and/or repurposing of data through parallel social services such as food banks. To ensure the anonymity of end recipients while reducing data duplication, RDN uses a "unique identifier" (essentially a hash) which concatenates portions of an end recipient's first name, last name, year of birth, self-identifying gender, and a sum of the integers in the end recipient's day of birth.

J	O	S	M	0	6	6	4	M
A1	A2	B1	B2	C1	C2	D1	D2	E

(A1, A2)	*(B1, B2)*	*(C1 + C2)*	*(D1, D2)*	*(E)*
First 2 letters of first name	*First two letters of last name*	*Sum of the numbers in birth date*	*Last two numbers in year of birth*	*M for male, F for female, or X for non-binary*

The advantage of this sort of human-encodable hash is that it can be created in real-time with pen and paper, requires almost no math, and is immediately sortable to prevent data duplication prior to further data capture. It also adheres to the principle of conservative implementation by being just unique enough to

accomplish its dual goals of anonymity and uniqueness; a more complex algorithm would require automation and provide little additional benefit.

4 - Reversible Encryption

The tradeoffs inherent in the above forms of anthroencryption will be unpalatable for many use cases. If identity data needs to be captured, stored and retrieved in full and in perpetuity, reversible encryption is the next logical approach.

Think of reversible encryption as a vault. You can place identity data in the vault (perhaps literally, if you've chosen to resort to hard copies as a strategy against electronic intrusion), you can lock the vault using some kind of key, and then when you need to retrieve that data, you can use that same key (or, if you've chosen to employ "public key" cryptography, a slightly different key) to reopen the vault and retrieve the original data.

That key can be a physical key, or a password, or a biometric attribute like a fingerprint, retina scan or voiceprint.

The risks of reversible encryption are myriad: encryption algorithms may be insufficiently secure to resist attack either now or in the future, physical keys can be lost or duplicated, passwords may be intercepted or deduced by a variety of means, biometrics can be unreliable (as with the fingerprints of manual laborers), impractical (particularly for smaller humanitarian organizations) or forged (as with audio deepfakes), and source data may be retained forever, representing an unending risk. Databases of encrypted data are routinely breached, which is disconcerting, but the good news is that many of these breaches are the result of preventable lapses in security posture, so with sufficient skill and forethought a database of encrypted identity data ought to be safe from most threats.

Another disadvantage to the "locked vault" approach, however, is shareability. Unlike other forms of anthroencryption, sharing reversibly encrypted full identity data rather than aggregated, hashed or otherwise minimized or obfuscated data likely involves sharing keys or

creating and managing multiple keys. While keys can be imbued with greater or lesser degrees of authority (a security strategy commonly known as "access control"), every key is a potential threat if it falls into the wrong hands. To mitigate this risk, many organizations prudently limit their keys to the absolute minimum, even though this reduces the utility of data which would otherwise be of benefit to other organizations, and to end recipients. An example of this sort of lost opportunity is witness data: statements taken from victims of crime or mass atrocities which are deemed too sensitive to share in any form because of the potential harm resulting from accidental identity disclosure.

The above risks notwithstanding, it's important to note that there are use cases of reversible encryption which aren't as vulnerable because the underlying data isn't as sensitive. An example of such a use case is the protection of gender or ethnicity data from job applicants in an effort to make hiring practices more equitable. While it would be counterproductive for this data to be compromised, the impacts of disclosure are less severe because no existential (physical) harm will come to individual candidates whose gender or ethnicity is unintentionally revealed.

5 - Deniable Encryption

An overview of anthroencryption strategies would be incomplete without touching on an area of research at the Dark Data Project which further mitigates some of the risks of reversible encryption.

One of the intrinsic weaknesses of most forms of encryption is that the output is clearly intentionally encrypted, and is therefore a target for those who wish to acquire the source data. A locked vault exists to protect things of presumed value.

The solution to this problem is to employ invisible or "deniable" encryption, a strategy whereby data is securely encrypted in such a way as to resemble unencrypted data. An example of deniable encryption would be to create an algorithm that transforms John Smith, born 1-Jan-1959, into Mary Jones, born 31-Dec-1960. The latter data can be saved in a database and freely shared because it presents the illusion of being unprotected and/or of no value. (More strategically, it can be saved with a weaker form of encryption, since the risk has been removed that if the secretly encrypted data is disclosed the actual source data will be in jeopardy.)

The rudimentary word substitution cyphers used by drug dealers to optimistically defeat electronic monitoring is a form of deniable encryption, as were the invisible inks used for espionage during World War II to embed real messages beneath false messages. More recently, subtly coded political language (colloquially known as "dog whistles") have been used by organizers of extremist movements in the United States and elsewhere as a form of deniable encryption, conveying hidden ideology using language which might otherwise appear innocuous.

The challenge of deniable identity encryption is that it requires either a particularly sophisticated algorithm to encode and decode a wide variety of personally identifiable datapoints in such a way as to make an encrypted identity appear plausibly unencrypted, or a secondary "lookup"database in which Identity A is linked to Identity B, which itself can then become a point of vulnerability. An ideal implementation would be to employ strong reversible encryption on the lookup database and then use deniable encryption for sharing data, which removes the need to share keys.

Summary

- Identity data is both more nuanced and more sensitive than most other forms of data employed by governments and humanitarian organizations, and cannot be treated as just part of an organization's overall online security process.
- Different levels of identity security are appropriate to different use cases. What differentiates levels of identity security is both the type of data stored and the degree to which it can be shared without jeopardy.
- At all times, data caretakers should be guided by the principle of conservative implementation.
- Data caretakers must also be mindful of the risks of correlation inference, particularly in contexts where state actors have access to a large volume of data.
- Even the most secure forms of encryption today are vulnerable to future brute force attacks as computer processing technologies evolve.