CS2080: Applied Privacy for Data Science
Contextual Integrity Meets Differential Privacy

School of Engineering & Applied Sciences
Harvard University

March 3, 2025

Lecture co-developed with Alice Huang, EmbeddedEthiCS postdoctoral fellow & some content drawn from 2022's EmbeddedEthiCS lecture (by Sophie Gibert)

Housekeeping

1. Based on the poll during last class, we have moved HW deadlines to Fridays 11:59pm. This new deadline applies to HW5 (due this Friday).

2. Please do filling out the mid-semester feedback form (see Henry’s post on Ed).

Examples of potential privacy violations

A hacker breaches a healthcare system's databases and gains access to millions of people's personal data

A video-conferencing software tracks how much attention meeting attendees are paying and sends this information to the meeting host

An employee feels that his boss shouldn't have been able to learn his political leaning based on a recent study

Examples of potential privacy violations

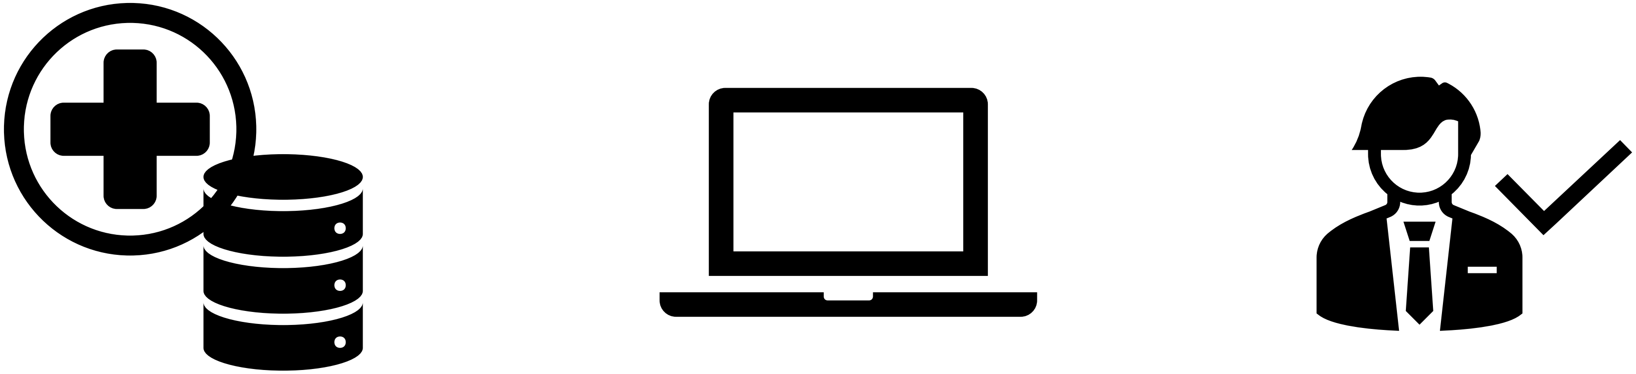

DP is only applicable in scenarios where we are analyzing a population or dataset.

Even when DP *can* be applied, it may not always satisfy our social understanding of privacy.

Contextual Integrity can help us identify when DP is appropriate given a social context.

Considering **context** can help us more formally reason about privacy in a social sense and identify when DP is an appropriate tool.

Can you think of a situation where your or someone else’s privacy was violated?

Can you think of a situation where your or someone else’s privacy was violated?

When an information flow was **inappropriate**
(i.e., contextual integrity was violated)

Nissenbaum 2009

Can you t here your
or someor s violated?

Wh was

(i.e., co violated)

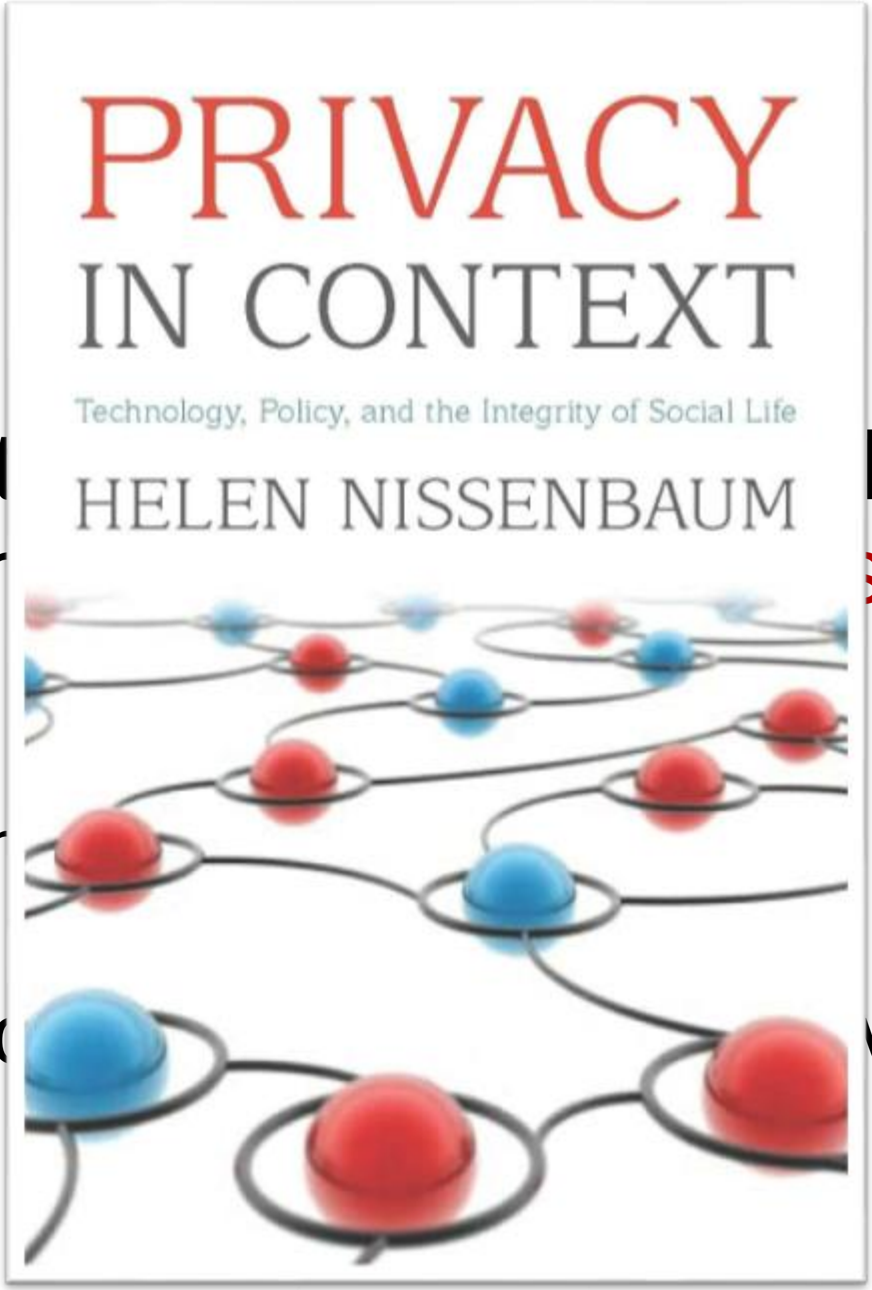

PRIVACY
IN CONTEXT

Technology, Policy, and the Integrity of Social Life

HELEN NISSENBAUM

Nissenbaum 2009

Building blocks of Contextual Integrity

1. **Contexts: “**structured social settings characterized by by canonical activities, roles, relationships, power structures, norms (or rules), and internal values”, for example education or healthcare

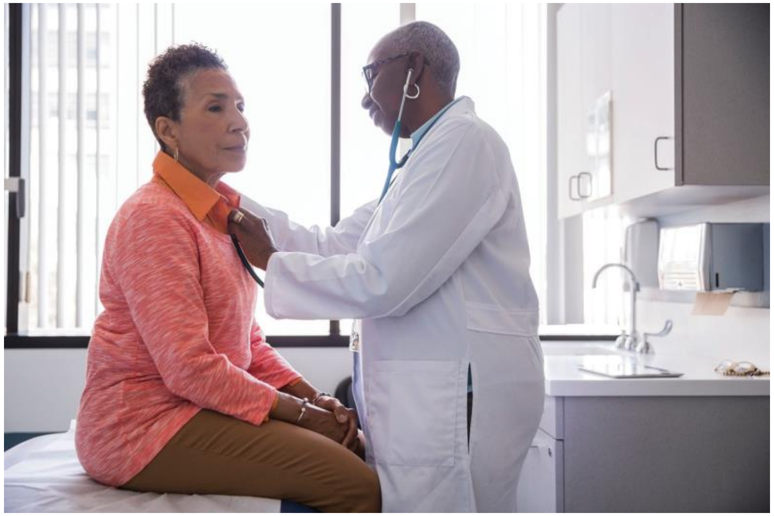

Images: The Crimson, Harvard Health

Building blocks of Contextual Integrity

1. **Contexts: “**structured social settings characterized by by canonical activities, roles, relationships, power structures, norms (or rules), and internal values”, for example education or healthcare
2. **Information norms:** context-specific rules or expectations that prescribe and proscribe certain actions and practices; expressed as a relation between five parameters

Building blocks of Contextual Integrity

1. **Contexts: “**structured social settings characterized by by canonical activities, roles, relationships, power structures, norms (or rules), and internal values”, for example education or healthcare
2. **Information norms:** context-specific rules or expectations that prescribe and proscribe certain actions and practices; expressed as a relation between five parameters

A university professor can share the grades of Alice
with her parent with Alice’s consent.

Residents of the U.S. must provide responses to the
census, under conditions of confidentiality.

A priest should not share a congregant’s confessions
with other people under any circumstances.

Five parameters of an information norm

1. Sender

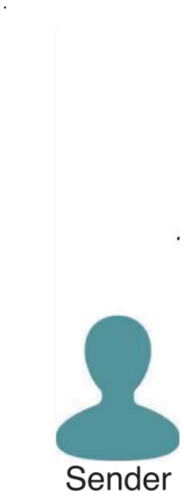

Diagram adapted from Malkin 2023

Five parameters of an information norm

1. Sender
2. Recipient

Recipient
Sender

Diagram adapted from Malkin 2023

Five parameters of an information norm

1. Sender
2. Recipient
3. Subject: Who the information is about

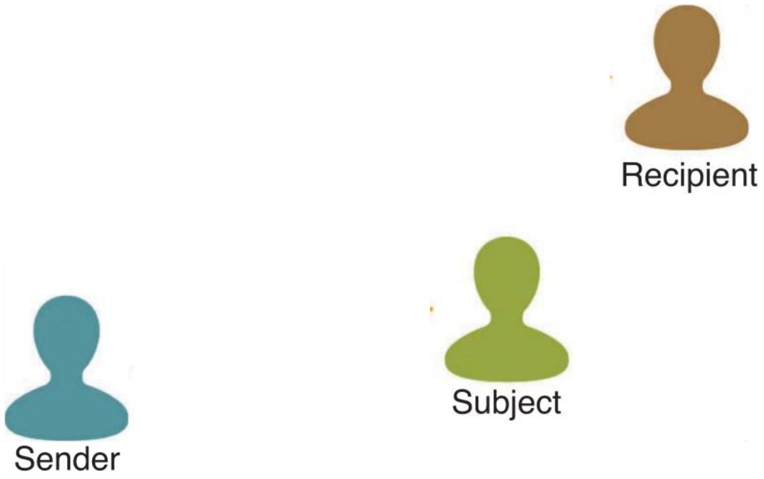

Diagram adapted from Malkin 2023

Five parameters of an information norm

1. Sender
2. Recipient
3. Subject: Who the information is about
4. Information type: Medical? Financial?

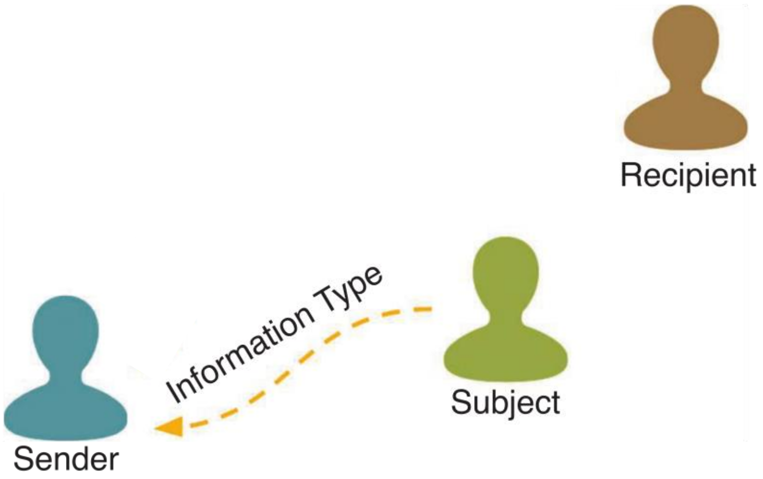

Diagram adapted from Malkin 2023

Five parameters of an information norm

1. Sender
2. Recipient
3. Subject: Who the information is about
4. Information type: Medical? Financial?
5. Transmission principles: Expectations or constraints governing how information flows, such as confidentiality agreements, informed consent, or legal obligations

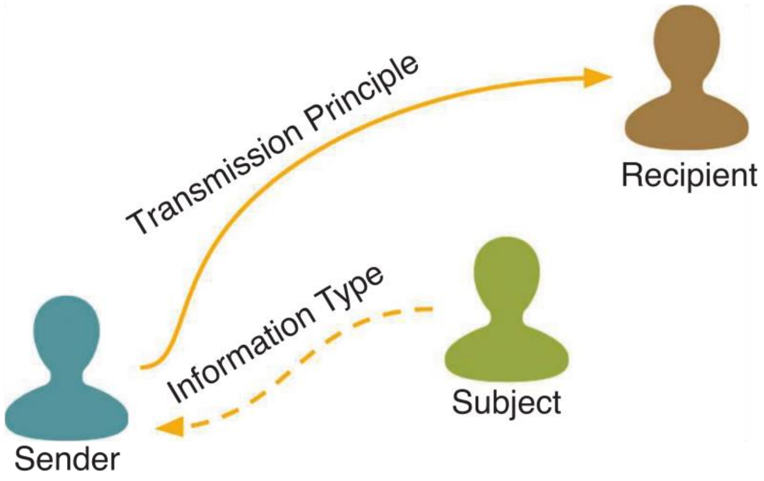

Diagram adapted from Malkin 2023

Try identifying the five parameters

A university professor can share the grades of Alice with her parent with Alice’s consent.

Residents of the U.S. must provide responses to the Census Bureau, under conditions of confidentiality.

A priest should not share a congregant’s confessions with other people under any circumstances.

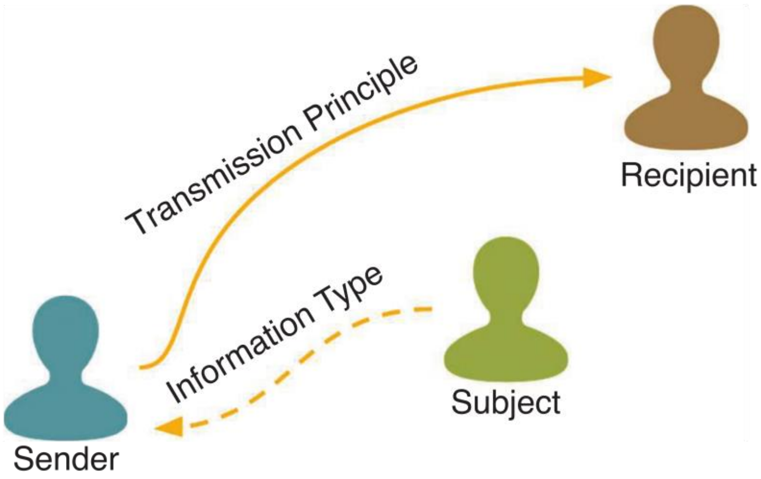

Try identifying the five parameters

A university professor can share the grades of Alice with her parent with Alice’s consent.

Residents of the U.S. must provide responses to the Census Bureau, under conditions of confidentiality.

A priest should not share a congregant’s confessions with other people under any circumstances.

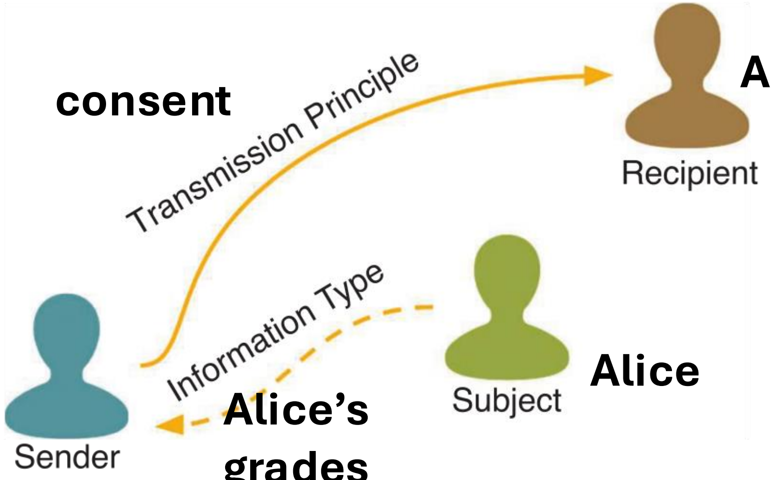

consent

Alice’s parent

Alice

University professor

Alice’s grades

Try identifying the five parameters

A university professor should share the grades of Alice with her parent with Alice's consent.

Residents of the U.S. must provide responses to the Census Bureau, under conditions of confidentiality.

A priest should not share a congregant's confessions with other people under any circumstances.

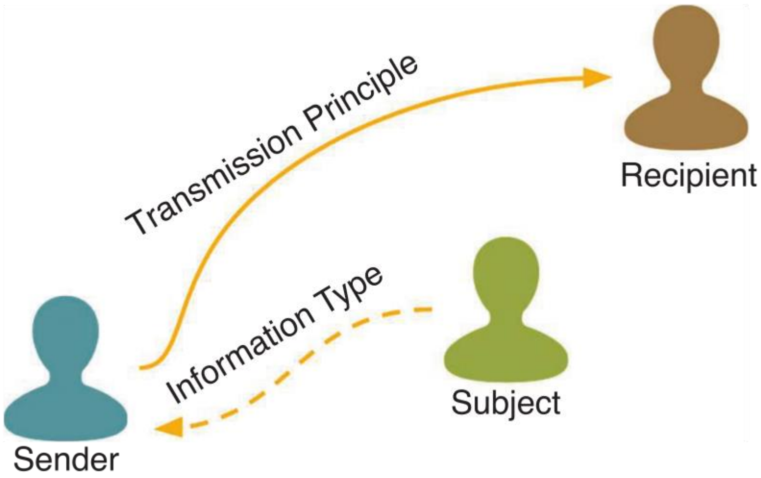

Try identifying the five parameters

A university professor should share the grades of Alice with her parent with Alice’s consent.

Residents of the U.S. must provide responses to the Census Bureau, under conditions of confidentiality.

A priest should not share a congregant’s confessions with other people under any circumstances.

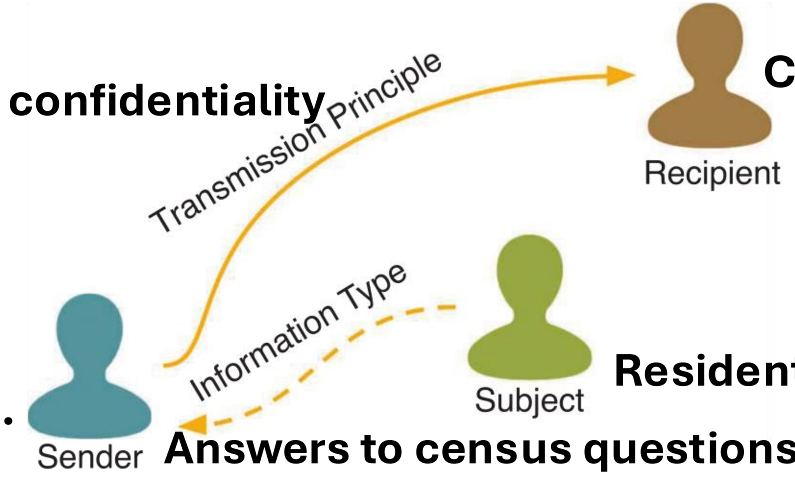

confidentiality

Census Bureau

Residents of the U.S.

Residents of the U.S.

Answers to census questions (e.g., demographics)

Try identifying the five parameters

A university professor should share the grades of Alice with her parent with Alice's consent.

Residents of the U.S. must provide responses to the Census Bureau, under conditions of confidentiality.

A priest should not share a congregant's confessions with other people under any circumstances.

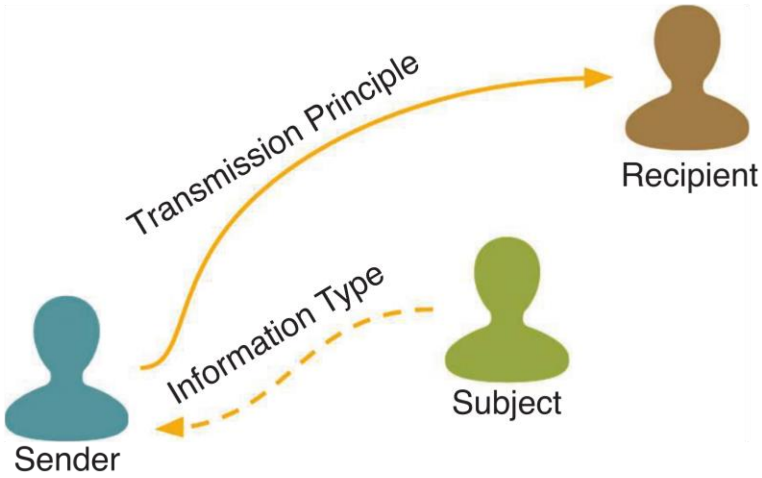

Try identifying the five parameters

A university professor should share the grades of Alice with her parent with Alice's consent.

Residents of the U.S. must provide responses to the Census Bureau, under conditions of confidentiality.

A priest should not share a congregant's confessions with other people under any circumstances.

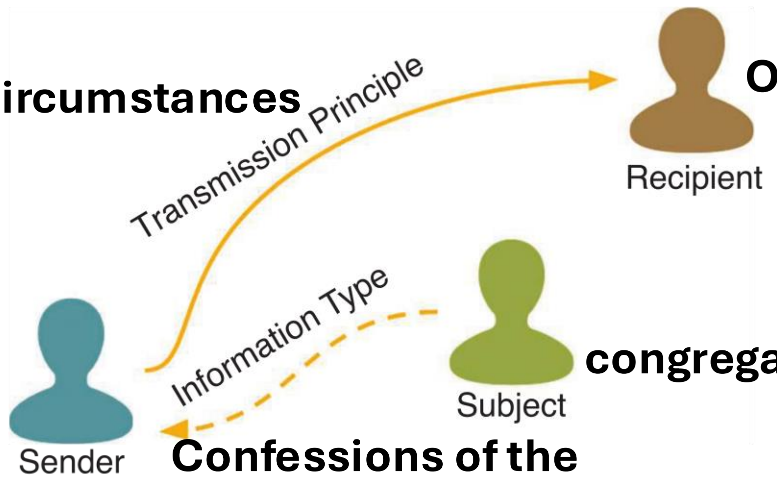

Under no circumstances

Other people

congregant

Priest

Confessions of the congregant

Building blocks of Contextual Integrity

1. **Contexts: "**structured social settings characterized by by canonical activities, roles, relationships, power structures, norms (or rules), and internal values", for example education or healthcare
2. **Information norms:** context-specific standards that prescribe and proscribe certain actions and practices; expressed as a relation between five parameters

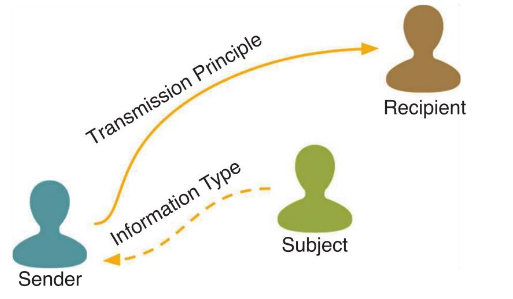

3. **Contextual values:** goals, purposes, or ends of the context; "objectives around which a context is oriented" --- for example to prepare students for the workforce or to provide quality medical care

What is privacy under contextual integrity?

Some information flows are appropriate to a given context, and some are not

- What makes a flow appropriate to a context is (in part) that it supports the attainment of the context-specific values and goals

What is privacy under contextual integrity?

Some information flows are appropriate to a given context, and some are not

- What makes a flow appropriate to a context is (in part) that it supports the attainment of the context-specific values and goals

Privacy is the context-appropriate flow of information

- **Privacy is violated when and because there is a disruption in the context-appropriate information norm** --- i.e., a disruption in how information *should* flow in the context
- A technology or information practice raises *legitimate* privacy concerns when and because it disrupts the context-appropriate information norm --- i.e., it causes information to flow in a way that it *shouldn't* in the context

Using Contextual Integrity to assess information flows

1. Establish the relevant social context and its goals (look for similarity in function between a seemingly new context and its more familiar counterpart)

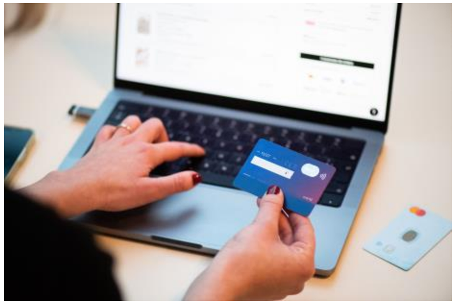

Online shopping is an extension of shopping in person

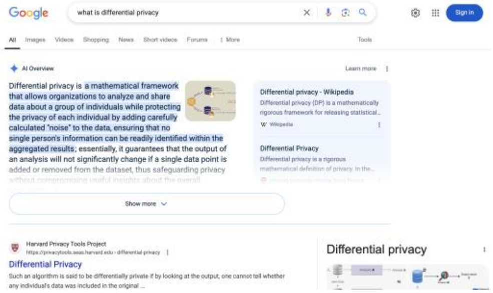

Differential privacy is a mathematical framework that allows organizations to analyze and share data about a group of individuals while protecting the privacy of each individual by adding carefully calculated "noise" to the data, ensuring that no single person's information can be readily identified within the aggregated results; essentially, it guarantees that the output of an analysis will not significantly change if a single data point is added or removed from the dataset, thus safeguarding privacy without compromising useful insights about the overall

Differential privacy - Wikipedia

Differential privacy (DP) is a mathematically rigorous framework for releasing statistical...

Wikipedia

Differential Privacy

Differential privacy is a rigorous mathematical definition of privacy. In the...

Harvard Privacy Tools Project

https://privacytools.seas.harvard.edu › differential-privacy

Differential Privacy

Such an algorithm is said to be differentially private if by looking at the output, one cannot tell whether any individual's data was included in the original ...

Searching the web is an extension of visiting a library

Talking with conversational AI is an extension of ___?

Images: Wikipedia, Business Insider

Using Contextual Integrity to assess information flows

1. Establish the relevant social context and its goals
2. Identify information norms in the context by specifying the five parameters (sender, subject, recipient, information type, transmission principle)

Using Contextual Integrity to assess information flows

1. Establish the relevant social context and its goals
2. Identify information norms in the context by specifying the five parameters (sender, subject, recipient, information type, transmission principle)
3. Identify ways in which the technology or new practice disrupts that information norm

Using Contextual Integrity to assess information flows

1. Establish the relevant social context and its goals
2. Identify information norms in the context by specifying the five parameters (sender, subject, recipient, information type, transmission principle)
3. Identify ways in which the technology or new practice disrupts that information norm
4. Evaluate whether the disruptions identified in Step 3 undermine the goals of the context

Using Contextual Integrity to assess information flows

1. Establish the relevant social context and its goals
2. Identify information norms in the context by specifying the five parameters (sender, subject, recipient, information type, transmission principle)
3. Identify ways in which the technology or new practice disrupts that information norm
4. Evaluate whether the disruptions identified in Step 3 undermine the goals of the context
5. Make a judgment about whether an information practice should be abandoned or changed

Example: Medical referral

Context: healthcare

Example: Medical referral

Context: healthcare

Contextual purposes: Improve patient health, efficient division of labor, fair and wide access to healthcare, transparency to patients, etc.

Example: Medical referral

Context: healthcare

Contextual purposes: Improve patient health, efficient division of labor, fair and wide access to healthcare, transparency to patients, etc.

Information norm: The patient’s medical record should flow from the family doctor to the specialist doctor as long as the patient consents to this sharing, and their records are not shared beyond the specialist doctor.

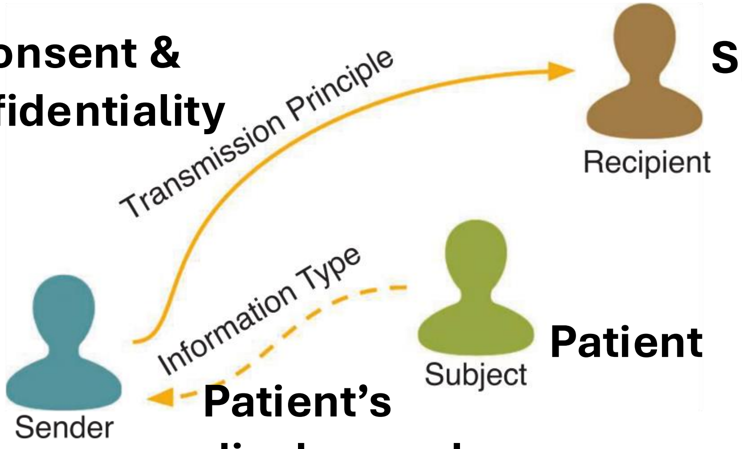

Consent & confidentiality

Specialist doctor

Family doctor

Patient

Patient’s medical record

Diagram adapted from Malkin 2023

Example: Medical referral

Context: healthcare

Contextual purposes: Improve patient health, efficient division of labor, fair and wide access to healthcare, transparency to patients, etc.

Information norm: The patient’s medical record should flow from the family doctor to the specialist doctor as long as the patient consents to this sharing, and their records are not shared beyond the specialist doctor.

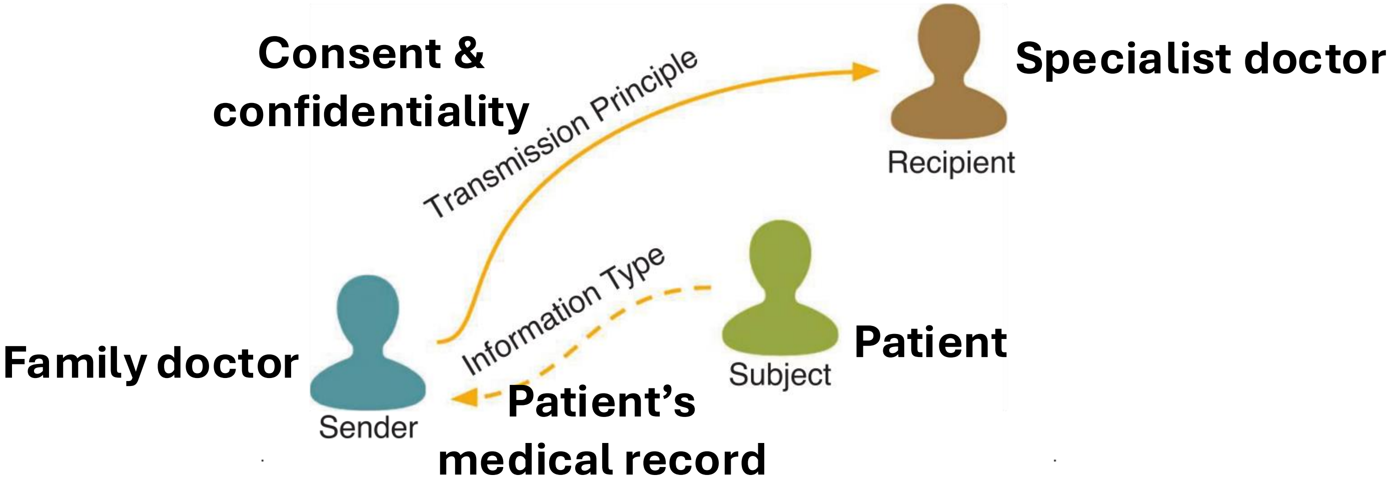

Consent & confidentiality

Specialist doctor

Family doctor

Patient

Patient’s medical record

What if the family doctor instead shares the patient’s medical record with the **patient’s employer**?

Would this disruption undermine any contextual purposes?

Diagram adapted from Malkin 2023

Example: Medical referral

Context: healthcare

Contextual purposes: Improve patient health, efficient division of labor, fair and wide access to healthcare, transparency to patients, etc.

Information norm: The patient's medical record should flow from the family doctor to the specialist doctor as long as the patient consents to this sharing, and their records are not shared beyond the specialist doctor.

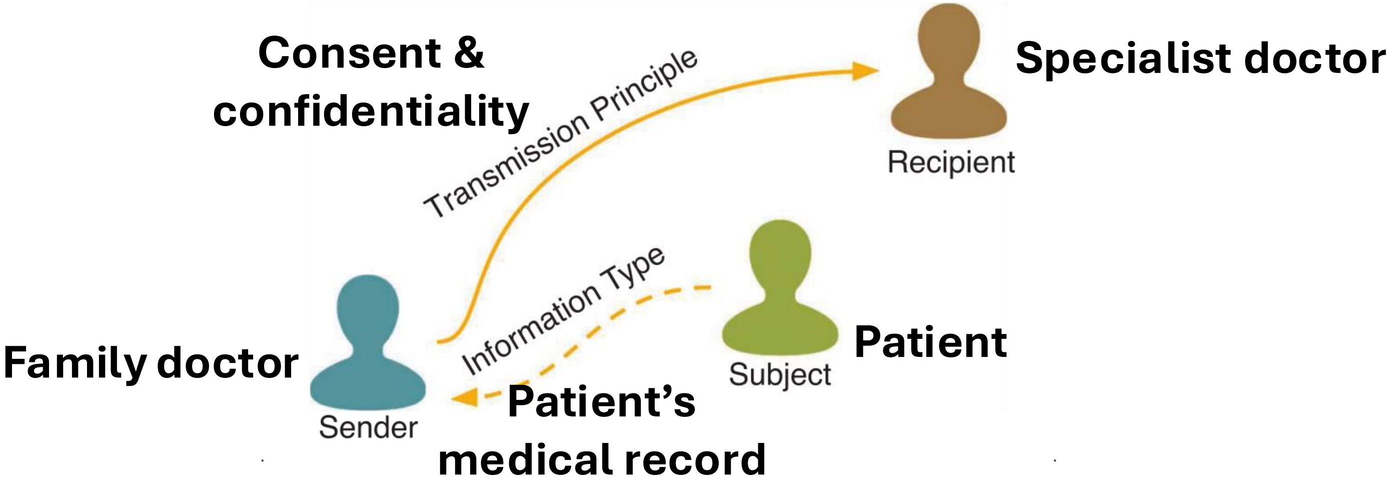

Consent & confidentiality

Specialist doctor

Family doctor

Patient

Patient's medical record

What if the family doctor uses a new app to share medical records with the specialist doctor, but the app developer also sees and stores medical records?

Would this disruption undermine any contextual purposes?

Diagram adapted from Malkin 2023

Activity: Zoom attention tracking (20 minutes)

In-person → Zoom classrooms (during the COVID-19 pandemic)

Zoom introduced an “attention tracking” feature – showed an icon next to meeting attendees who had not been in the Zoom application for over 30 seconds and at the end of the meeting generated a percentage breakdown for each participant for how long they had the Zoom window open during the meeting

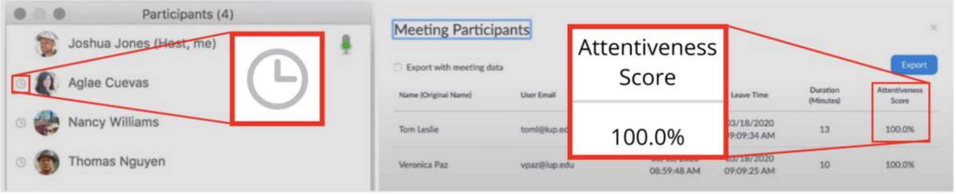

Participants (4)

Joshua Jones (Host, me)
Aglae Cuevas
Nancy Williams
Thomas Nguyen

Meeting Participants

□ Export with meeting data

Name (Original Name)	User Email		Leave Time	Duration (Minutes)	Attentiveness Score
Tom Leslie	toml@iup.ed		3/18/2020 9:09:34 AM	13	100.0%
Veronica Paz	vpaz@iup.edu	08:59:48 AM	03/18/2020 09:09:25 AM	10	100.0%

Attentiveness Score

100.0%

First image from Li Arya Jin 2024
Diagram adapted from Malkin 2023

Activity: Zoom attention tracking (20 minutes)

In-person → Zoom classrooms (during the COVID-19 pandemic)

Zoom introduced an “attention tracking” feature – showed an icon next to meeting attendees who had not been in the Zoom application for over 30 seconds and at the end of the meeting generated a percentage breakdown for each participant for how long they had the Zoom window open during the meeting

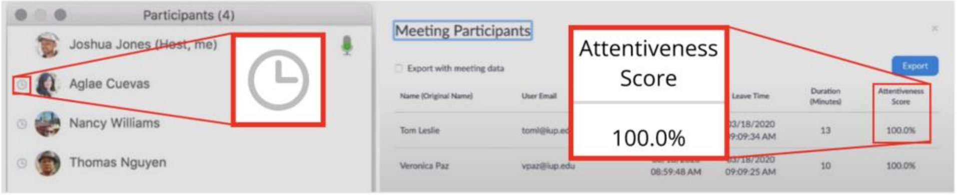

Participants (4)

- Joshua Jones (Host, me)
- Aglae Cuevas
- Nancy Williams
- Thomas Nguyen

Meeting Participants

□ Export with meeting data

Name (Original Name)	User Email		Leave Time	Duration (Minutes)	Attentiveness Score
Tom Leslie	toml@iup.ed		03/18/2020 09:09:34 AM	13	100.0%
Veronica Paz	vpaz@iup.edu	08:59:48 AM	03/18/2020 09:09:25 AM	10	100.0%

1. Establish the relevant social context and its goals
2. Identify information norms in the context by specifying the five parameters
3. Identify ways in which the technology or new practice disrupts that information norm
4. Evaluate whether the disruptions identified in Step 3 undermine the goals of the context
5. Make a judgment about whether an information practice should be abandoned or changed

Image from Li Arya Jin 2024

Taking a step back

1. How does the definition of DP differ from the concept of privacy as the context-appropriate flow of information? When might the two come apart?

2. When and to what extent does DP address legitimate privacy concerns? What kinds of privacy concerns does it address?

Activity: Story of Mr. E (10 min)

Imagine a study asks a random sample of voters in a town about their race and whether they are registered as Republican, Democrat, or Independent. *A local resident named Mr. E is included in this sample.* The investigators *use differentially-private mechanisms* to compute proportions of party registration by race, along with confidence intervals around these proportions.

Their results indicate that Asian people in the town are almost always registered as Democrats, however people of other racial backgrounds in the town are almost always registered as Republicans.

Mr. E is an Asian person living in the town. His boss, a registered Republican, knows Mr. E's race and uses the results of the study to conclude that Mr. E is likely a Democrat.

Mr. E finds out what his boss has learned and is upset, feeling that his privacy was violated.

Activity: Story of Mr. E

1. Establish the relevant social context and its goals
2. Identify information norms in the context by specifying the five parameters
3. Identify ways in which the technology or new practice disrupts that information norm
4. Evaluate whether the disruptions identified in Step 3 undermine the goals of the context
5. Make a judgment about whether an information practice should be abandoned or changed

Case study: Applying Contextual Integrity to the 2020 U.S. Census

Brief History of Census Disclosure Avoidance Protection

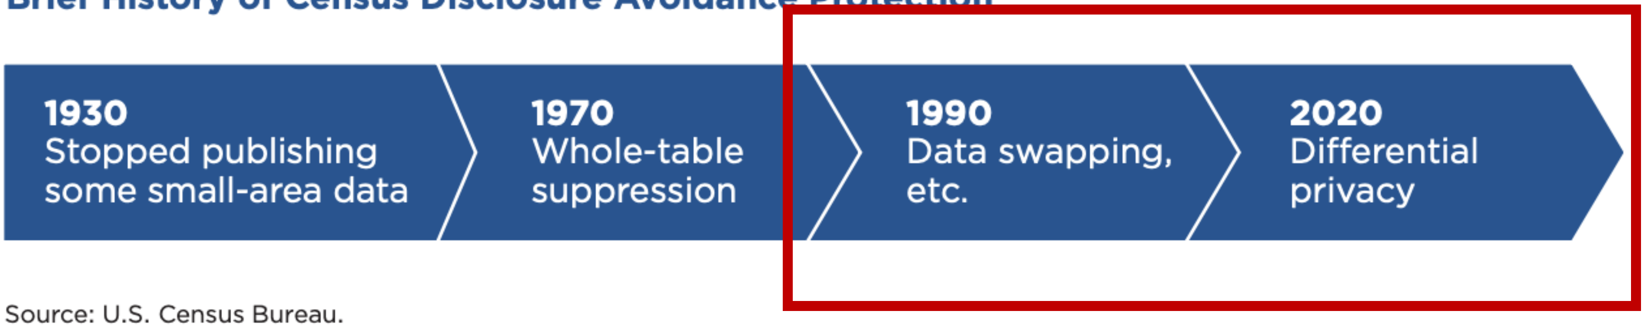

1930
Stopped publishing some small-area data

1970
Whole-table suppression

1990
Data swapping, etc.

2020
Differential privacy

Source: U.S. Census Bureau.

This example is based on the case study in Benthall & Cummings ‘24

Case study: Applying Contextual Integrity to the 2020 U.S. Census

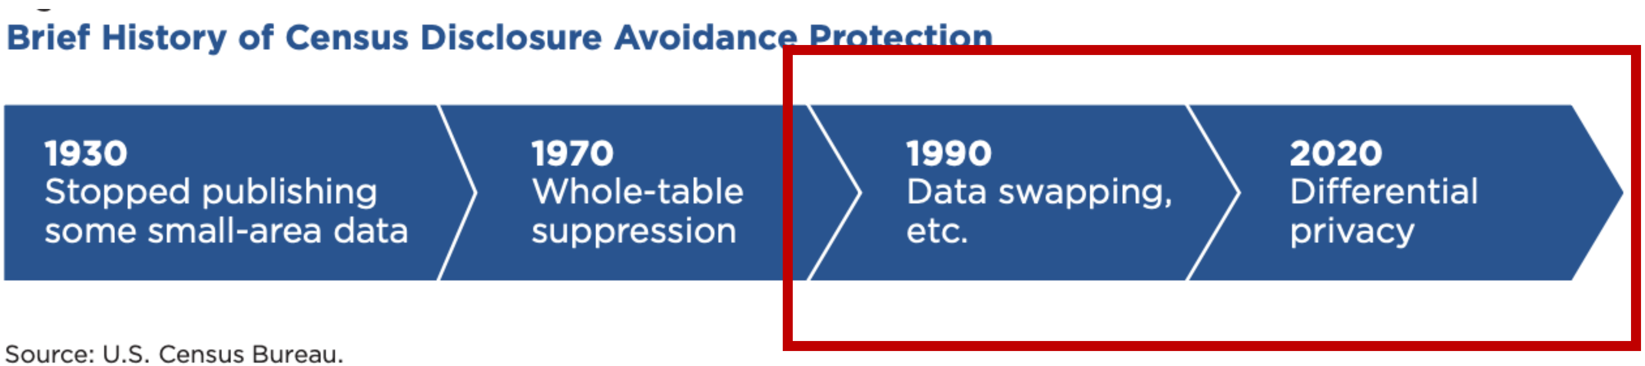

Brief History of Census Disclosure Avoidance Protection

1930 Stopped publishing some small-area data

1970 Whole-table suppression

1990 Data swapping, etc.

2020 Differential privacy

Source: U.S. Census Bureau.

Is swapping or differential privacy more appropriate for the Census?

In other words: Which approach better "[respects] the established legal norms for protecting the privacy of individuals, and [satisfies] the contextual purposes of the Decennial Census"? (BC '24)

This example is based on the case study in Benthall & Cummings '24

Case study: Applying Contextual Integrity to the 2020 U.S. Census

- **Contextual purposes:** counts help allocate seats in the House of Representatives, draw new state & local districts, and allocate federal funding

This example is based on the case study in Benthall & Cummings ‘24

Case study: Applying Contextual Integrity to the 2020 U.S. Census (for redistricting)

- **Contextual purposes:** draw new state & local districts

This example is based on the case study in Benthall & Cummings '24

Case study: Applying Contextual Integrity to the 2020 U.S. Census (for redistricting)

- **Contextual purposes:** draw new state & local districts
- **Information norms:** race, ethnicity, age, etc. can flow from the Census Bureau to the public & redistricters granted that the information was collected lawfully and is shared in such a way that subjects are not identifiable

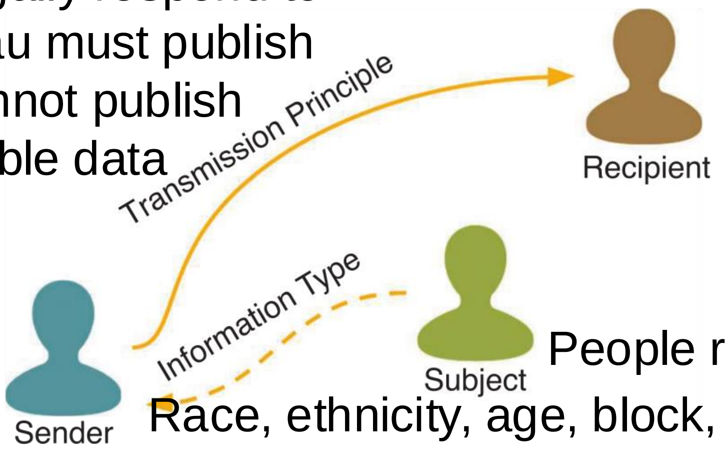

People in the U.S. must legally respond to the Census; Census Bureau must publish useful statistics and cannot publish individually-identifiable data

Redistricters, the public

People residing in U.S.

Census Bureau

Race, ethnicity, age, block, housing type

This example is based on the case study in Benthall & Cummings ‘24

Case study: Applying Contextual Integrity to the 2020 U.S. Census (for redistricting)

- **Contextual purposes:** draw new state & local districts
- **Information norms:** race, ethnicity, age, etc. can flow from the Census Bureau to the public & redistricters granted that the information was collected lawfully and is shared in such a way that subjects are not identifiable

People in the U.S. must legally respond to the Census; Census Bureau must publish useful statistics and cannot publish individually-identifiable data

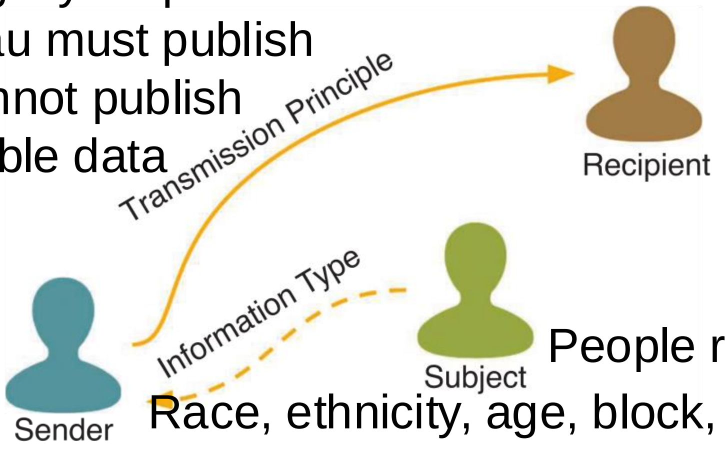

Redistricters, the public

People residing in U.S.

Census Bureau

Race, ethnicity, age, block, housing type

Does swapping or DP help maintain contextual integrity?

This example is based on the case study in Benthall & Cummings ‘24

Proposal by Benthall & Cummings ’24 to integrate CI + DP with a sixth parameter

- **Contextual purposes:** draw new state & local districts
- **Information norms:** race, ethnicity, age, etc. can flow from the Census Bureau to the public & redistricters granted that the information was collected lawfully and is shared in such a way that subjects are not identifiable

People in the U.S. must legally respond to the Census; Census Bureau must publish useful statistics and cannot publish individually-identifiable data

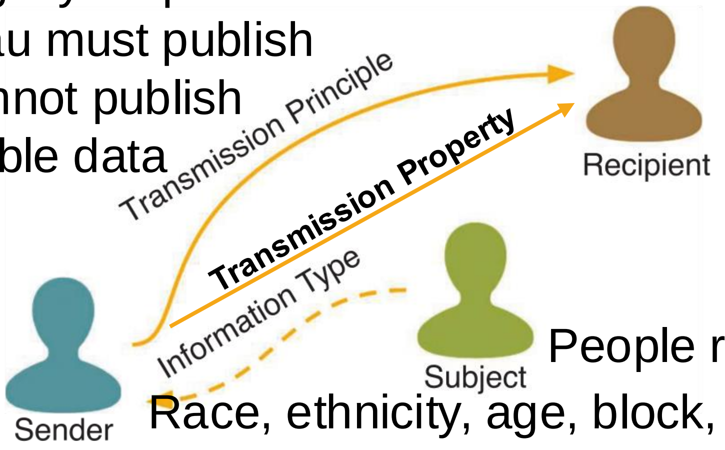

Redistricters, the public

People residing in U.S.

Census Bureau

Race, ethnicity, age, block, housing type

Does swapping or DP help maintain contextual integrity?

This example is based on the case study in Benthall & Cummings ‘24

Proposal by Benthall & Cummings ’24 to integrate CI + DP with a sixth parameter

One direct way to bridge CI + DP is to **add** a parameter to contextual integrity.

1. Sender
2. Recipient
3. Subject
4. Information type
5. Transmission principle
6. **Transmission property**

Normative aspects of how the information is transmitted

Descriptive aspects of how the information is transmitted (e.g., with epsilon-DP, with swapping)

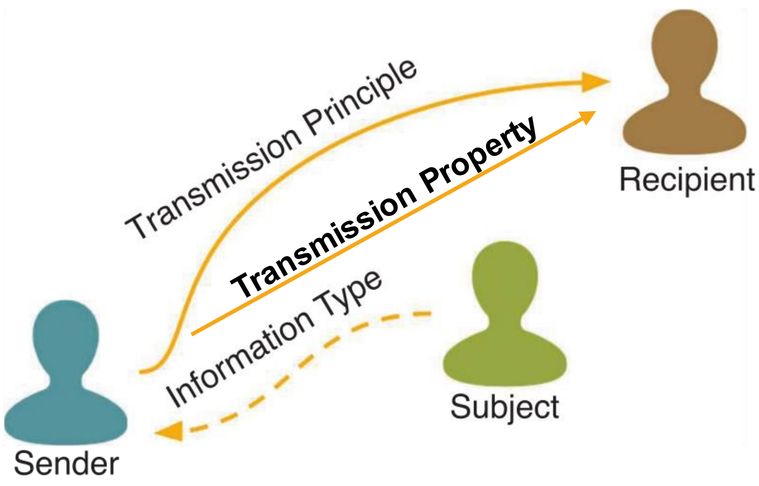

Benthall & Cummings 2024

Examples of transmission *properties* vs. transmission *principles*

Transmission Properties (Descriptive)	**Transmission Principles (Normative)**
Flow with no PET	With subject consent
With Gaussian noise N(0,1) added	Under reciprocity agreement
Within a 95% confidence interval	Constitutes disclosure
Using public-key encryption	With a warrant
With secure multi-party computation	As mandated by law

Table reproduced from Benthall & Cummings 2024

Takeaways

- Contextual Integrity defines privacy as the appropriate flow of information in a given context. Applying Contextual Integrity requires defining a context, its contextual purposes, and information norms (a relation between subject, sender, recipient, information type, transmission principle)

- Contextual Integrity can help us more formally reason about when privacy violations have occurred in a social sense (vs. a mathematical sense, as DP helps us do)

- DP can help achieve contextual integrity in certain contexts; contextual integrity can help us decide when to apply DP (and how)