Differential Privacy in Action

Federico Mazzone

Cybersecurity Seminars

04 April, 2025

UNIVERSITY OF TWENTE. | DIGITAL SOCIETY INSTITUTE

SCS

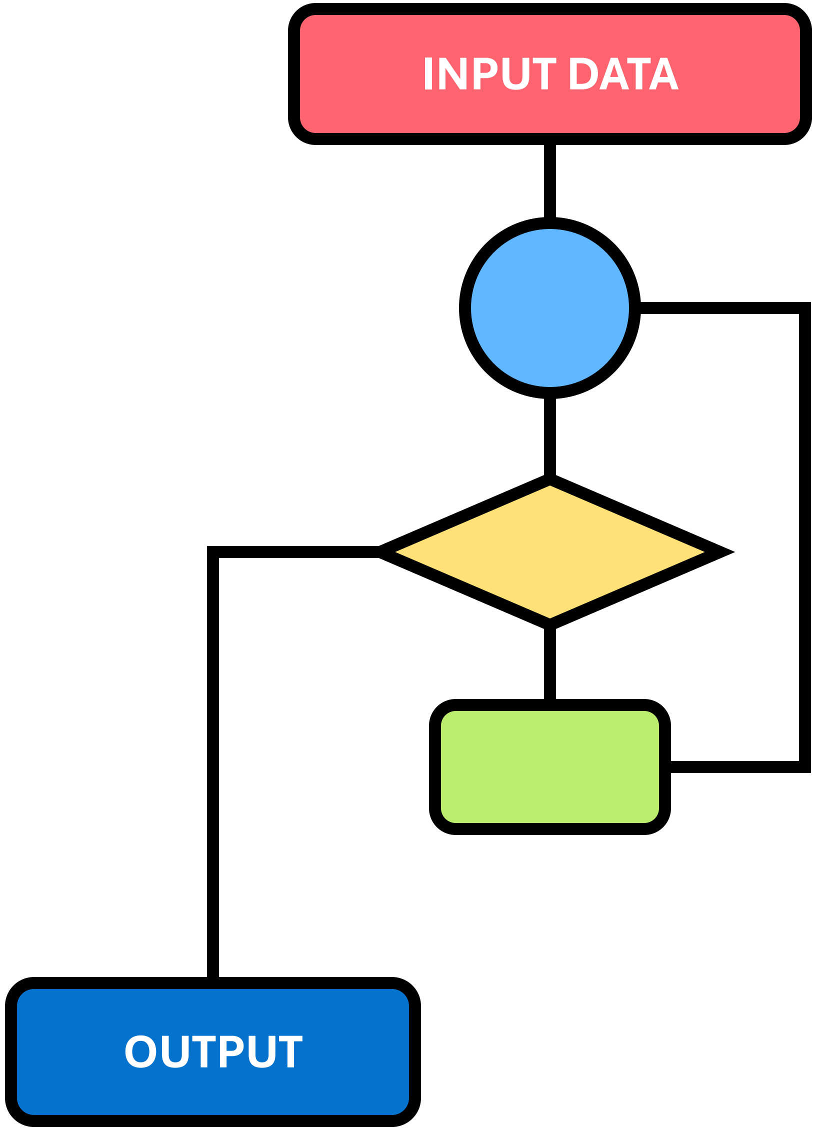
INPUT DATA
OUTPUT

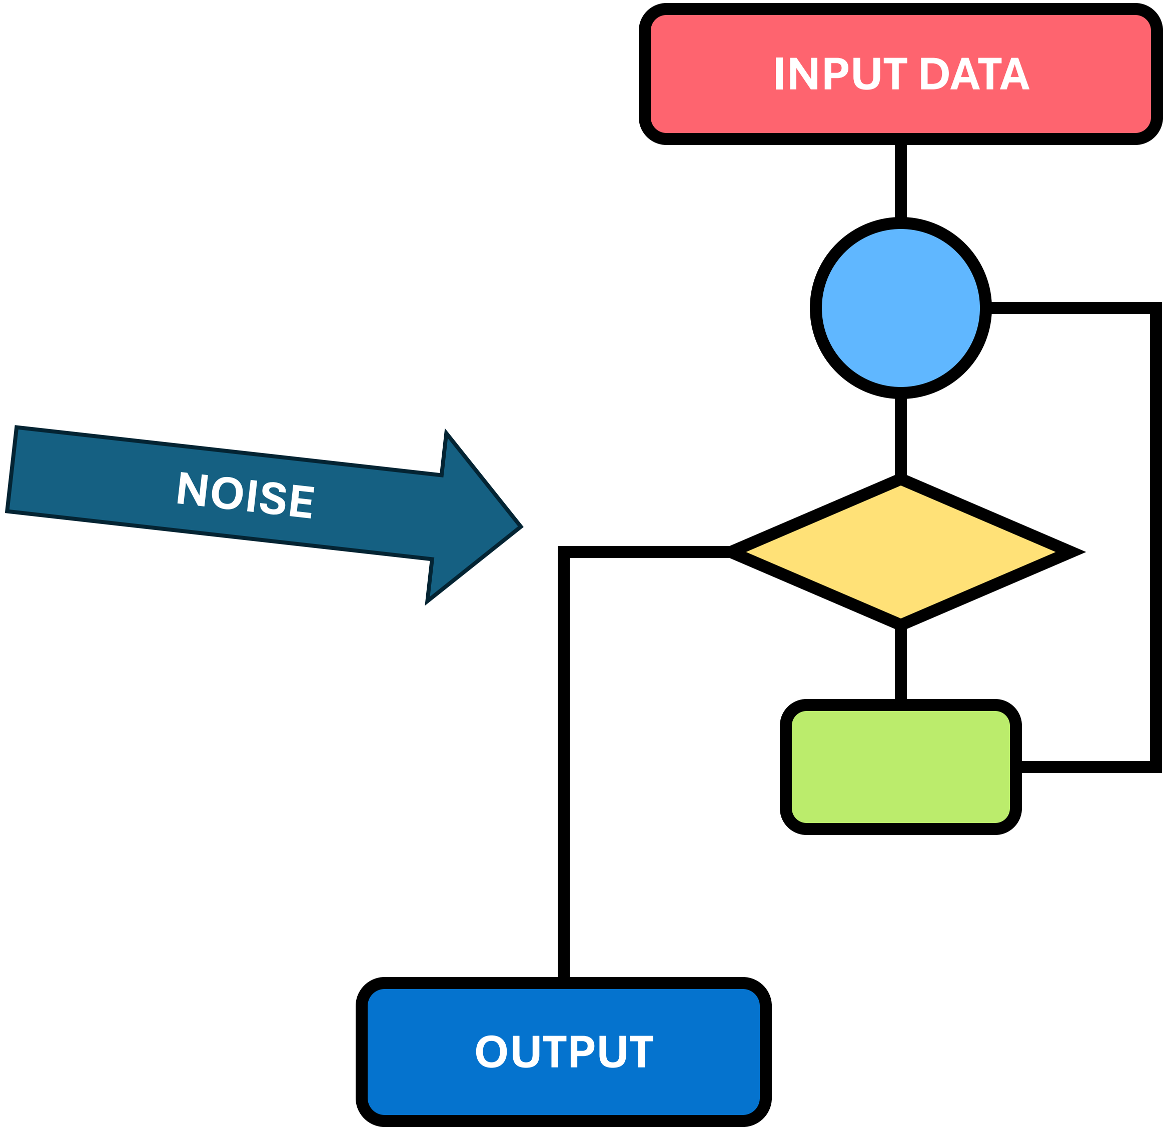
INPUT DATA
NOISE
OUTPUT

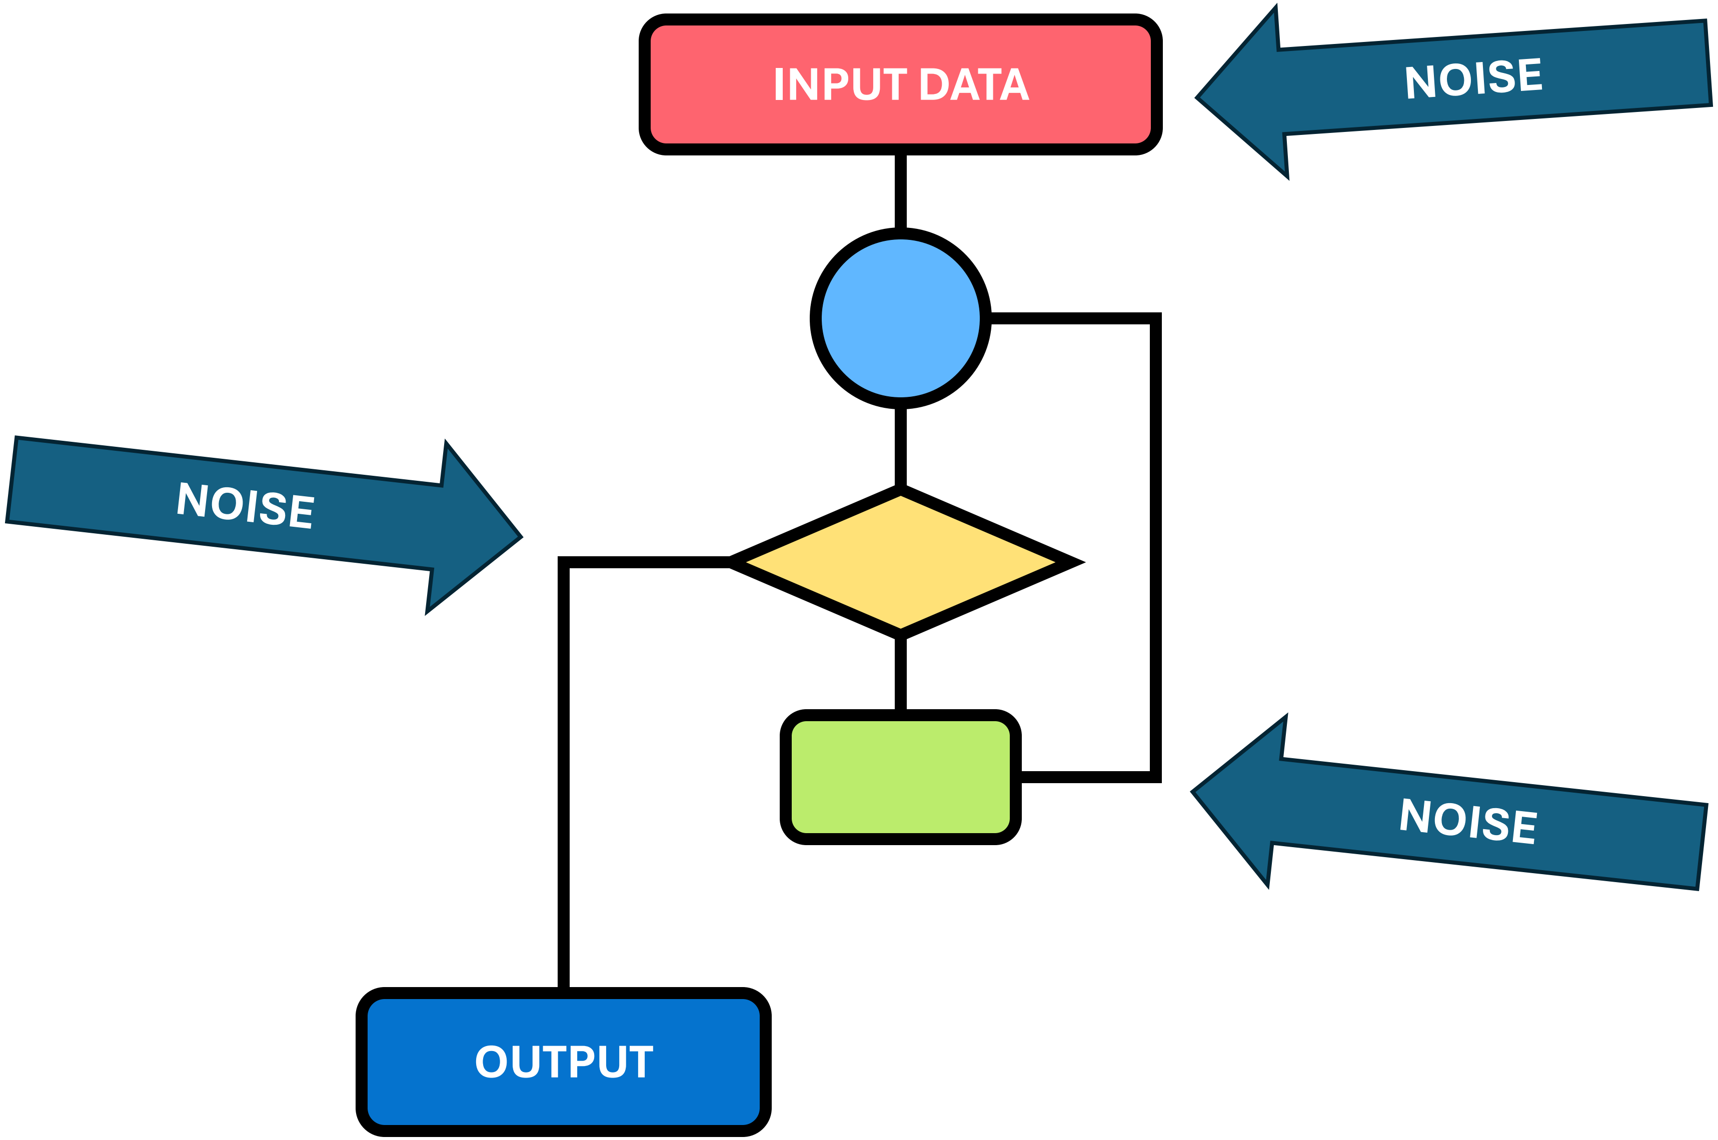
INPUT DATA
NOISE
NOISE
NOISE
OUTPUT

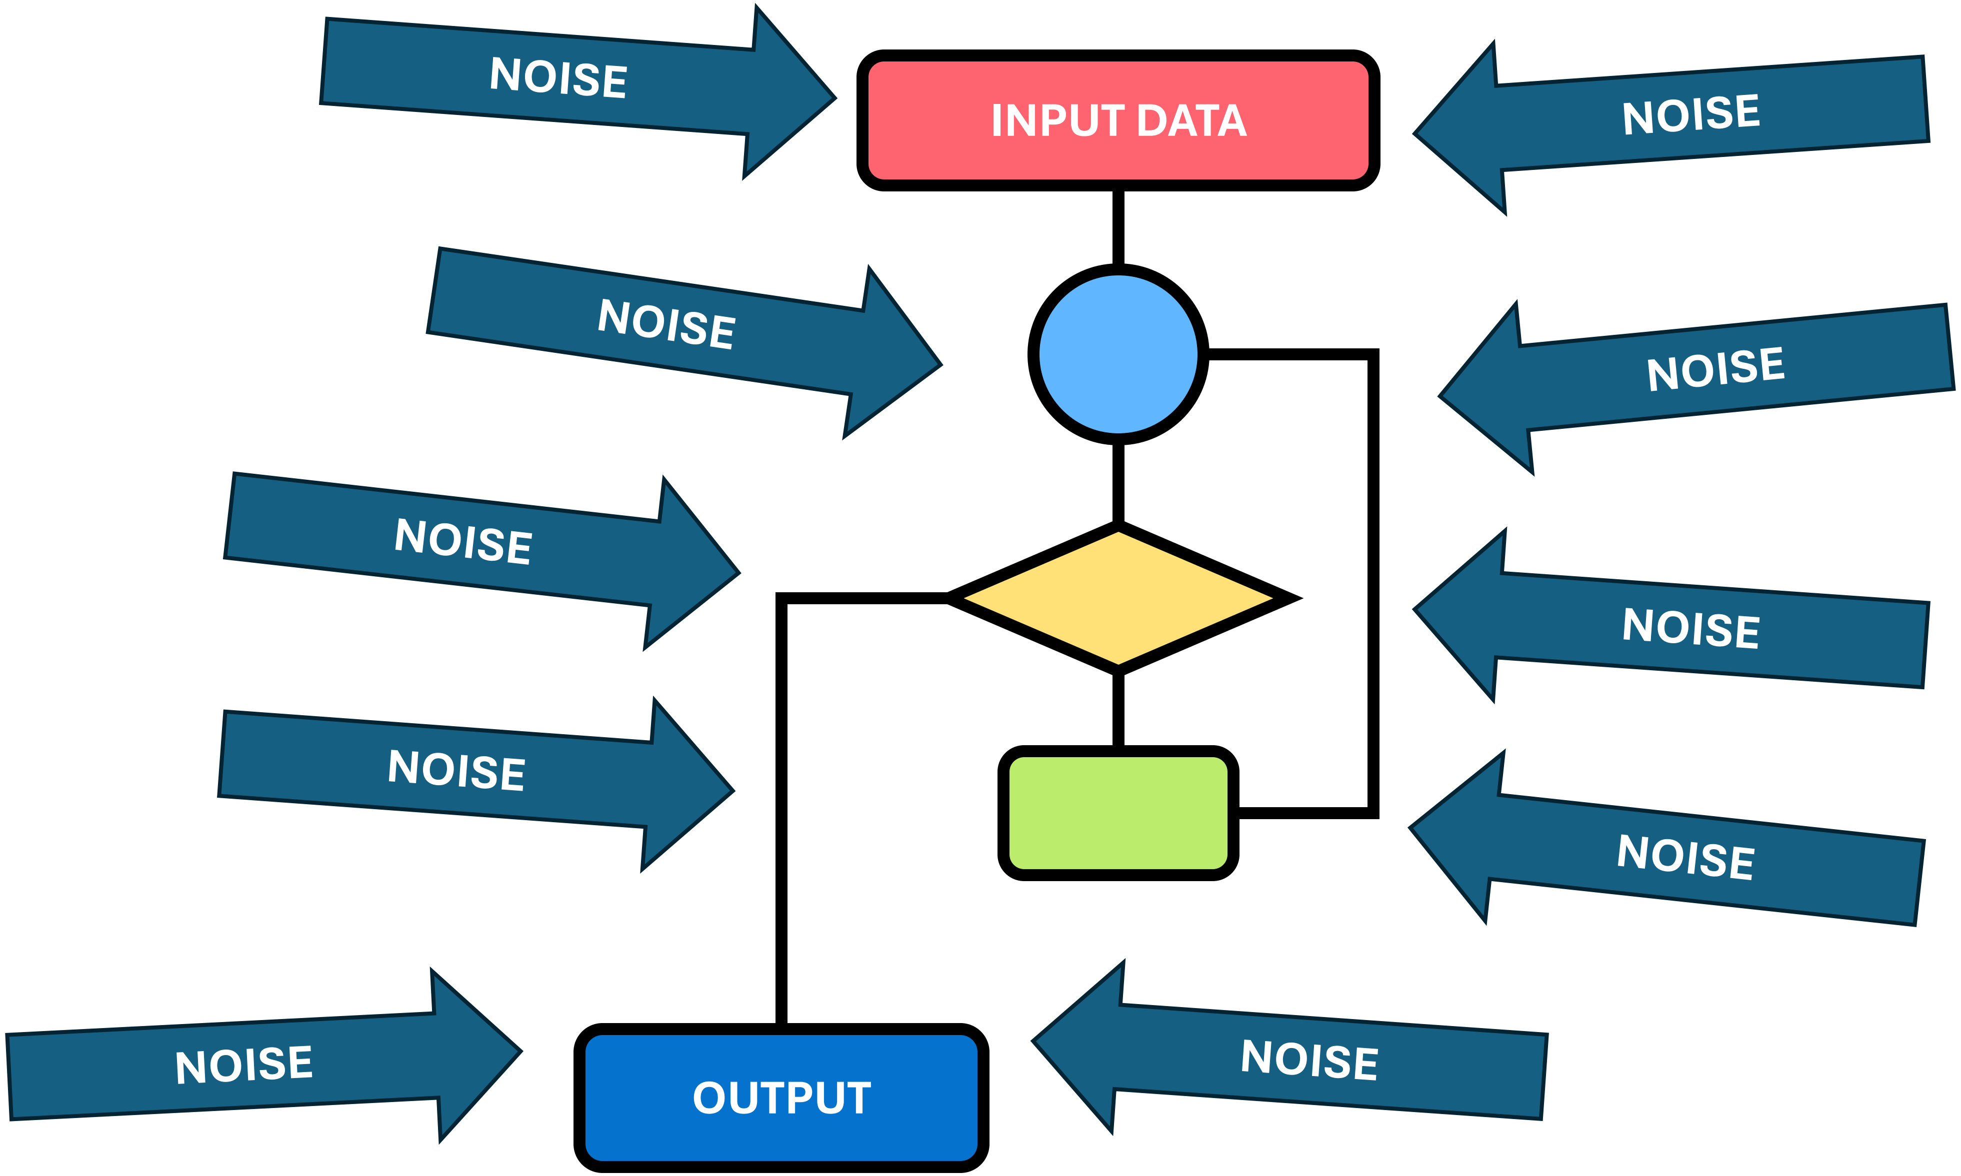
NOISE
INPUT DATA
NOISE
NOISE
NOISE
NOISE
NOISE
NOISE
NOISE
NOISE
OUTPUT
NOISE


NOISE
70
iStock
Credit: Jevtic

From Heuristics to Formal Privacy

Road to Differential Privacy

1965

Randomized Response (Warner)
Simple noise-based technique for survey privacy

From Heuristics to Formal Privacy

Road to Differential Privacy

1965

Randomized Response (Warner)
Simple noise-based technique for survey privacy

1980s-90s

Data Perturbation Methods
Ad-hoc noise injection in statistical databases without formal guarantees

From Heuristics to Formal Privacy

Road to Differential Privacy

1965

Randomized Response (Warner)
Simple noise-based technique for survey privacy

1980s-90s

Data Perturbation Methods
Ad-hoc noise injection in statistical databases without formal guarantees

1993

Statistical Disclosure Control (SDC)
Early framework for protecting sensitive data via obfuscation

From Heuristics to Formal Privacy

Road to Differential Privacy

1965

Randomized Response (Warner)
Simple noise-based technique for survey privacy

1980s-90s

Data Perturbation Methods
Ad-hoc noise injection in statistical databases without formal guarantees

1993

Statistical Disclosure Control (SDC)
Early framework for protecting sensitive data via obfuscation

2000s

Cryptographic Approaches to Privacy
Noise introduced in secure computations (e.g., LWE-based encryption)

From Heuristics to Formal Privacy

Road to Differential Privacy

1965

Randomized Response (Warner)
Simple noise-based technique for survey privacy

1980s-90s

Data Perturbation Methods
Ad-hoc noise injection in statistical databases without formal guarantees

1993

Statistical Disclosure Control (SDC)
Early framework for protecting sensitive data via obfuscation

2000s

Cryptographic Approaches to Privacy
Noise introduced in secure computations (e.g., LWE-based encryption)

2006

Differential Privacy (Dwork et al.)
Formalization of privacy with rigorous mathematical guarantees

From Heuristics to Formal Privacy
Road to Differential Privacy

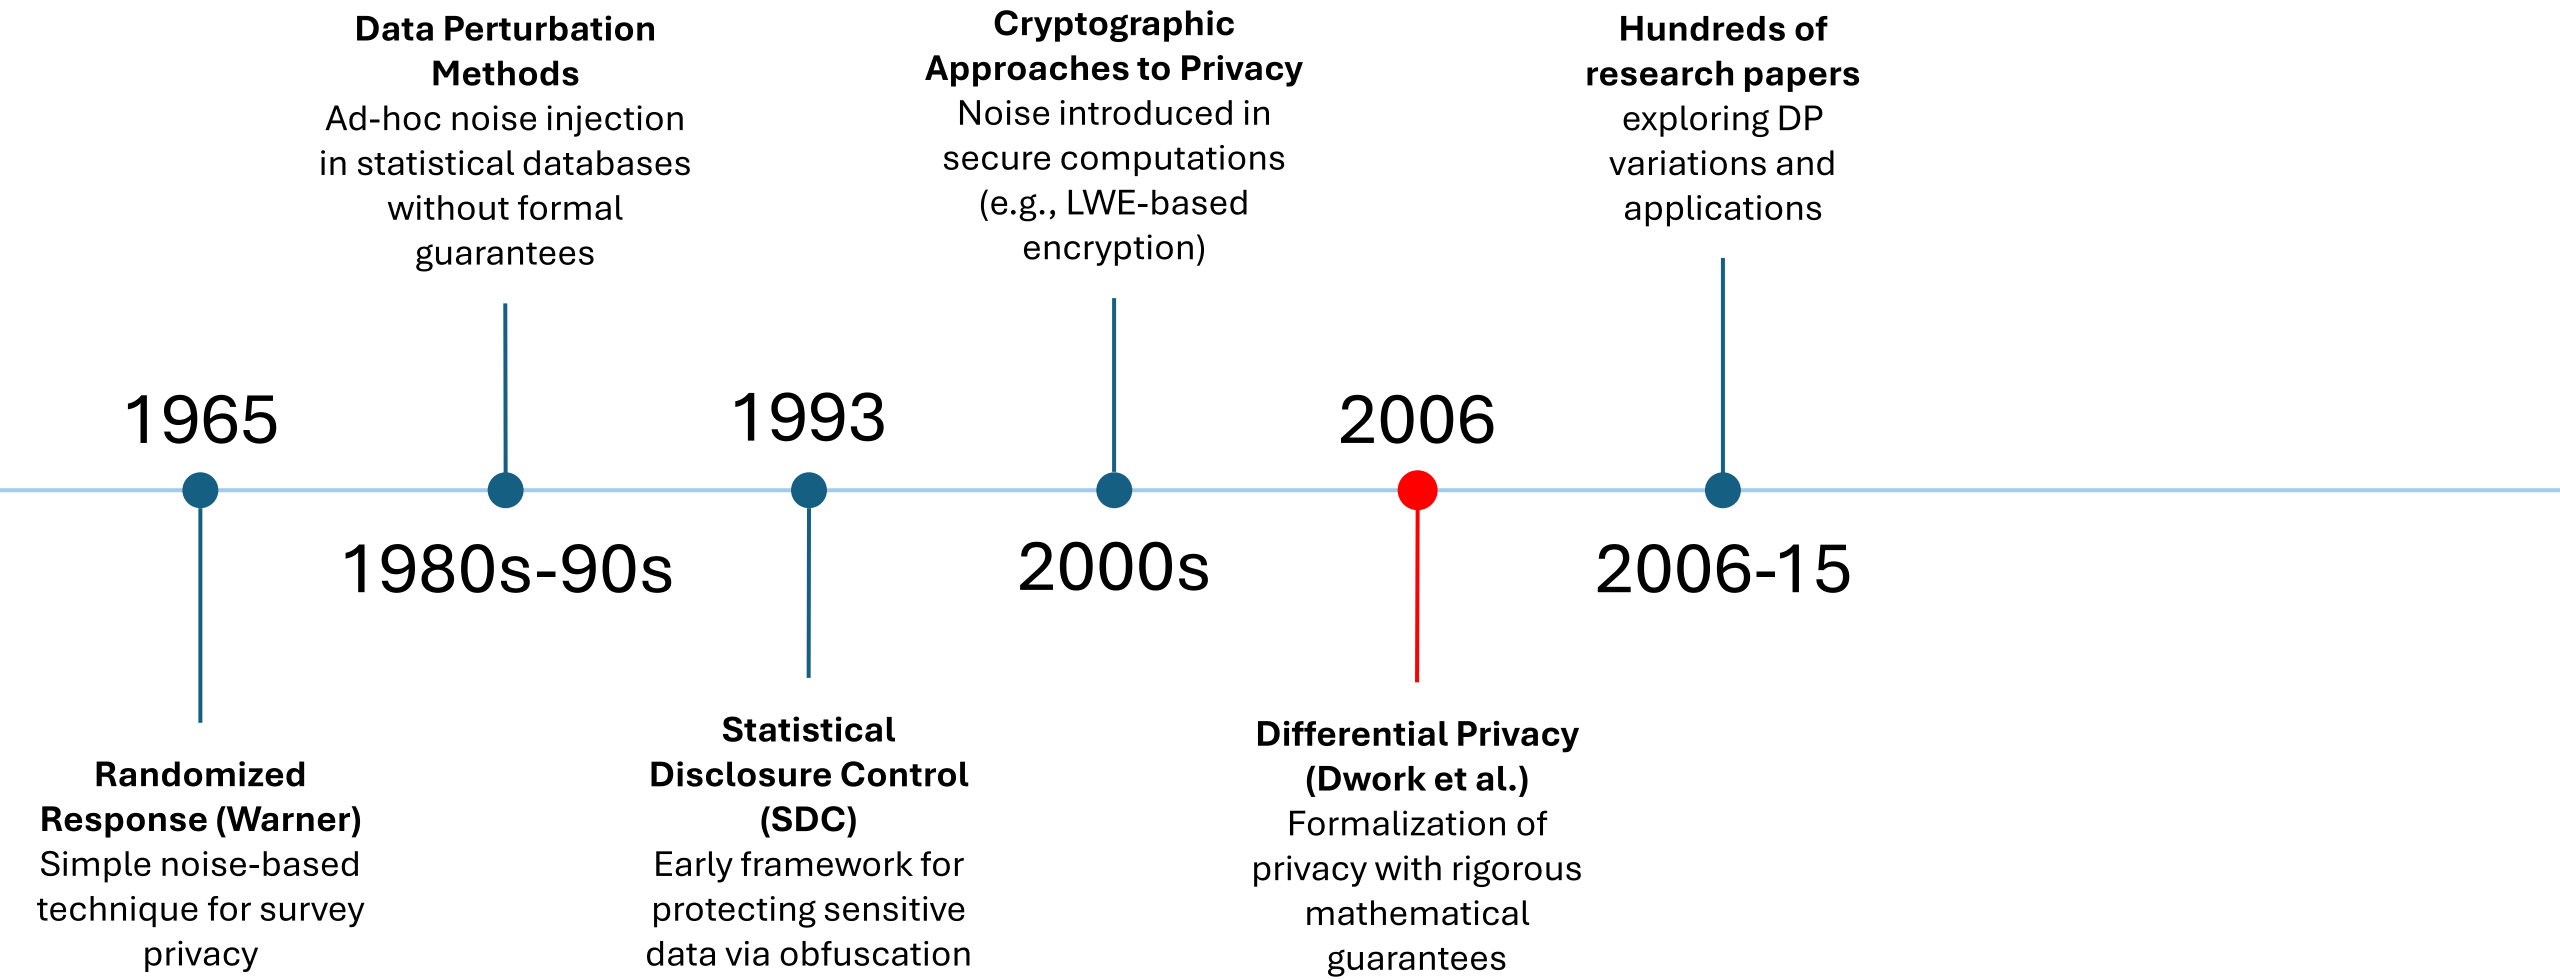

From Heuristics to Formal Privacy

Road to Differential Privacy

- **1965** — **Randomized Response (Warner)** Simple noise-based technique for survey privacy
- **1980s-90s** — **Data Perturbation Methods** Ad-hoc noise injection in statistical databases without formal guarantees
- **1993** — **Statistical Disclosure Control (SDC)** Early framework for protecting sensitive data via obfuscation
- **2000s** — **Cryptographic Approaches to Privacy** Noise introduced in secure computations (e.g., LWE-based encryption)
- **2006** — **Differential Privacy (Dwork et al.)** Formalization of privacy with rigorous mathematical guarantees
- **2006-15** — **Hundreds of research papers** exploring DP variations and applications
- **2016-20** — **Thousands of papers**, industry adoption (Apple, Google, US Census Bureau)

From Heuristics to Formal Privacy

Road to Differential Privacy

1965
Randomized Response (Warner)
Simple noise-based technique for survey privacy

1980s-90s
Data Perturbation Methods
Ad-hoc noise injection in statistical databases without formal guarantees

1993
Statistical Disclosure Control (SDC)
Early framework for protecting sensitive data via obfuscation

2000s
Cryptographic Approaches to Privacy
Noise introduced in secure computations (e.g., LWE-based encryption)

2006
Differential Privacy (Dwork et al.)
Formalization of privacy with rigorous mathematical guarantees

2006-15
Hundreds of research papers exploring DP variations and applications

2016-20
Thousands of papers, industry adoption (Apple, Google, US Census Bureau)

2021-now
10,000+ papers citing DP, expanding into machine learning and security

How Does DP Work?

- Differential privacy adds noise to a function, hiding how much an individual data point can influence the result.

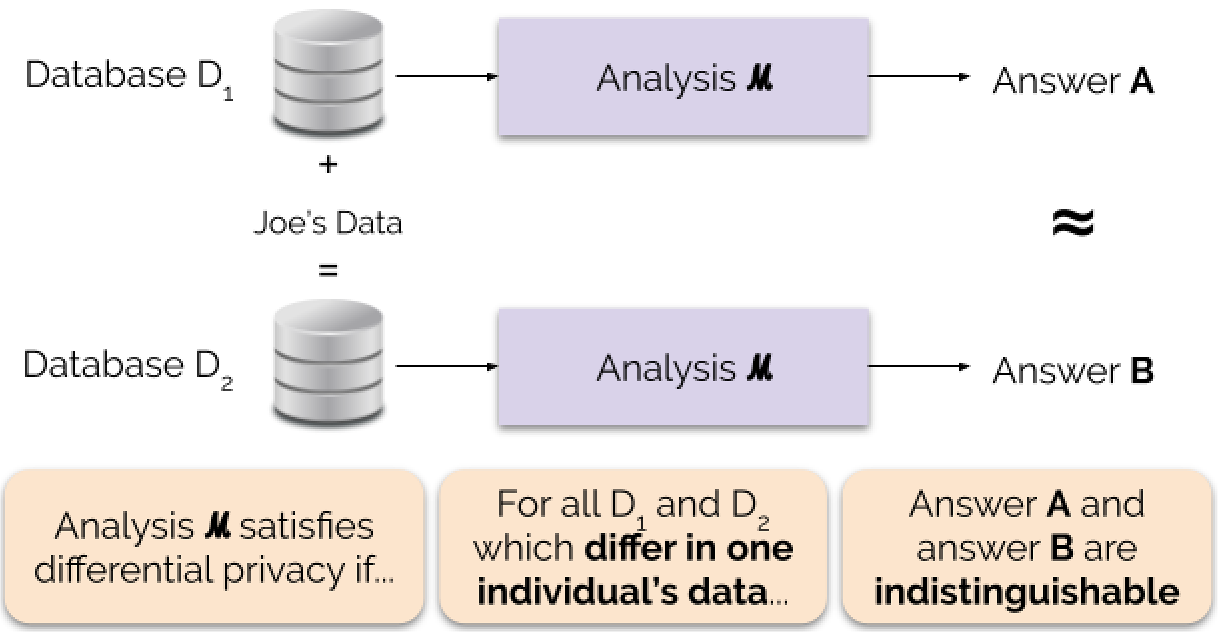

Image by Joseph Near, David Darais and Kaitlin Boeckl - https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-ourCopyright info at https://www.nist.gov/oism/copyrights, Public Domain, https://commons.wikimedia.org/w/index.php?curid=150281038

Trade-Off with Utility

- Example: Counting Query
- $f(D)$ = number of people in D with a given disease
- $\tilde{f}(D) = f(D) + \text{Lap}(1/\epsilon)$

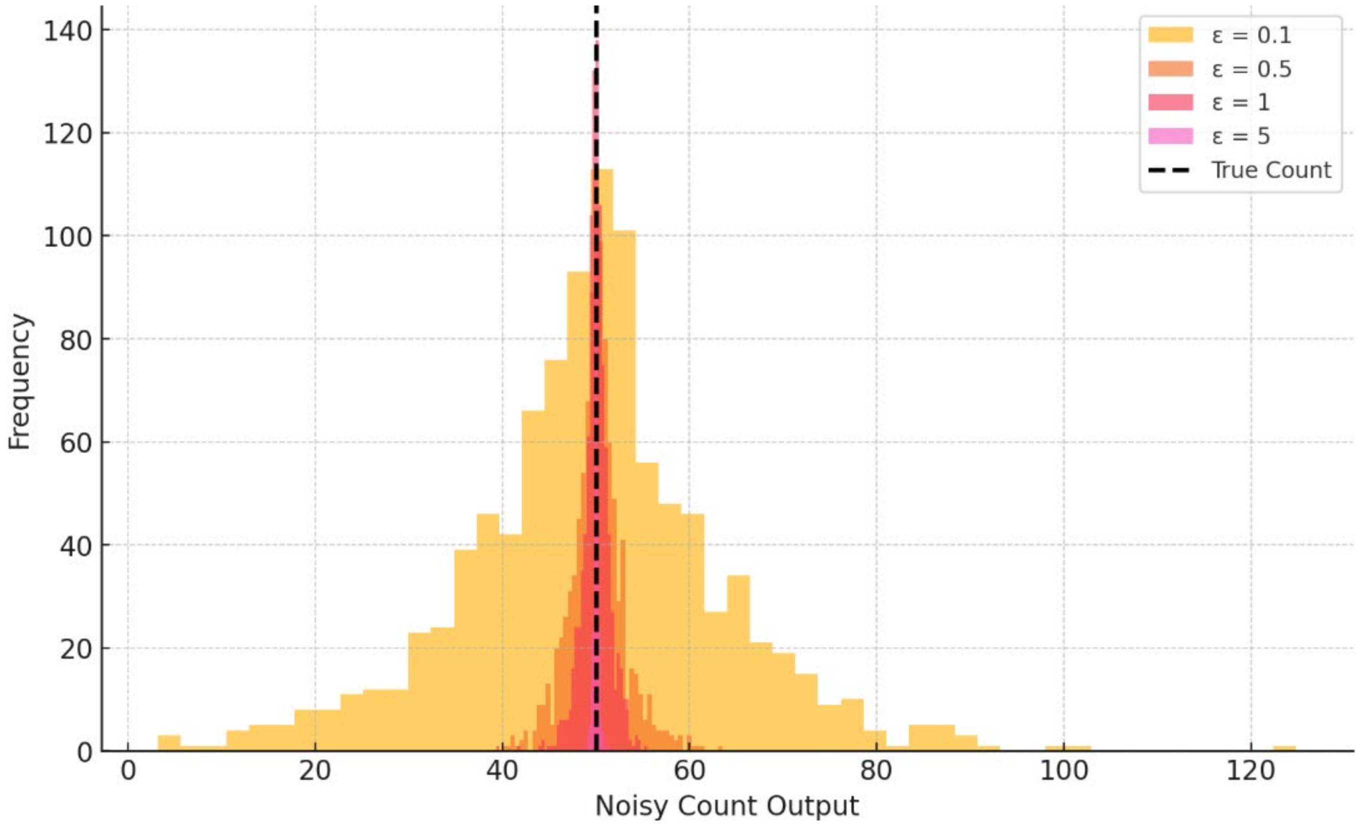

What is DP formalization actually providing?

- With different ϵ you get different points in the trade-off, but how much privacy am I getting from this?

$$\Pr[f(D_1) \in O] \leq e^{\epsilon} \Pr[f(D_2) \in O]$$

- Not an absolute linking between noise and concrete privacy, that is too much application dependant.
- It helps to measure how much noise to provide to ensure the <u>same "level of privacy" across different instances</u>.

What is DP formalization actually providing?

- Example: Sum Query
- $f(D) =$ sum of people's ages
- $\tilde{f}(D) = f(D) + \text{Lap}(120/\epsilon)$

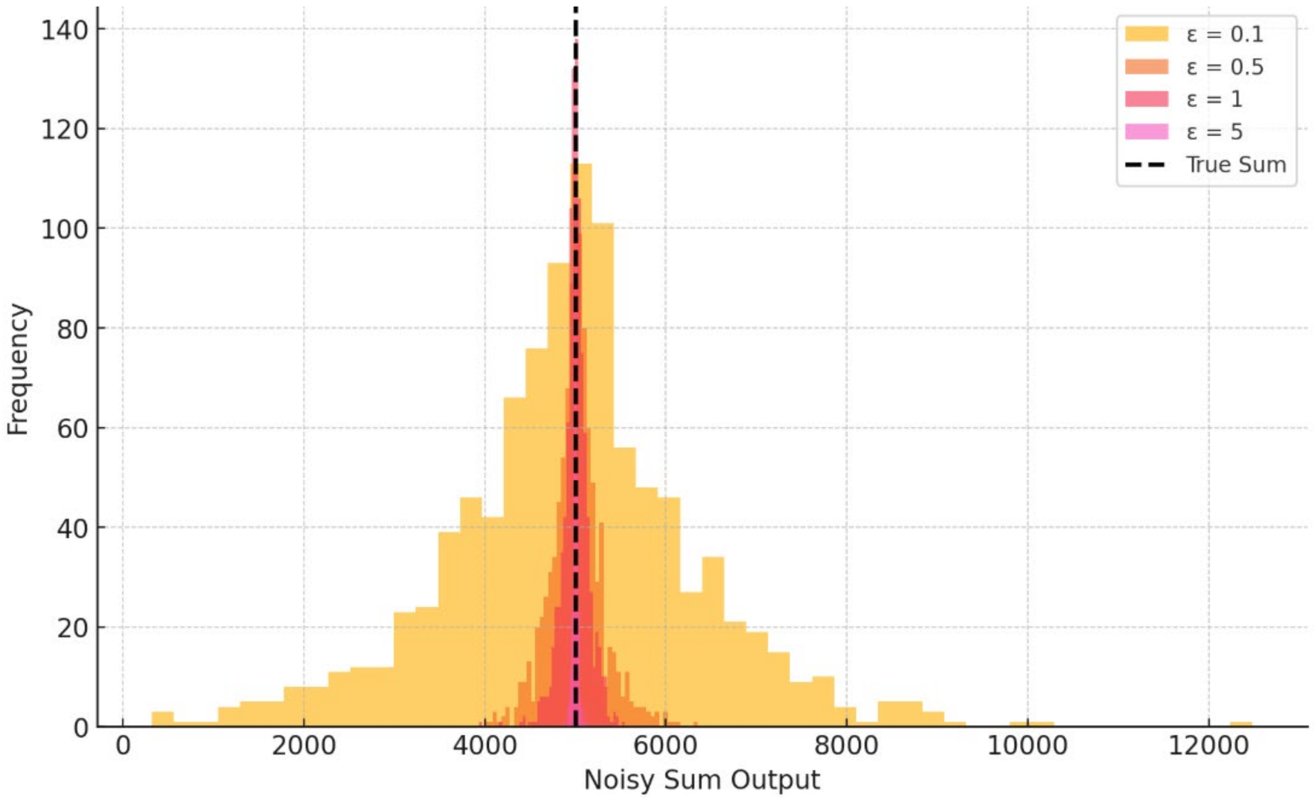

What is DP formalization actually providing?

- Example: Sum Query
- $f(D)$ = sum of people's ages
- $\tilde{f}(D) = f(D) + \mathrm{Lap}(120/\epsilon)$
- $\tilde{f}(D) = f(D) + \mathrm{Lap}(\Delta/\epsilon)$
- If a function is highly sensitive, it means the output is heavily dependent on individual inputs.

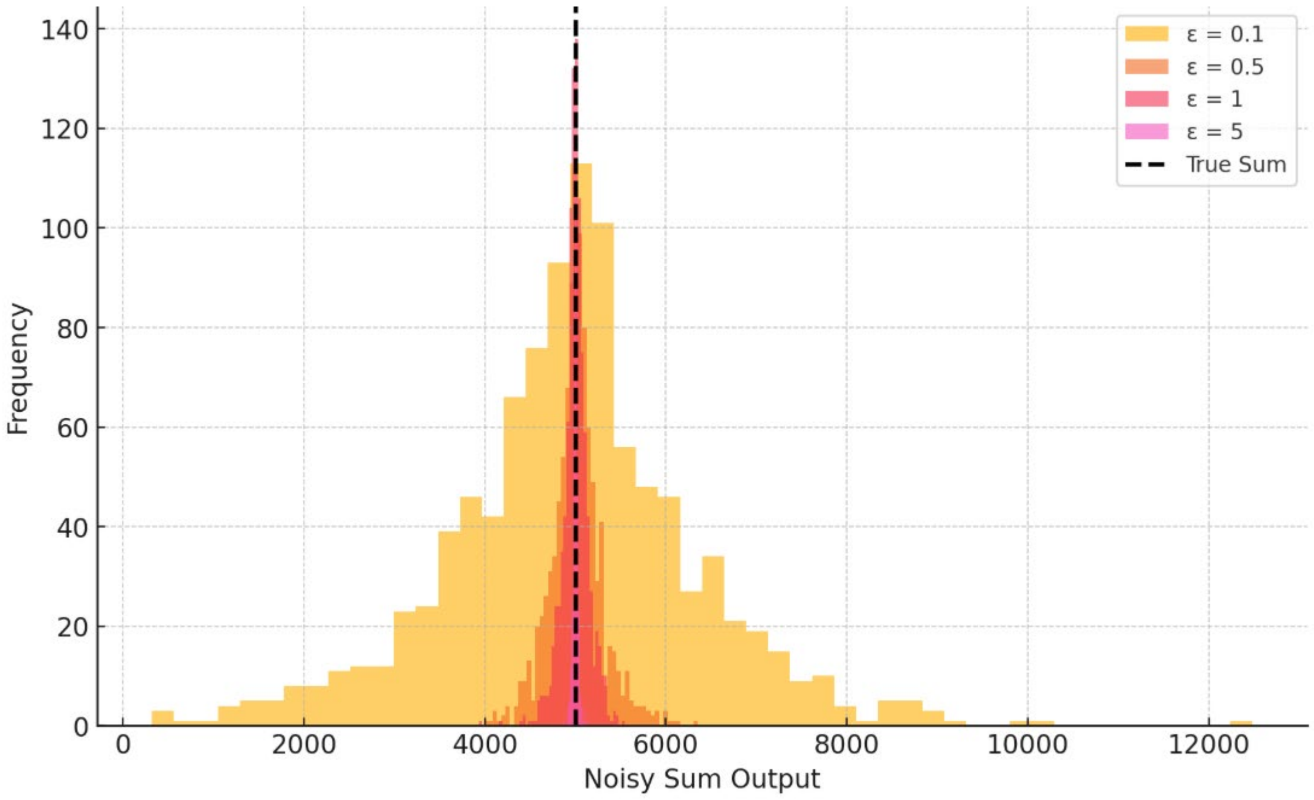

- It also helps to measure how much noise to provide to ensure the <u>same "level of privacy" across different points in the same algorithm</u>.

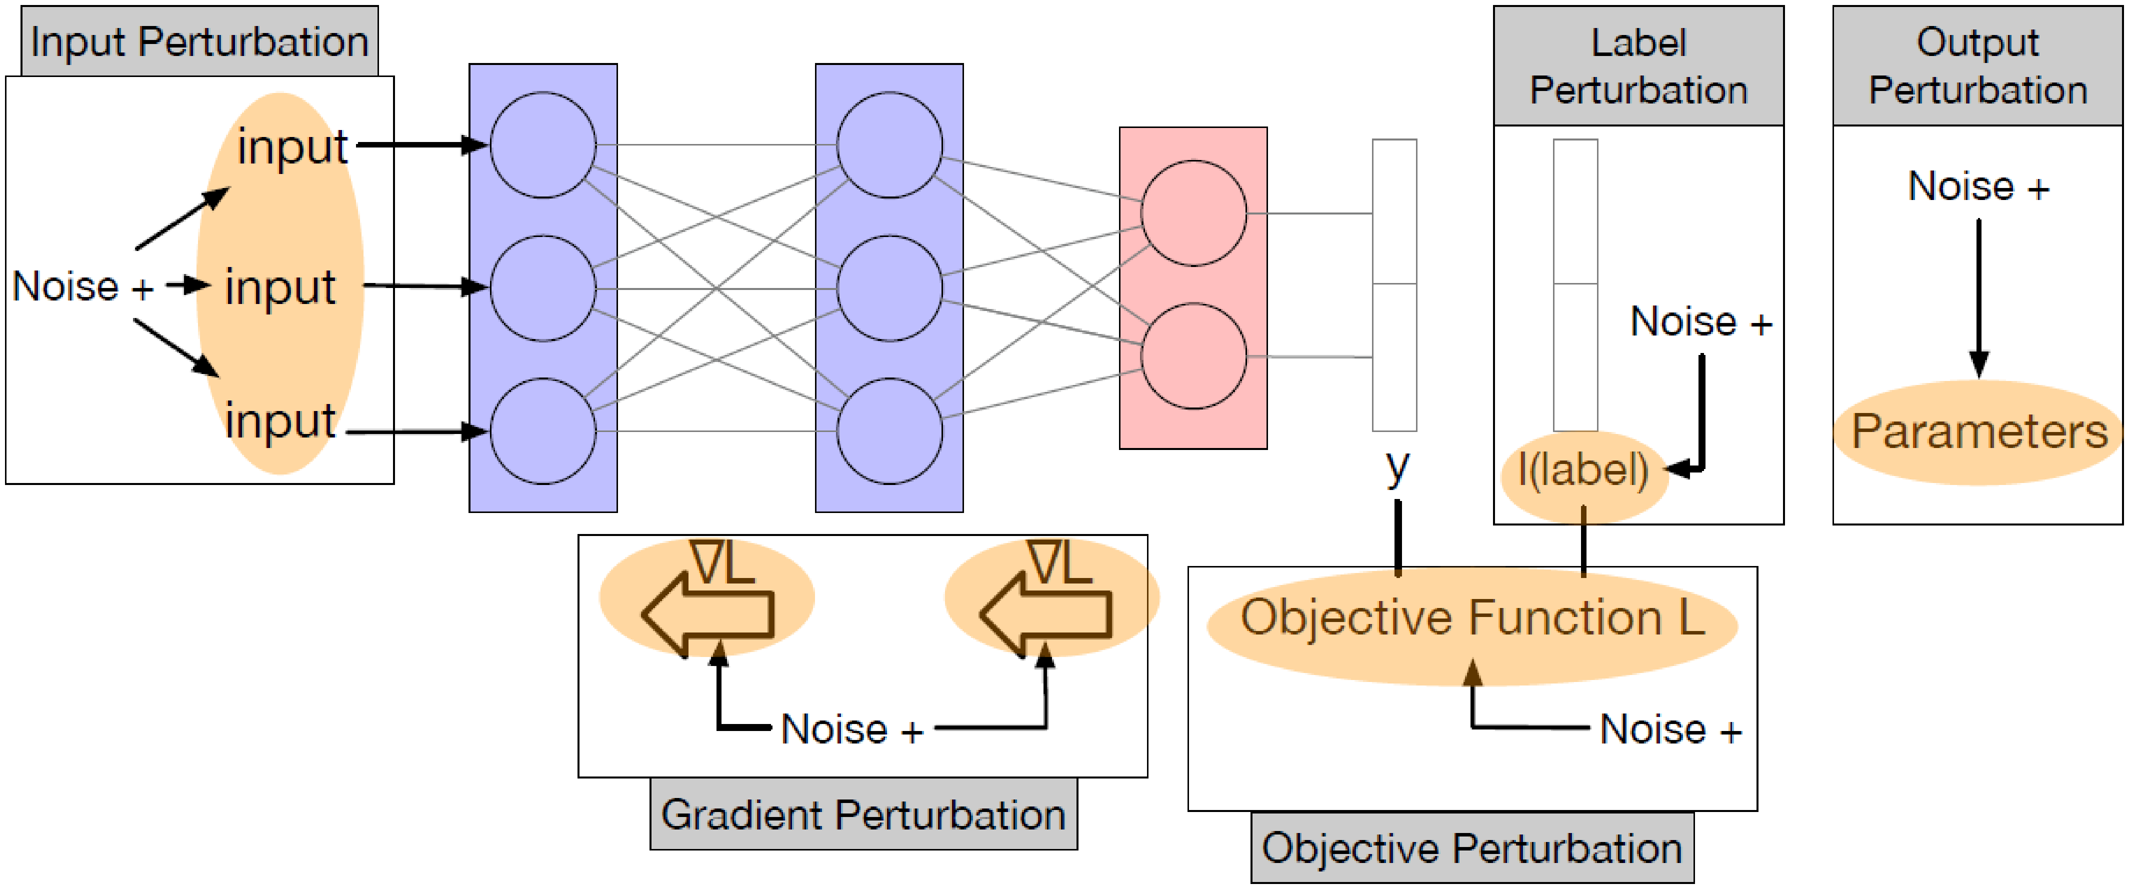

DP in Protocols

- In the context of protocols choosing where to inject noise can affect how much performance you gain or lose.

- Example of vertically-partitioned clustering.

Vertically-Partitioned Clustering

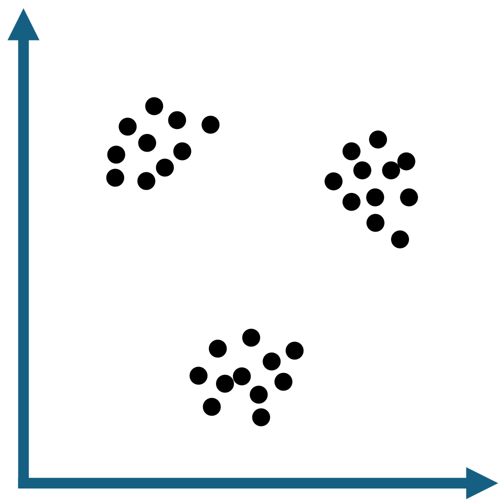

Vertically-Partitioned Clustering

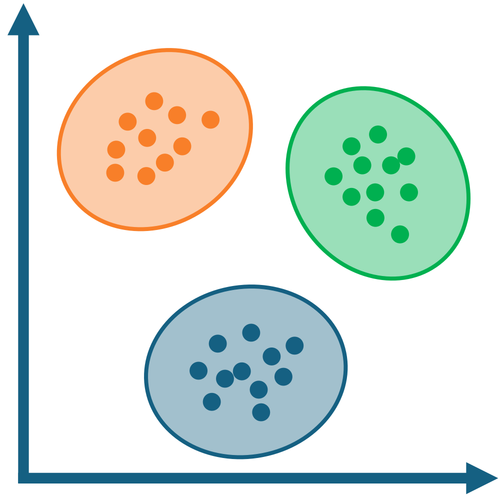

Vertically-Partitioned Clustering

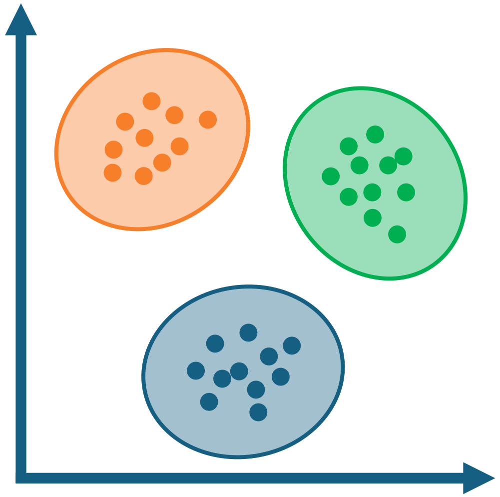

ID	age	visits	score	freq.	bonus
1	23	1240	5	0.70	true
2	19	2256	7	0.55	false
3	47	3210	9	0.12	false
4	32	889	3	0.98	true

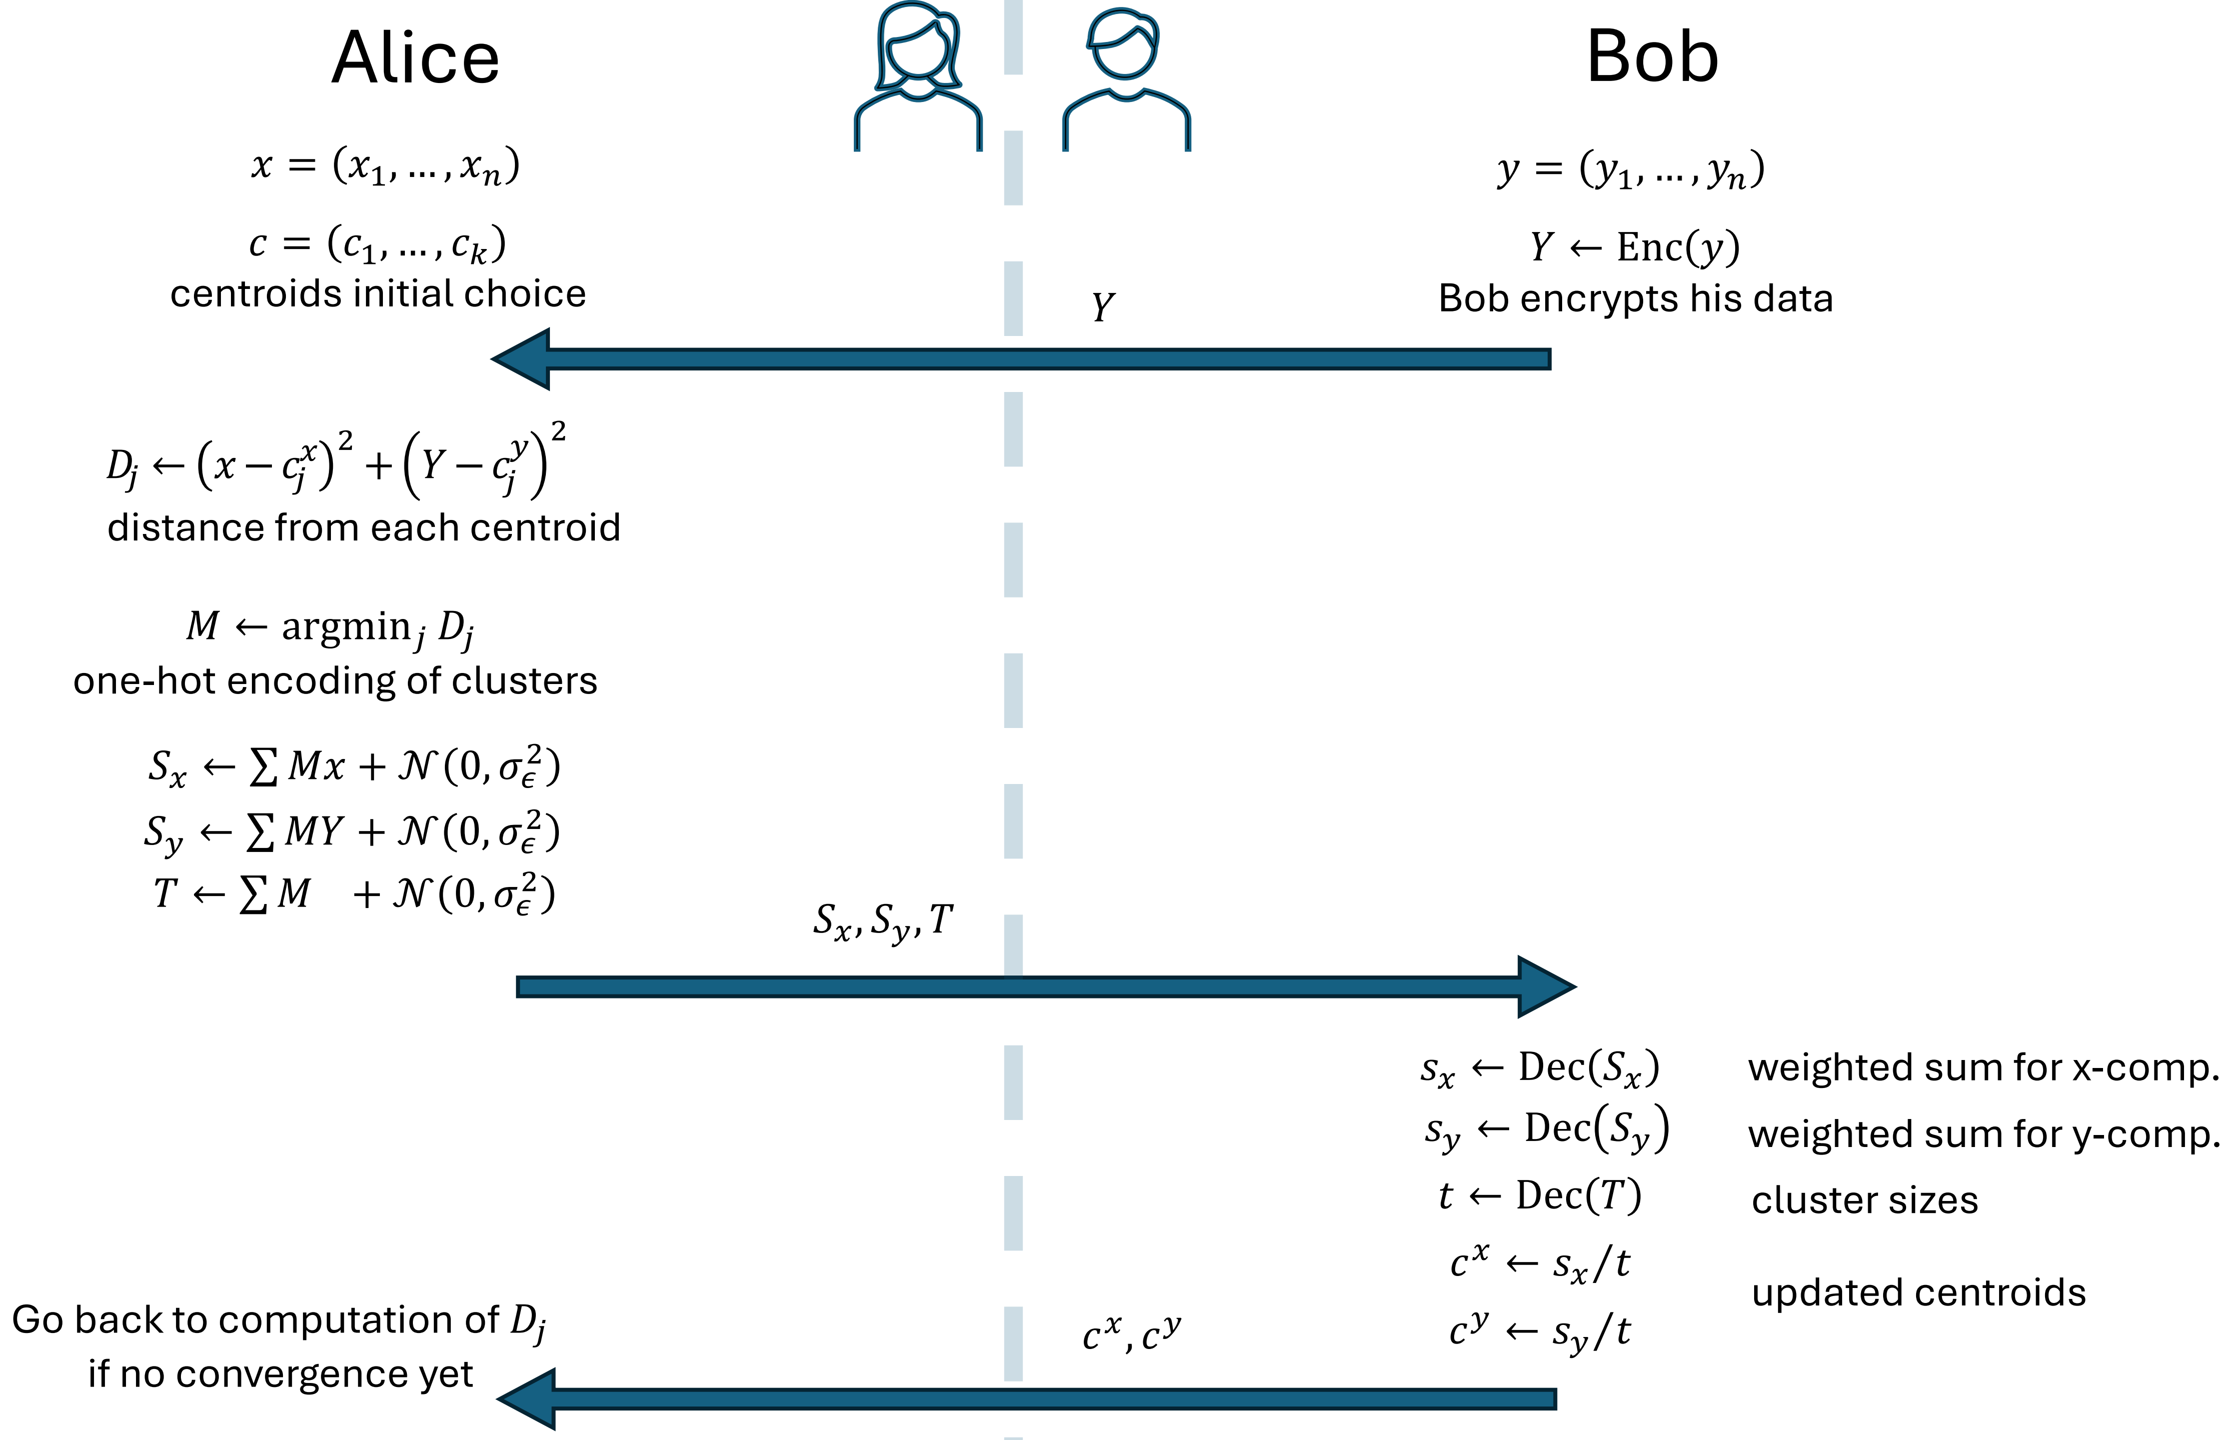

Alice
Bob
$x = (x_1, \ldots, x_n)$
$c = (c_1, \ldots, c_k)$
centroids initial choice
$y = (y_1, \ldots, y_n)$
$Y \leftarrow \mathrm{Enc}(y)$
Bob encrypts his data
Y
$D_j \leftarrow (x - c_j^x)^2 + (Y - c_j^y)^2$
distance from each centroid
$M \leftarrow \mathrm{argmin}_j D_j$
one-hot encoding of clusters
$S_x \leftarrow \sum Mx + \mathcal{N}(0, \sigma_\epsilon^2)$
$S_y \leftarrow \sum MY + \mathcal{N}(0, \sigma_\epsilon^2)$
$T \leftarrow \sum M + \mathcal{N}(0, \sigma_\epsilon^2)$
S_x, S_y, T
$s_x \leftarrow \mathrm{Dec}(S_x)$
weighted sum for x-comp.
$s_y \leftarrow \mathrm{Dec}(S_y)$
weighted sum for y-comp.
$t \leftarrow \mathrm{Dec}(T)$
cluster sizes
$c^x \leftarrow s_x / t$
$c^y \leftarrow s_y / t$
updated centroids
Go back to computation of D_j
if no convergence yet
c^x, c^y

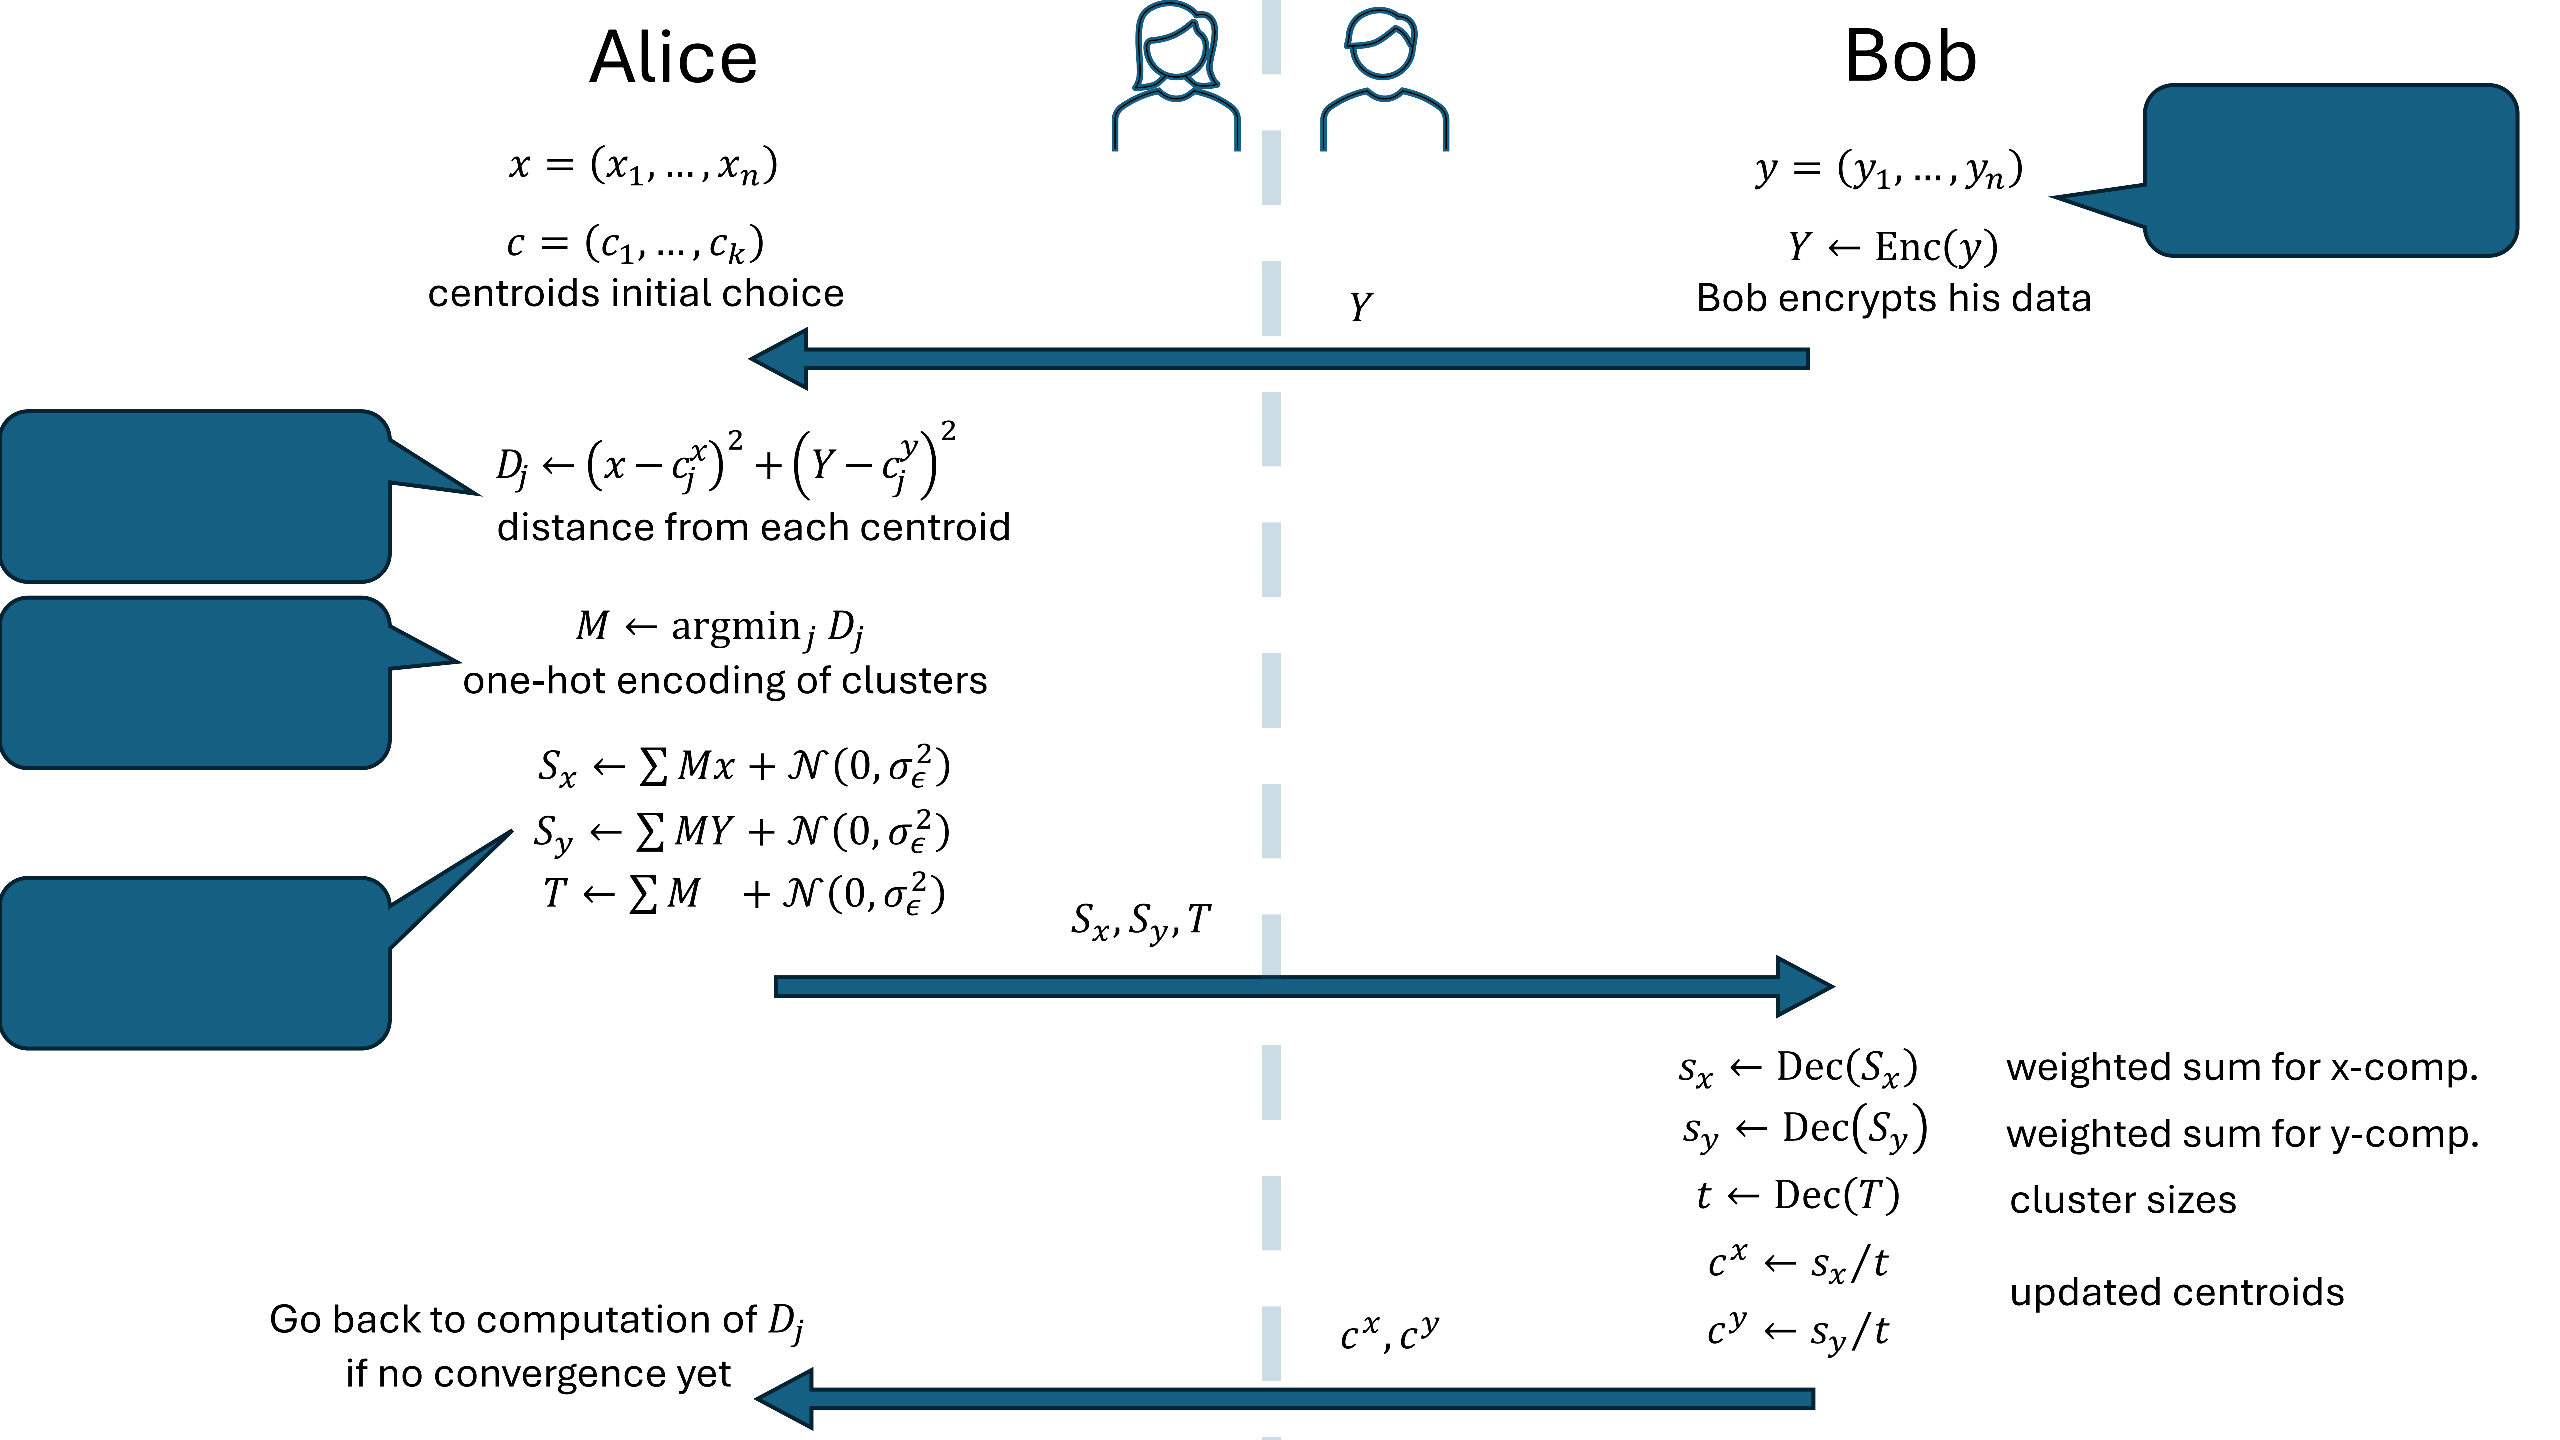

Alice
Bob
$x = (x_1, \ldots, x_n)$
$c = (c_1, \ldots, c_k)$
centroids initial choice
$y = (y_1, \ldots, y_n)$
$Y \leftarrow \mathrm{Enc}(y)$
Bob encrypts his data
Y
$D_j \leftarrow (x - c_j^x)^2 + (Y - c_j^y)^2$
distance from each centroid
$M \leftarrow \mathrm{argmin}_j\, D_j$
one-hot encoding of clusters
$S_x \leftarrow \sum Mx + \mathcal{N}(0, \sigma_\epsilon^2)$
$S_y \leftarrow \sum MY + \mathcal{N}(0, \sigma_\epsilon^2)$
$T \leftarrow \sum M + \mathcal{N}(0, \sigma_\epsilon^2)$
S_x, S_y, T
$s_x \leftarrow \mathrm{Dec}(S_x)$
weighted sum for x-comp.
$s_y \leftarrow \mathrm{Dec}(S_y)$
weighted sum for y-comp.
$t \leftarrow \mathrm{Dec}(T)$
cluster sizes
$c^x \leftarrow s_x / t$
$c^y \leftarrow s_y / t$
updated centroids
c^x, c^y
Go back to computation of D_j
if no convergence yet

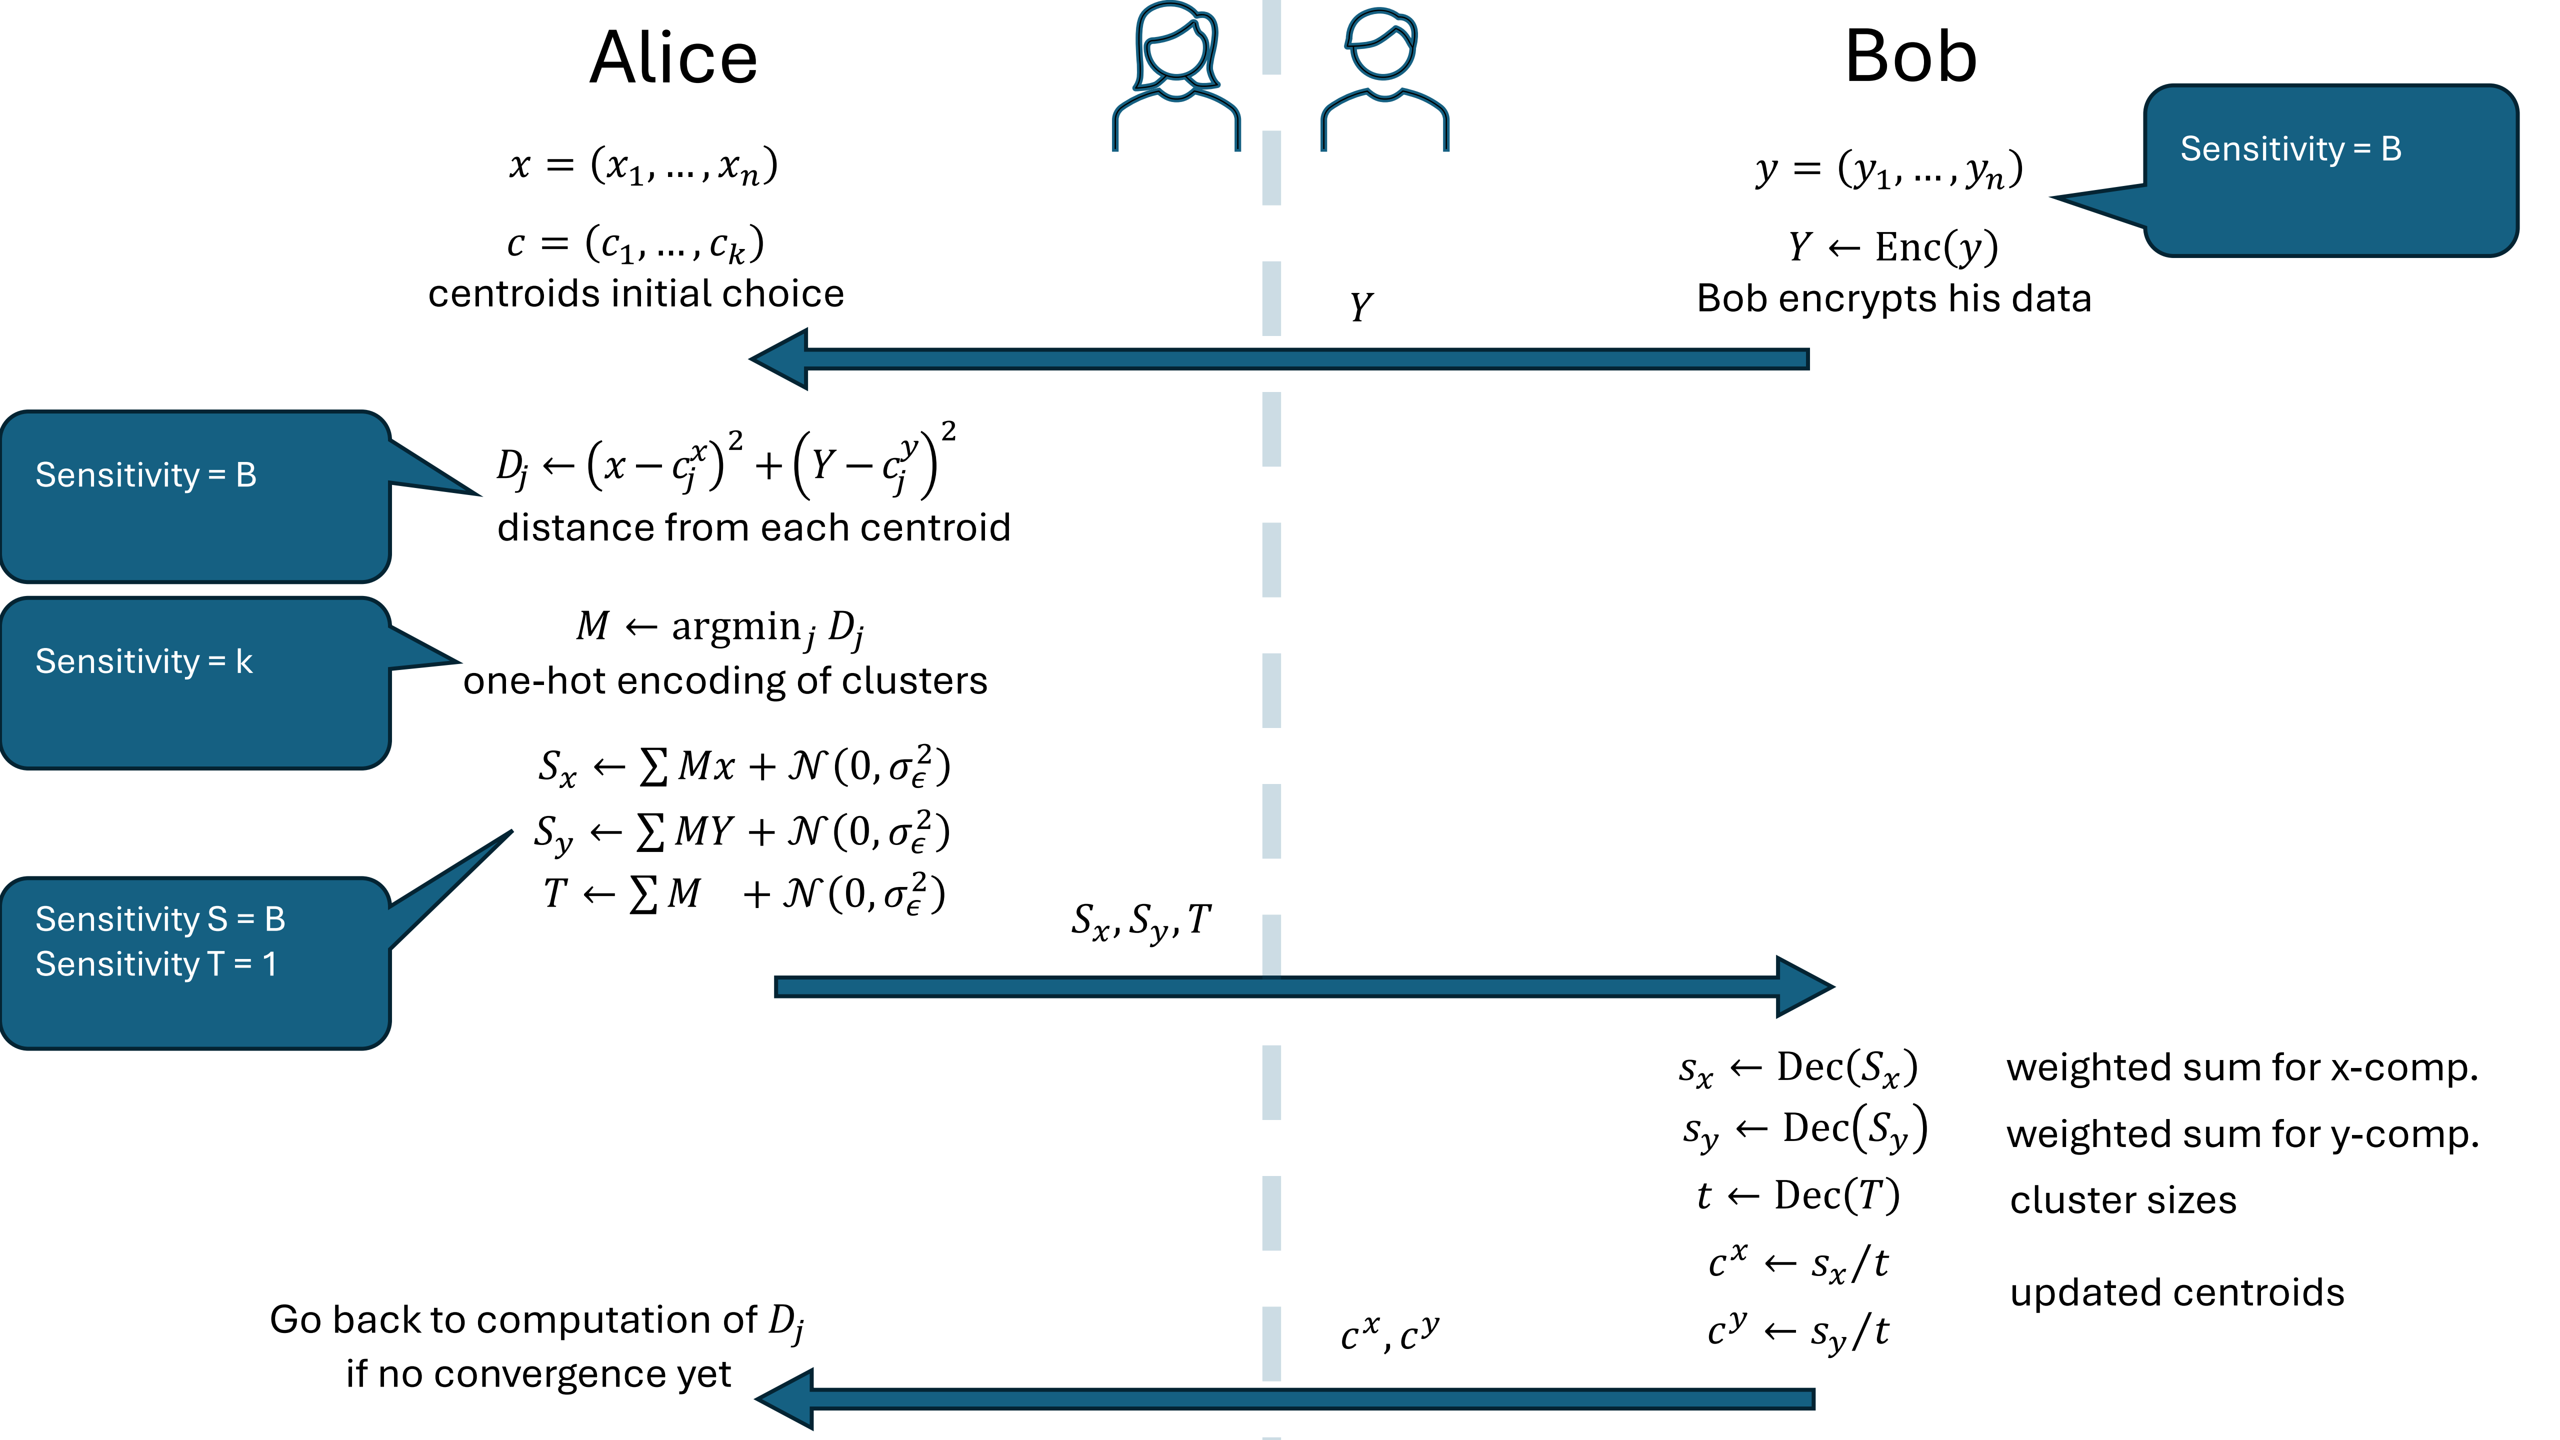
Alice
Bob
$x = (x_1, \ldots, x_n)$
$c = (c_1, \ldots, c_k)$
centroids initial choice
$y = (y_1, \ldots, y_n)$
Sensitivity = B
$Y \leftarrow \mathrm{Enc}(y)$
Bob encrypts his data
Y
Sensitivity = B
$D_j \leftarrow (x - c_j^x)^2 + (Y - c_j^y)^2$
distance from each centroid
Sensitivity = k
$M \leftarrow \mathrm{argmin}_j \, D_j$
one-hot encoding of clusters
$S_x \leftarrow \sum Mx + \mathcal{N}(0, \sigma_\epsilon^2)$
$S_y \leftarrow \sum MY + \mathcal{N}(0, \sigma_\epsilon^2)$
$T \leftarrow \sum M \quad + \mathcal{N}(0, \sigma_\epsilon^2)$
Sensitivity S = B
Sensitivity T = 1
S_x, S_y, T
$s_x \leftarrow \mathrm{Dec}(S_x)$ weighted sum for x-comp.
$s_y \leftarrow \mathrm{Dec}(S_y)$ weighted sum for y-comp.
$t \leftarrow \mathrm{Dec}(T)$ cluster sizes
$c^x \leftarrow s_x / t$
$c^y \leftarrow s_y / t$
updated centroids
c^x, c^y
Go back to computation of D_j
if no convergence yet

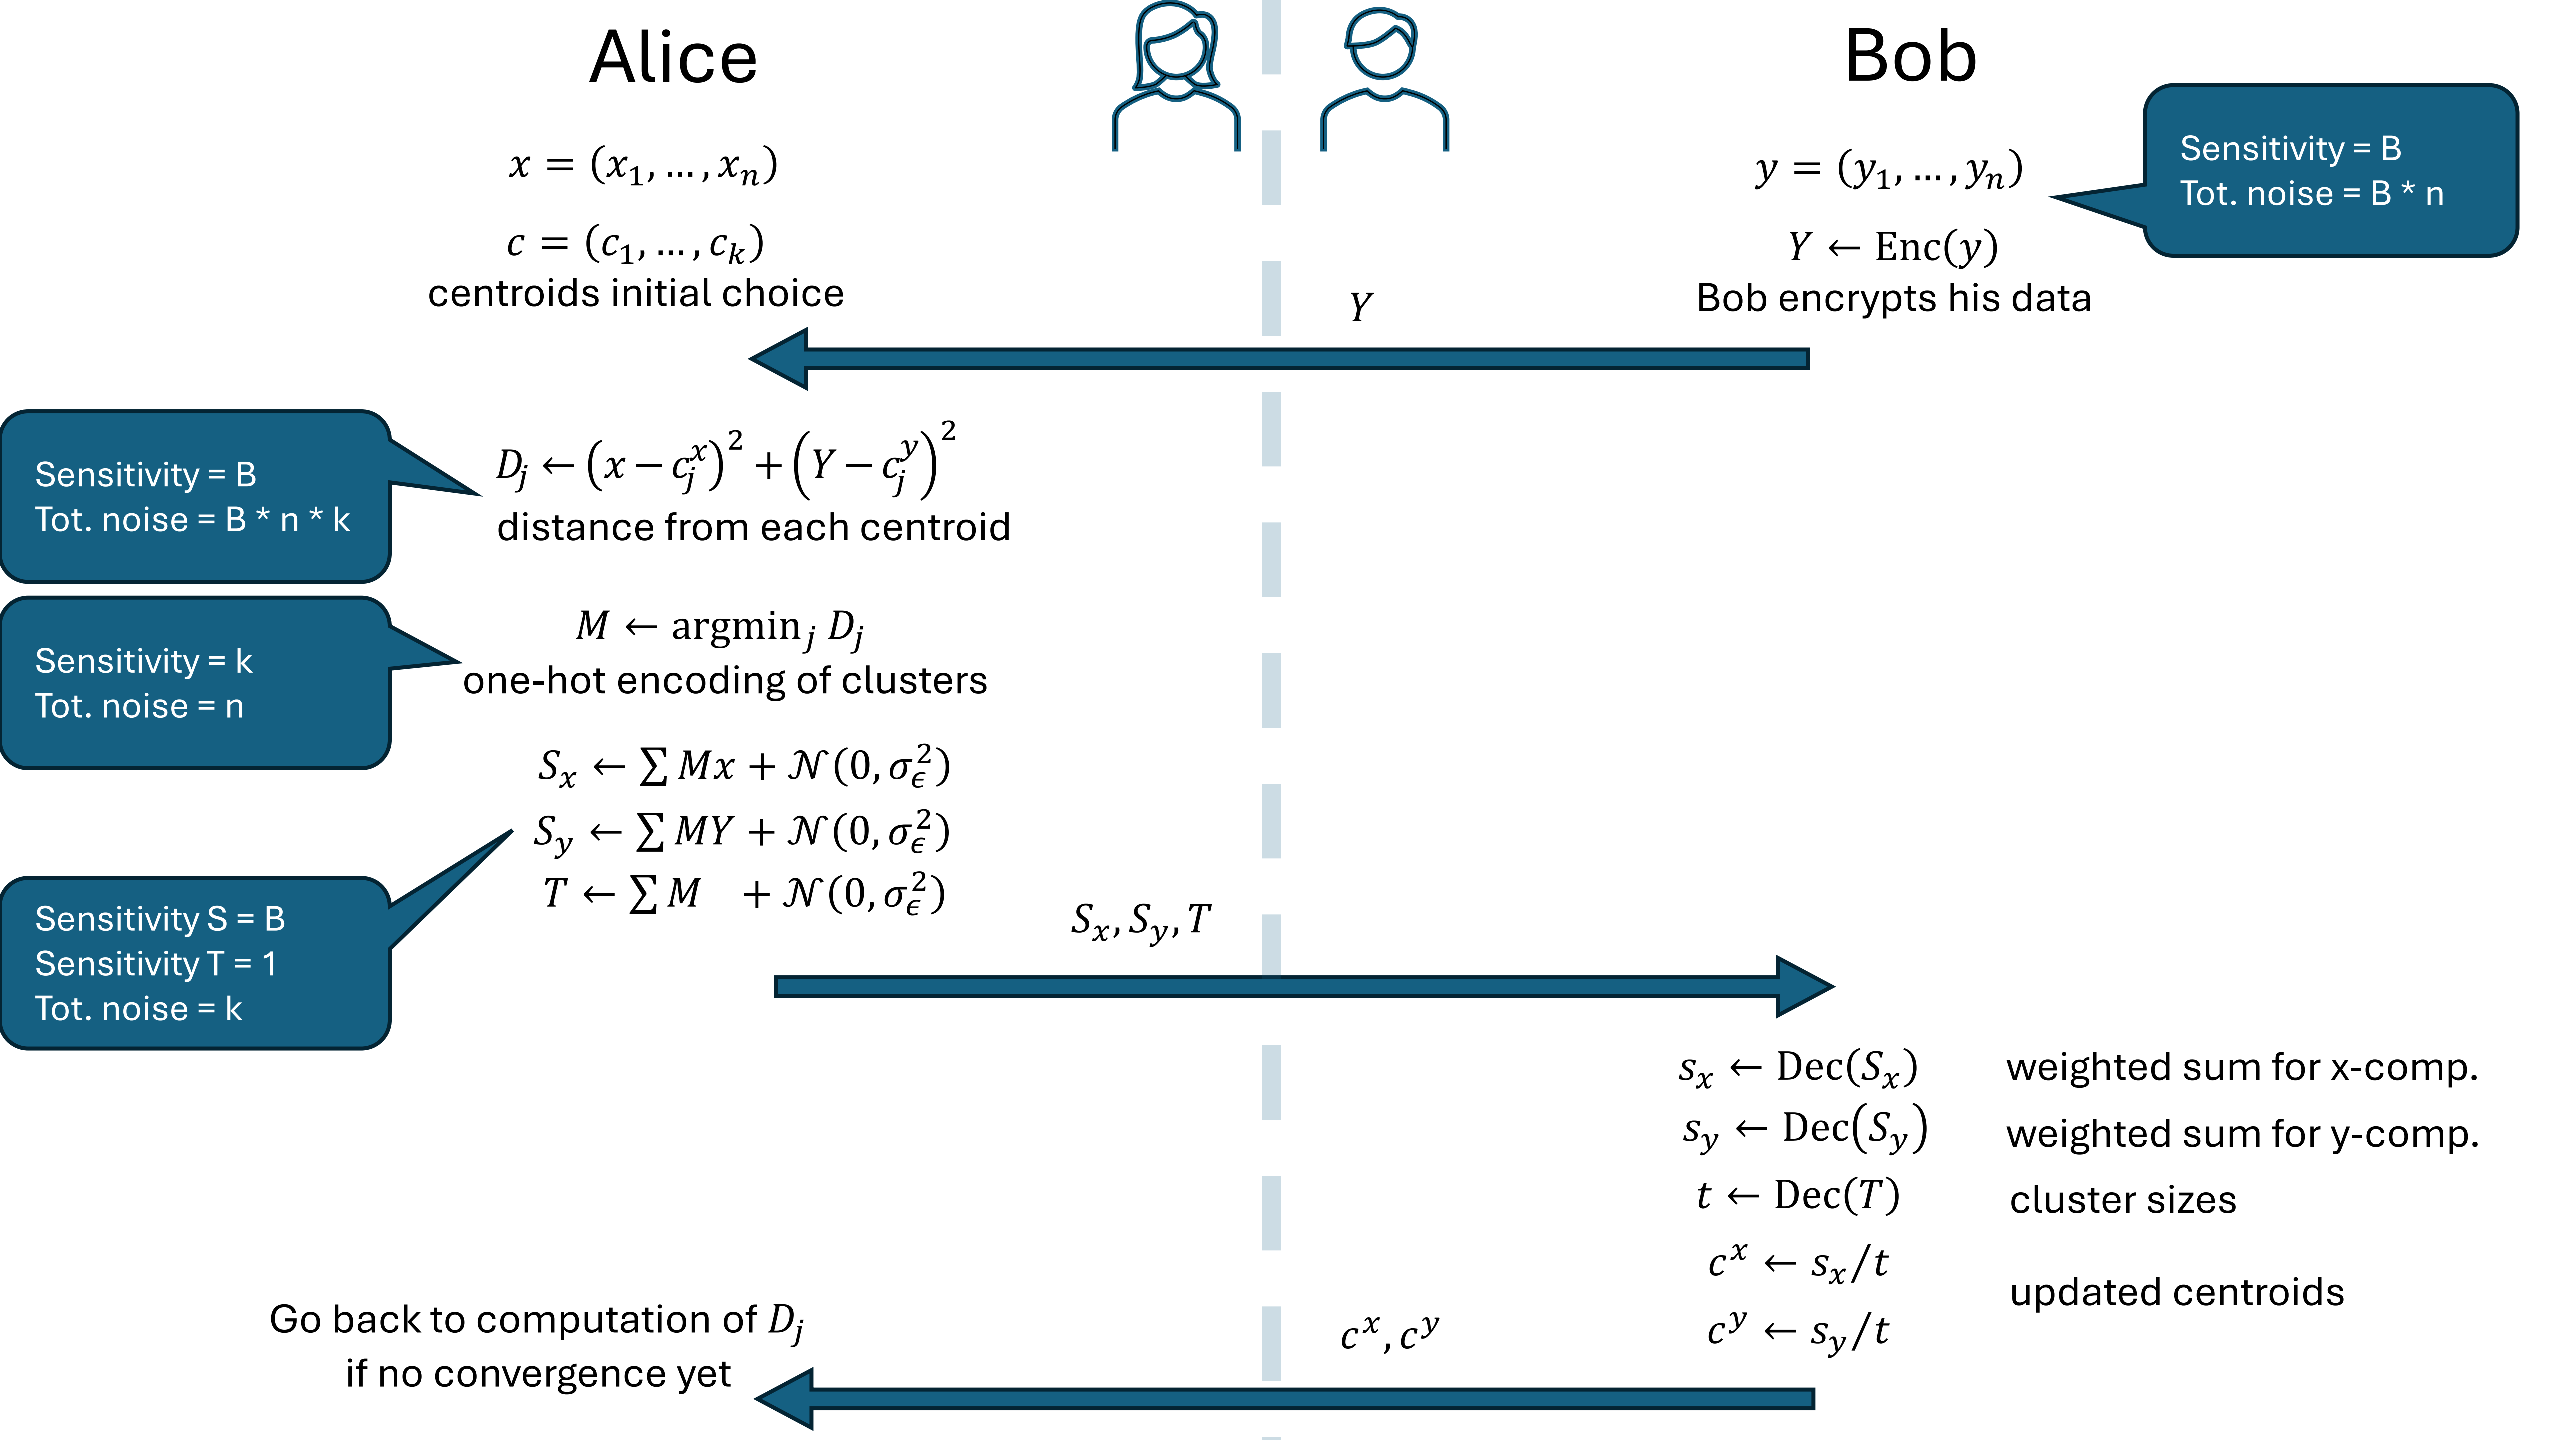
Alice
Bob
$x = (x_1, \dots, x_n)$
$c = (c_1, \dots, c_k)$
centroids initial choice
$y = (y_1, \dots, y_n)$
$Y \leftarrow \mathrm{Enc}(y)$
Bob encrypts his data
Sensitivity = B
Tot. noise = B * n
Y
$D_j \leftarrow (x - c_j^x)^2 + (Y - c_j^y)^2$
distance from each centroid
Sensitivity = B
Tot. noise = B * n * k
$M \leftarrow \mathrm{argmin}_j\, D_j$
one-hot encoding of clusters
Sensitivity = k
Tot. noise = n
$S_x \leftarrow \sum Mx + \mathcal{N}(0, \sigma_\epsilon^2)$
$S_y \leftarrow \sum MY + \mathcal{N}(0, \sigma_\epsilon^2)$
$T \leftarrow \sum M \quad + \mathcal{N}(0, \sigma_\epsilon^2)$
Sensitivity S = B
Sensitivity T = 1
Tot. noise = k
S_x, S_y, T
$s_x \leftarrow \mathrm{Dec}(S_x)$
weighted sum for x-comp.
$s_y \leftarrow \mathrm{Dec}(S_y)$
weighted sum for y-comp.
$t \leftarrow \mathrm{Dec}(T)$
cluster sizes
$c^x \leftarrow s_x / t$
$c^y \leftarrow s_y / t$
updated centroids
c^x, c^y
Go back to computation of D_j
if no convergence yet