Understanding Information Disclosure from Secure Computation Output

Analytical and Data-Driven Analysis

Alessandro Baccarini, PhD

abaccarini@proton.me
abaccarini.github.io

March 11, 2025

Outline

Motivation

Information disclosure analysis

- Framework for quantifying disclosure
- Case study: average salary computation
- Advanced statistical functions

Conclusions and future work

(Secure) multi-party computation, in a nutshell

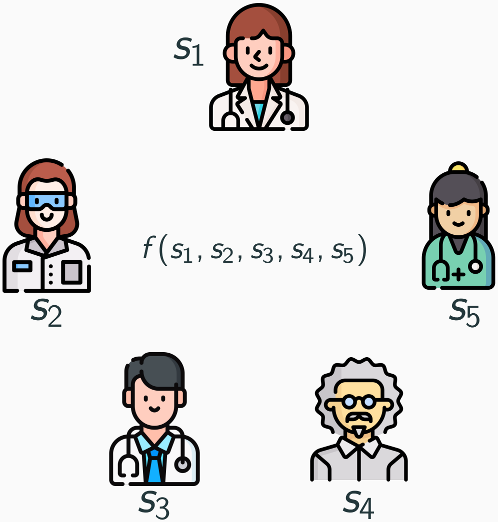

Multi-party computation (MPC)

Multiple participants **jointly** evaluating an **arbitrary** function on private inputs while revealing only the output(s).

- FHE, garbled circuits, secret sharing
- Variety of practical applications

(Secure) multi-party computation, in a nutshell

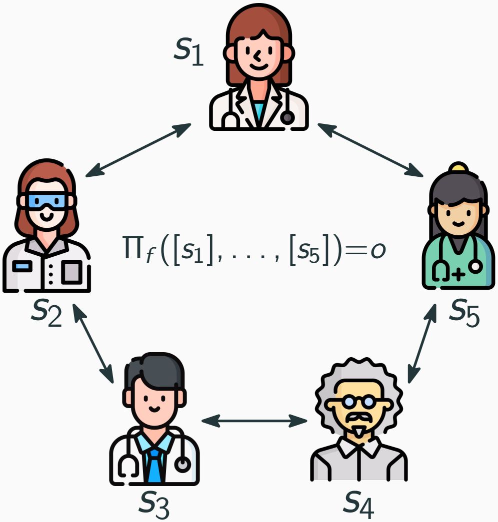

Multi-party computation (MPC)

Multiple participants **jointly** evaluating an **arbitrary** function on private inputs while revealing only the output(s).

- FHE, garbled circuits, secret sharing
- Variety of practical applications

(Secure) multi-party computation, in a nutshell

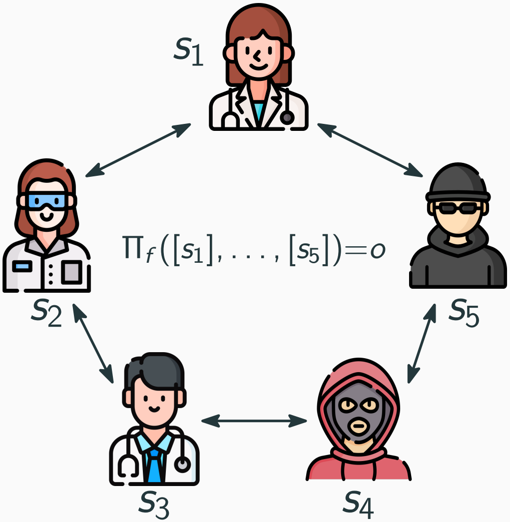

Multi-party computation (MPC)

Multiple participants **jointly** evaluating an **arbitrary** function on private inputs while revealing only the output(s).

- FHE, garbled circuits, secret sharing
- Variety of practical applications

What do we *really* mean by "secure"?

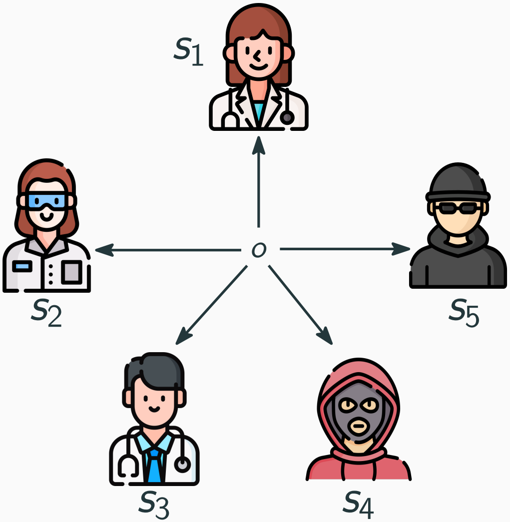

– No information disclosed throughout computation, **other than the output**

What do we *really* mean by "secure"?

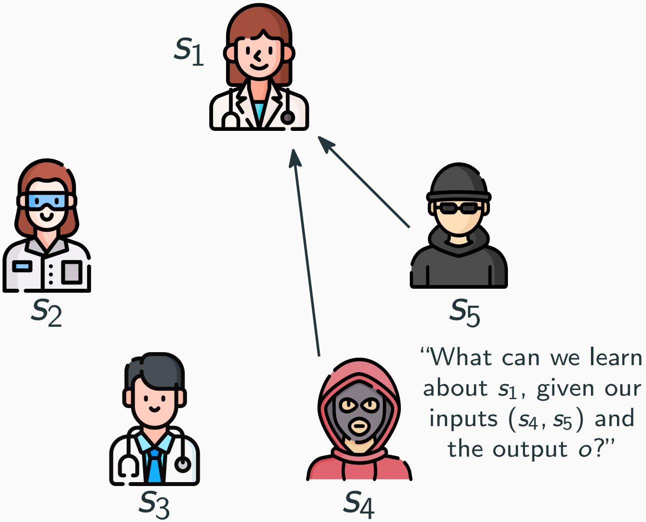

- No information disclosed throughout computation, **other than the output**
- But does the **output itself** contain sensitive information?
- Can we **quantify** this disclosure in a meaningful way?

Information disclosure analysis [Bac24, Part II]

- Develop an information-theoretic approach to measure disclosure
- Apply technique to a practically significant function (the **average**)
- Extend analysis to complex statistical functions
- Determine and apply appropriate mitigation strategies

Function f evaluated on private data

Information disclosure analysis [Bac24, Part II]

- Develop an information-theoretic approach to measure disclosure
- Apply technique to a practically significant function (the **average**)
- Extend analysis to complex statistical functions
- Determine and apply appropriate mitigation strategies

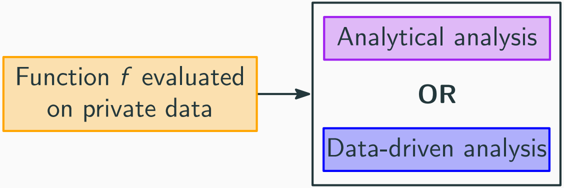

Information disclosure analysis [Bac24, Part II]

- Develop an information-theoretic approach to measure disclosure
- Apply technique to a practically significant function (the **average**)
- Extend analysis to complex statistical functions
- Determine and apply appropriate mitigation strategies

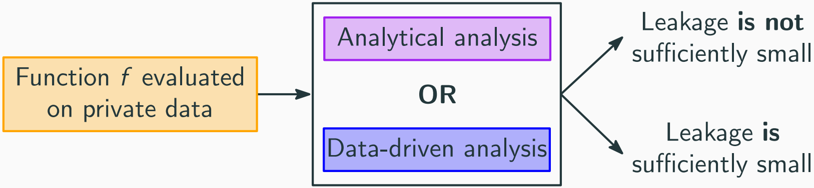

Information disclosure analysis [Bac24, Part II]

- Develop an information-theoretic approach to measure disclosure
- Apply technique to a practically significant function (the **average**)
- Extend analysis to complex statistical functions
- Determine and apply appropriate mitigation strategies

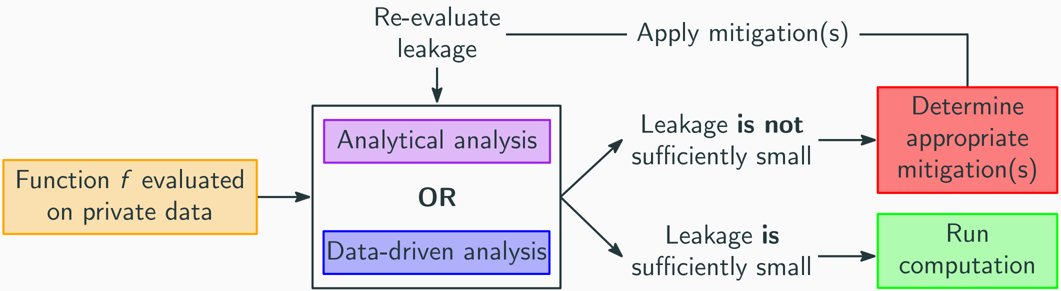

Formal setting

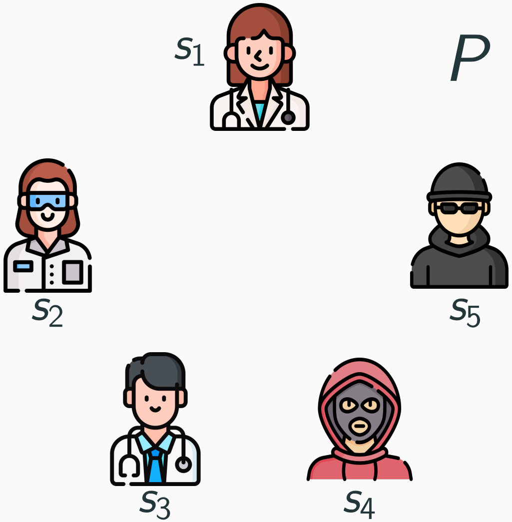

- How do we **distinguish** participants from each other?
- Partition parties P into:

Formal setting

- How do we **distinguish** participants from each other?
- Partition parties P into:
  - attackers A

Formal setting

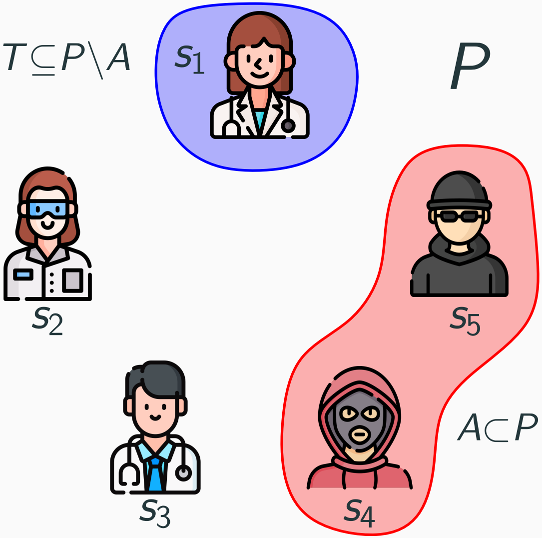

- How do we **distinguish** participants from each other?
- Partition parties P into:
  - attackers A
  - targets T

Formal setting

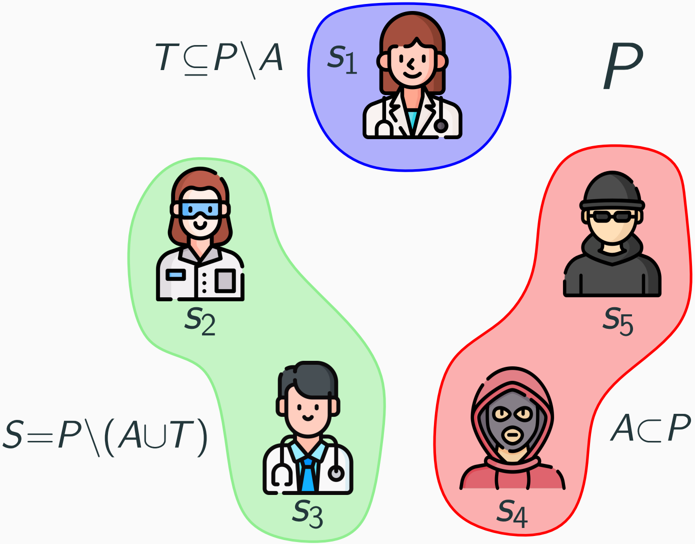

- How do we **distinguish** participants from each other?
- Partition parties P into:
  - attackers A
  - targets T
  - spectators S

Metric?

– Model participant i's inputs by a **random variable** X_{P_i} $(\mathbf{X}_P = (X_{P_1}, \ldots, X_{P_m}))$

Metric?

- Model participant i's inputs by a **random variable** X_{P_i} $(\mathbf{X}_P = (X_{P_1}, \ldots, X_{P_m}))$
- How to **measure** the **information** disclosed by the output?

Metric?

- Model participant i's inputs by a **random variable** X_{P_i} $\quad (\mathbf{X}_P = (X_{P_1}, \ldots, X_{P_m}))$
- How to **measure** the **information** disclosed by the output?

Entropy!

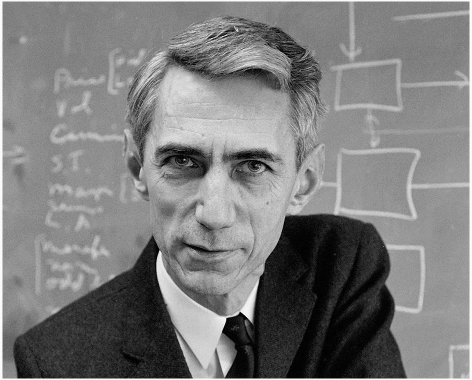

C. Shannon. Photo: Alfred Eisenstaedt, 1963

Metric?

- Model participant i's inputs by a **random variable** X_{P_i} $(\mathbf{X}_P = (X_{P_1}, \ldots, X_{P_m}))$
- How to **measure** the **information** disclosed by the output?

Entropy!

Shannon (discrete):

$$H(X) = -\sum_{x \in \mathcal{X}} \Pr(X = x) \log \Pr(X = x)$$

Differential (continuous):

$$h(X) = -\int_{\mathcal{X}} f(x) \log f(x) \, dx$$

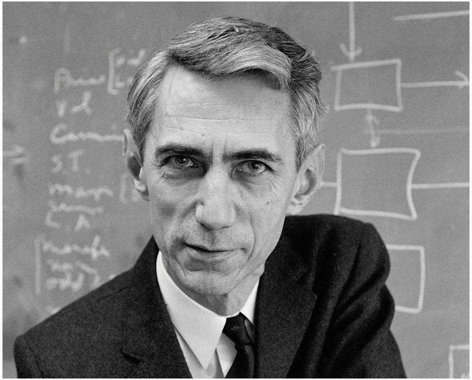

C. Shannon. Photo: Alfred Eisenstaedt, 1963

Putting everything together

- Attackers $\mathbf{X}_A$, targets $\mathbf{X}_T$, and spectators $\mathbf{X}_S$
- Treat the **output** as a random variable: $f(\mathbf{X}_A, \mathbf{X}_T, \mathbf{X}_S) = O$

Putting everything together

- Attackers $\mathbf{X}_A$, targets $\mathbf{X}_T$, and spectators $\mathbf{X}_S$
- Treat the **output** as a random variable: $f(\mathbf{X}_A, \mathbf{X}_T, \mathbf{X}_S) = O$

Attackers' weighted average entropy [AH17]

$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) \implies$ "how much information A learns about the target, given $\mathbf{x}_A$ and O"

Absolute entropy loss [BBZ24a; BBZ24b]

$H(\mathbf{X}_T) - H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) \implies$ "the total amount of information disclosed about the target, given $\mathbf{x}_A$ and O"

Absolute entropy loss $\iff$ **mutual information** between $\mathbf{X}_T$ and O (conditioned on $\mathbf{X}_A = \mathbf{x}_A$)

Case study: the average salary computation

- 2016 Boston gender pay gap survey
- Analyzed the **private** wages based on gender and race **using MPC**
- **Average salary computation**

Mayor Walsh & Boston Women's Workforce Council Release 2016 Gender Wage Gap Report; New Partnership with BU

Thursday, January 5, 2017 – Mayor Martin J. Walsh and the Boston Women's Workforce Council (BWWC) released the 2016 gender wage report and announced a new academic partnership with Boston University, where the BWWC will now be hosted within the BU Hariri Institute for Computing.

Source: Boston University, 2017

Case study: the average salary computation

- 2016 Boston gender pay gap survey
- Analyzed the **private** wages based on gender and race **using MPC**
- **Average salary computation**

Mayor Walsh & Boston Women's Workforce Council Release 2016 Gender Wage Gap Report; New Partnership with BU

Thursday, January 5, 2017 – Mayor Martin J. Walsh and the Boston Women's Workforce Council (BWWC) released the 2016 gender wage report and announced a new academic partnership with Boston University, where the BWWC will now be hosted within the BU Hariri Institute for Computing.

Source: Boston University, 2017

However, the average reduces to a **sum**: $f_\mu(\mathbf{x}) = \frac{1}{n}(x_1 + \cdots + x_n) \longrightarrow x_1 + \cdots + x_n$

Case study: the average salary computation

- 2016 Boston gender pay gap survey
- Analyzed the **private** wages based on gender and race **using MPC**
- **Average salary computation**

Mayor Walsh & Boston Women's Workforce Council Release 2016 Gender Wage Gap Report; New Partnership with BU

Thursday, January 5, 2017 – Mayor Martin J. Walsh and the Boston Women's Workforce Council (BWWC) released the 2016 gender wage report and announced a new academic partnership with Boston University, where the BWWC will now be hosted within the BU Hariri Institute for Computing.

Source: Boston University, 2017

However, the average reduces to a **sum**: $f_\mu(\mathbf{x}) = \frac{1}{n}(x_1 + \cdots + x_n) \longrightarrow x_1 + \cdots + x_n$

$$O = \sum_i X_{T_i} + \sum_j X_{A_j} + \sum_k X_{S_k}$$

Optimal attacker strategy

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \sum_i X_{T_i} + \sum_k X_{S_k})$$

Optimal attacker strategy

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \sum_i X_{T_i} + \sum_k X_{S_k})$$

– Intuition: an adversary can "remove" their influence

s_5

s_4

$$o = \frac{1}{5}(s_1 + s_2 + s_3 + s_4 + s_5)$$

Optimal attacker strategy

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \sum_i X_{T_i} + \sum_k X_{S_k})$$

– Intuition: an adversary can "remove" their influence

s_5

s_4

$$o = \cancel{1/5}(s_1 + s_2 + s_3 + s_4 + s_5)$$

Optimal attacker strategy

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \sum_i X_{T_i} + \sum_k X_{S_k})$$

– Intuition: an adversary can "remove" their influence

s_5

s_4

$o = s_1 + s_2 + s_3 + s_4 + s_5$

Optimal attacker strategy

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \textstyle\sum_i X_{T_i} + \sum_k X_{S_k})$$

– Intuition: an adversary can "remove" their influence

s_5

s_4

$$o = s_1 + s_2 + s_3 + s_4 + s_5$$

Optimal attacker strategy

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \textstyle\sum_i X_{T_i} + \sum_k X_{S_k})$$

– Intuition: an adversary can "remove" their influence

s_5

s_4

$$o = s_1 + s_2 + s_3$$

Claim

The information disclosure is **independent** of the attackers' input(s):

$$H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O) = H(\mathbf{X}_T \mid \sum_i X_{T_i} + \sum_k X_{S_k})$$

- Intuition: an adversary can "remove" their influence
- May not always be the case (depending on f)

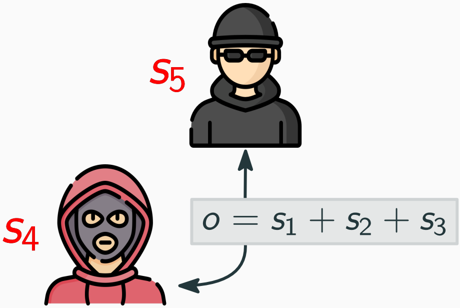

Single evaluation

- Model inputs by common distributions:
  - Poisson
  - Uniform
  - Gaussian
  - Log-normal [Cao+22]

Single evaluation

- Model inputs by common distributions:
  - Poisson
  - Uniform
  - **Gaussian**
  - Log-normal [Cao+22]
- For a single evaluation, information disclosure is **independent** of
  - the distribution parameters
  - the **distribution itself**
- Disclosure is proportional to the number of **spectators**

CLT

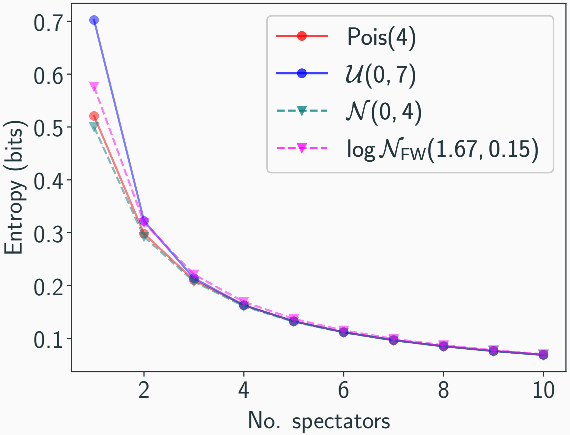

Figure 1: Absolute entropy loss (lower is better)

“What about differential privacy?”

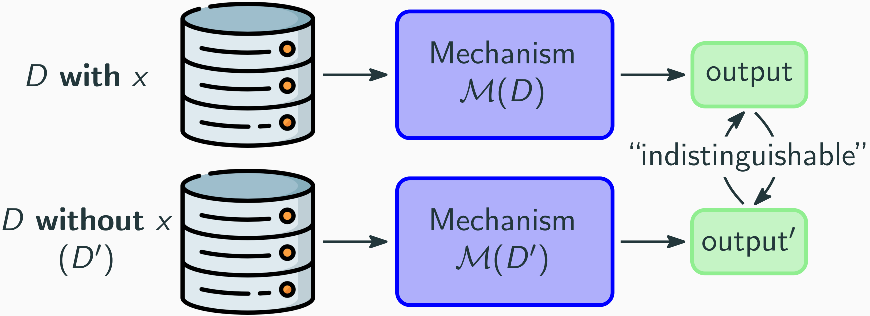

“What about differential privacy?”

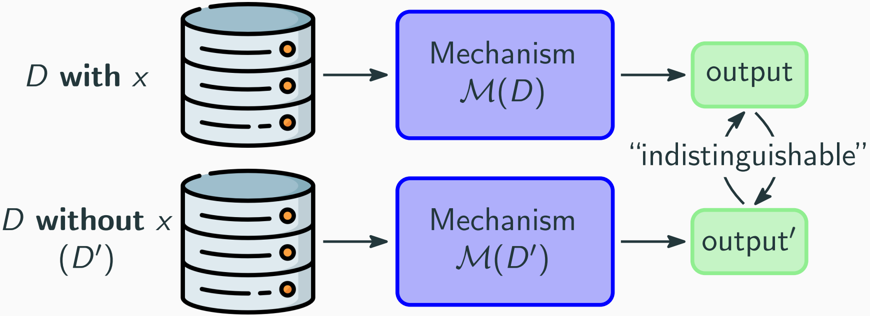

– Useful for large databases (e.g., $n \geq 10{,}000$) . . .

"What about differential privacy?"

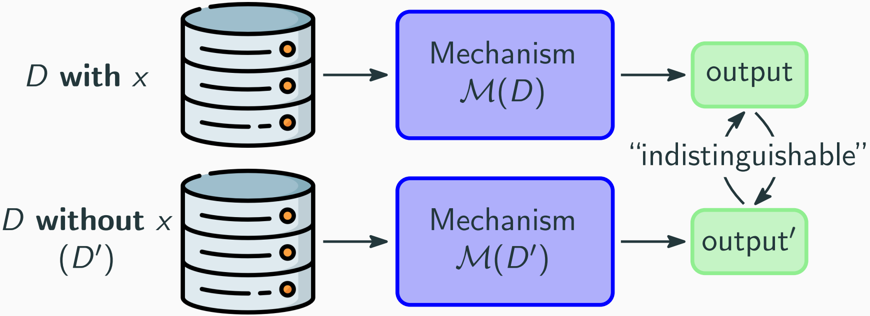

- Useful for large databases (e.g., $n \geq 10{,}000$) ...
- ... but **destroys** the utility of the result for small n (up to 100% error!)

“What about differential privacy?”

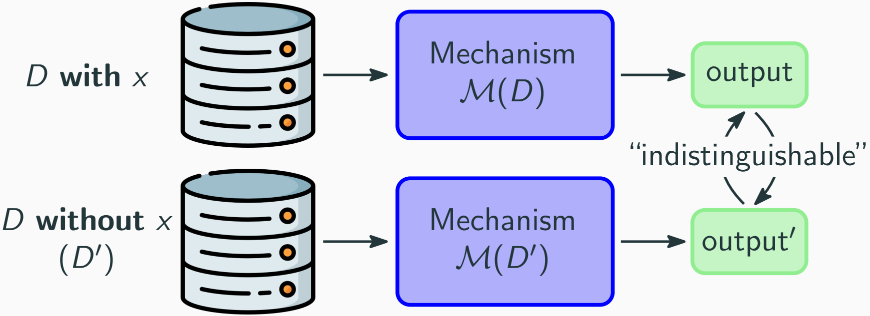

- Useful for large databases (e.g., $n \geq 10{,}000$) . . .
- . . . but **destroys** the utility of the result for small n (up to 100% error!)
- Our goal: first **determine** if a function discloses too much information
- We have an effective means of lowering disclosure for the **average** (increasing participants)

Computing the average salary *twice*

- First wage gap study was successful
- Conducted again the following year with an extended set of participants

Mayor Walsh & BWWC Release 2017 Wage Gap Report

The Boston Women's Workforce Council released its 2017 report this morning, which uses real employer wage information to assess the pay gap in Boston. The 2017 report

Source: Boston University, 2018

Computing the average salary *twice*

- First wage gap study was successful
- Conducted again the following year with an extended set of participants
- Combination of new **and** old parties participating in the computation
- Spectators present in the first, second, and both evaluation(s)


Mayor Walsh & BWWC Release 2017 Wage Gap Report

The Boston Women's Workforce Council released its 2017 report this morning, which uses real employer wage information to assess the pay gap in Boston. The 2017 report

BOSTON WOMEN'S WORKFORCE COUNCIL REPORT 2017

Source: Boston University, 2018

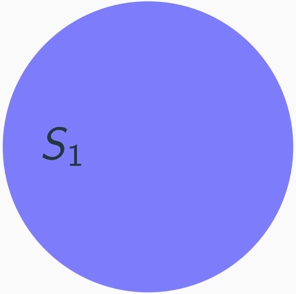

Computing the average salary *twice*

- First wage gap study was successful
- Conducted again the following year with an extended set of participants
- Combination of new **and** old parties participating in the computation
- Spectators present in the first, second, and both evaluation(s)

Mayor Walsh & BWWC Release 2017 Wage Gap Report

The Boston Women's Workforce Council released its 2017 report this morning, which uses real employer wage information to assess the pay gap in Boston. The 2017 report

BOSTON WOMEN'S WORKFORCE COUNCIL REPORT 2017

Source: Boston University, 2018

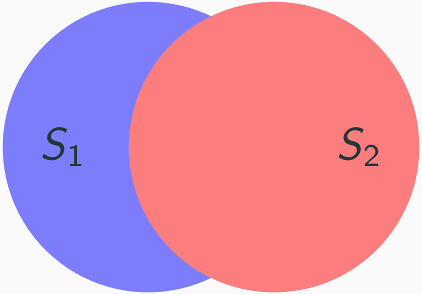

- First wage gap study was successful
- Conducted again the following year with an extended set of participants
- Combination of new **and** old parties participating in the computation
- Spectators present in the first, second, and both evaluation(s)

Mayor Walsh & BWWC Release 2017 Wage Gap Report

The Boston Women's Workforce Council released its 2017 report this morning, which uses real employer wage information to assess the pay gap in Boston. The 2017 report

Source: Boston University, 2018

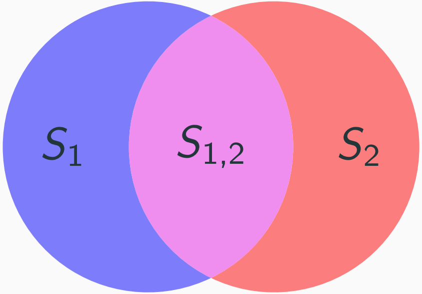

An interesting question

$$f(s_1, s_2, s_3, s_4, s_5) = o$$

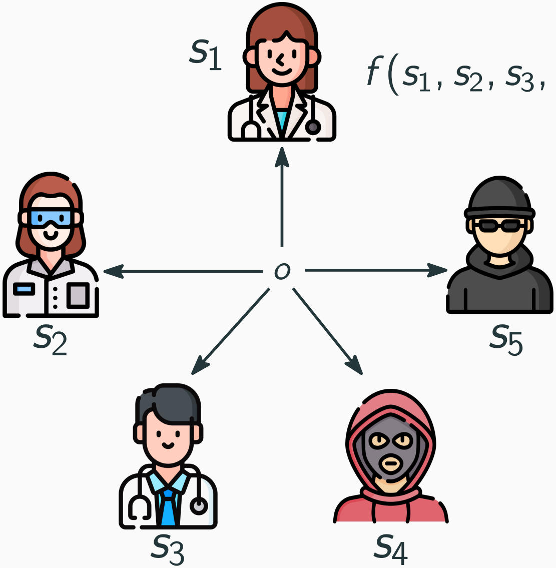

An interesting question

"What happens if everyone else participates *again*, but without me?"

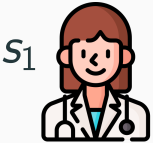

$$f(s_2, s_3, s_4, s_5) = o'$$

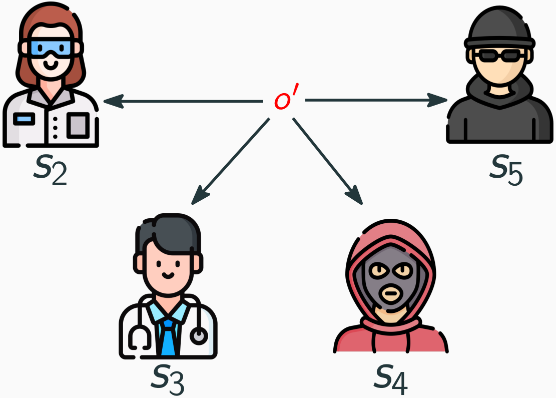

Target participates in one or both evaluations

– Vary the **ratio** of shared spectators $S_{1,2}$ to the (fixed) total number of spectators

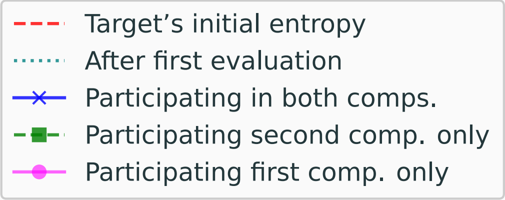

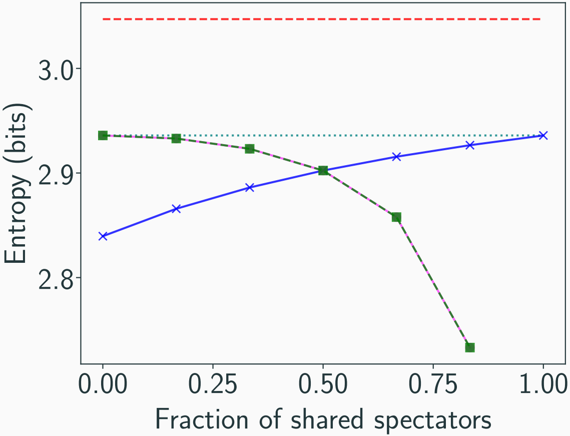

Figure 2: Conditional entropies, 6 total spectators

Target participates in one or both evaluations

- Vary the **ratio** of shared spectators $S_{1,2}$ to the (fixed) total number of spectators
- Largest protection at **50% overlap**
- Undesirable disclosure at extrema

Target's initial entropy
After first evaluation
Participating in both comps.
Participating second comp. only
Participating first comp. only

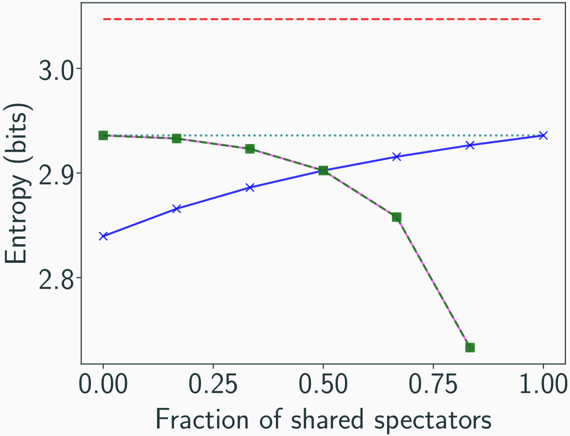

Figure 2: Conditional entropies, 6 total spectators

Next step: advanced statistical measures

What are some logical successors to the average?

Next step: advanced statistical measures

What are some logical successors to the average?

- Order statistics (max/min, median)

$$f_{\max}(\mathbf{x}) = \max_i x_i$$

- Variability measures (variance)

$$f_{\sigma^2}(\mathbf{x}) = \frac{1}{n}\sum_i (x_i - f_\mu(\mathbf{x}))^2$$

- *Multidimensional* functions

$$f_{(\mu,\sigma^2)}(\mathbf{x}) = (f_\mu(\mathbf{x}), f_{\sigma^2}(\mathbf{x}))$$

New functions $\implies$ new challenges

Prior analysis leveraged properties of **sums of RVs**, **closed-form expressions** of the entropy

Prior analysis leveraged properties of **sums of RVs**, **closed-form expressions** of the entropy

Data-driven estimators of entropy

Discrete
plug-in [Pan03]

Continuous
k-NN [GOV18]

New functions $\implies$ new challenges

Prior analysis leveraged properties of **sums of RVs**, **closed-form expressions** of the entropy

Data-driven estimators of entropy

Discrete
plug-in [Pan03]

Continuous
k-NN [GOV18]

Problem

Function could produce discrete outputs from continuous inputs, producing a "**mixture**"

New functions $\implies$ new challenges

Prior analysis leveraged properties of **sums of RVs**, **closed-form expressions** of the entropy

Data-driven estimators of entropy

Discrete
plug-in [Pan03]

Continuous
k-NN [GOV18]

Problem

Function could produce discrete outputs from continuous inputs, producing a "**mixture**"

Computer Science > Information Theory

[Submitted on 19 Sep 2017 (v1), last revised 9 Oct 2018 (this version, v3)]

Estimating Mutual Information for Discrete-Continuous Mixtures

Weihao Gao, Sreeram Kannan, Sewoong Oh, Pramod Viswanath

New functions $\Longrightarrow$ new challenges

Prior analysis leveraged properties of **sums of RVs**, **closed-form expressions** of the entropy

Data-driven estimators of entropy

Discrete
plug-in [Pan03]

Continuous
k-NN [GOV18]

Computer Science > Information Theory

[Submitted on 19 Sep 2017 (v1), last revised 9 Oct 2018 (this version, v3)]

Estimating Mutual Information for Discrete-Continuous Mixtures

Weihao Gao, Sreeram Kannan, Sewoong Oh, Pramod Viswanath

Problem

Function could produce discrete outputs from continuous inputs, producing a "**mixture**"

But **recall** (from slide 8)...

mutual information
$\Leftrightarrow$
absolute loss

Intuitive observations: maximum

Maximum

An adversary **maximizes** the information learned by **minimizing** their influence.*

*Inverse behavior for the minimum

Maximum

An adversary **maximizes** the information learned by **minimizing** their influence.*

*Inverse behavior for the minimum

- In fact, the information A learns is **bounded** by observing the output, **without participating** in $f_{\max}$

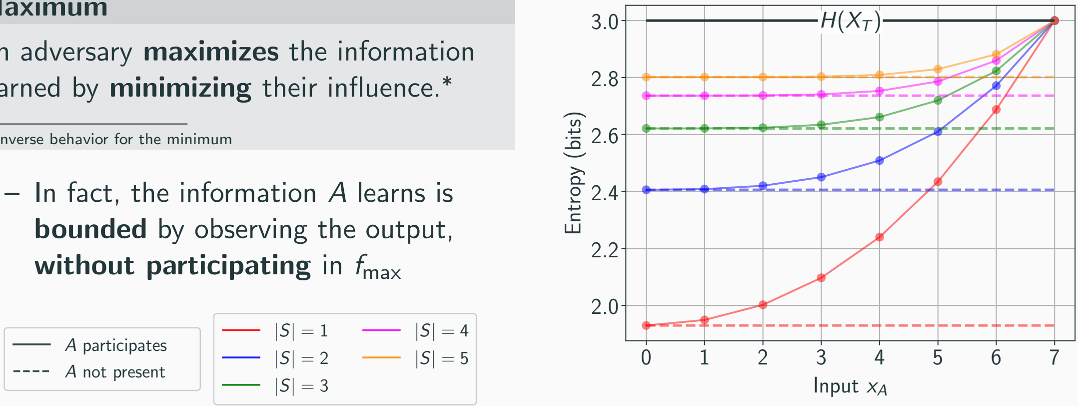

Figure 3: Uniform $\mathcal{U}(0,7)$, $H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O)$

Maximum

An adversary **maximizes** the information learned by **minimizing** their influence.*

*Inverse behavior for the minimum

- In fact, the information A learns is **bounded** by observing the output, **without participating** in $f_{\max}$

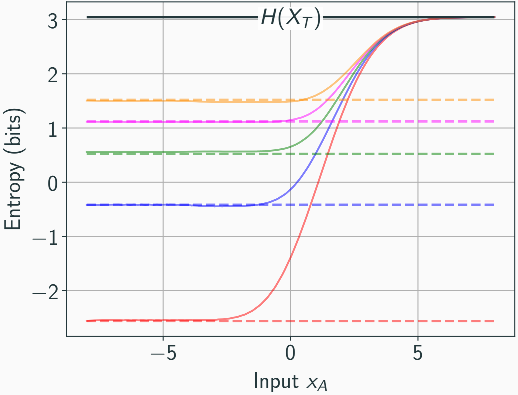

Figure 3: Normal $\mathcal{N}(0, 4.0)$, $H(\mathbf{X}_T \mid \mathbf{X}_A = \mathbf{x}_A, O)$

Surprising observations: simultaneous release

Variance and mean release

The total disclosure from **individual** function outputs f_μ and f_{σ^2} is **at least** the amount of information disclosed from a **joint release** $f_{(\mu,\sigma^2)}$?

$$H_{f_\mu} + H_{f_{\sigma^2}} \overset{?}{\geq} H_{(f_\mu, f_{\sigma^2})}$$

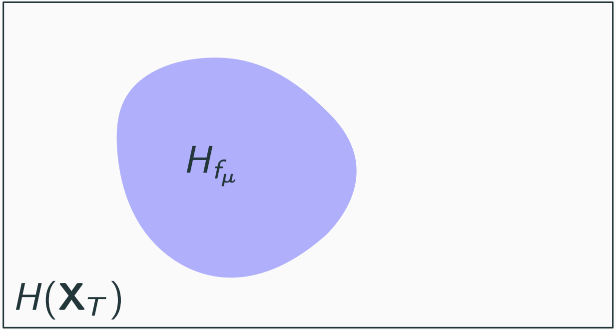

Surprising observations: simultaneous release

Variance and mean release

The total disclosure from **individual** function outputs f_μ and f_{σ^2} is **at least** the amount of information disclosed from a **joint release** $f_{(\mu,\sigma^2)}$?

$$H_{f_\mu} + H_{f_{\sigma^2}} \overset{?}{\geq} H_{(f_\mu, f_{\sigma^2})}$$

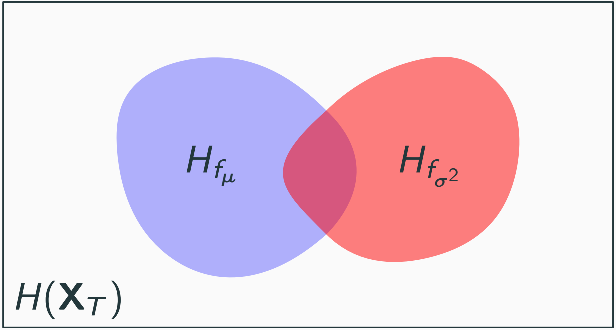

Surprising observations: simultaneous release

Variance and mean release

The total disclosure from **individual** function outputs f_μ and f_{σ^2} is **at least** the amount of information disclosed from a **joint release** $f_{(\mu,\sigma^2)}$?

$$H_{f_\mu} + H_{f_{\sigma^2}} \overset{?}{\geq} H_{(f_\mu, f_{\sigma^2})}$$

– The quantity $H_{f_\mu} + H_{f_{\sigma^2}}$ itself isn't practically significant

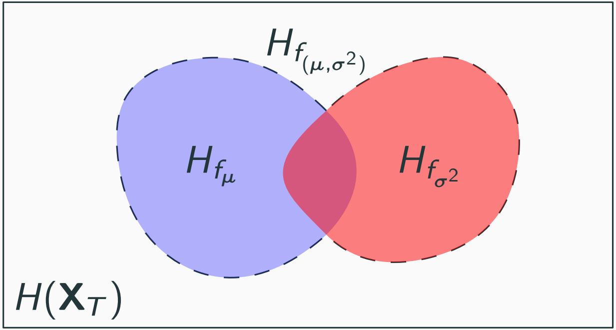

Surprising observations: simultaneous release

Variance and mean release

The total disclosure from **individual** function outputs f_μ and f_{σ^2} is **at least** the amount of information disclosed from a **joint release** $f_{(\mu,\sigma^2)}$?

$$H_{f_\mu} + H_{f_{\sigma^2}} \overset{?}{\geq} H_{(f_\mu, f_{\sigma^2})}$$

- The quantity $H_{f_\mu} + H_{f_{\sigma^2}}$ itself isn't practically significant
- **Gap** between the curves implies A can learn **more** information about the target

—●—	$H_{f_\mu} + H_{f_{\sigma^2}}$
- -×- -	$H_{f_{(\mu,\sigma^2)}}$

blue line	$\lvert S\rvert = 2$
orange line	$\lvert S\rvert = 5$

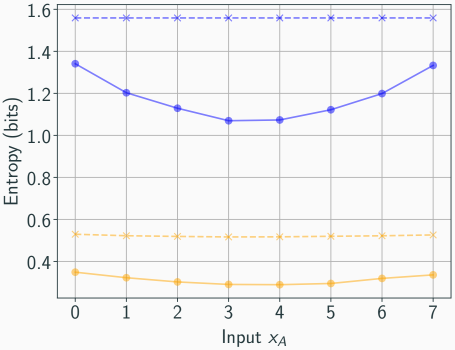

Figure 4: Abs. entropy loss, $\mathcal{U}(0, 7)$ (lower is better)

Surprising observations: simultaneous release

Variance and mean release

The total disclosure from **individual** function outputs f_μ and f_{σ^2} is **at least** the amount of information disclosed from a **joint release** $f_{(\mu,\sigma^2)}$?

$$H_{f_\mu} + H_{f_{\sigma^2}} \leq H_{(f_\mu, f_{\sigma^2})}$$

- The quantity $H_{f_\mu} + H_{f_{\sigma^2}}$ itself isn't practically significant
- **Gap** between the curves implies A can learn **more** information about the target

$H_{f_\mu} + H_{f_{\sigma^2}}$
$H_{f_{(\mu,\sigma^2)}}$

$|S| = 2$
$|S| = 5$

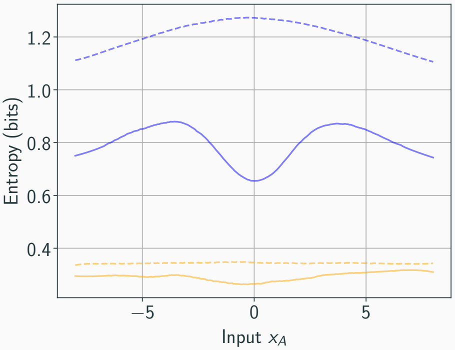

Figure 5: Abs. entropy loss, $\mathcal{N}(0, 2)$ (lower is better)

Variance and mean release

More information is revealed from the **joint release** $f_{(\mu,\sigma^2)}$ than from the **individual** function outputs f_μ and f_{σ^2}.

$$H_{f_\mu} + H_{f_{\sigma^2}} \leq H_{(f_\mu, f_{\sigma^2})}$$

- The quantity $H_{f_\mu} + H_{f_{\sigma^2}}$ itself isn't practically significant
- **Gap** between the curves implies A can learn **more** information about the target

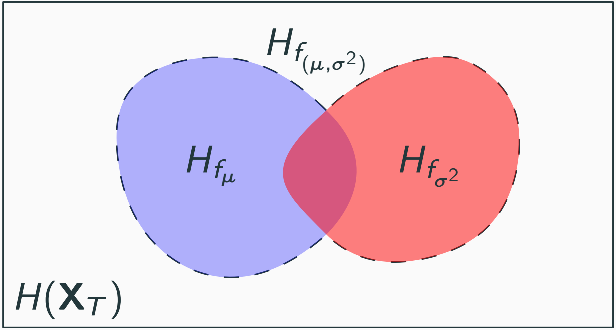

Surprising observations: simultaneous release

Variance and mean release

More information is revealed from the **joint release** $f_{(\mu,\sigma^2)}$ than from the **individual** function outputs f_μ and f_{σ^2}.

$$H_{f_\mu} + H_{f_{\sigma^2}} \leq H_{(f_\mu, f_{\sigma^2})}$$

- The quantity $H_{f_\mu} + H_{f_{\sigma^2}}$ itself isn't practically significant
- **Gap** between the curves implies A can learn **more** information about the target

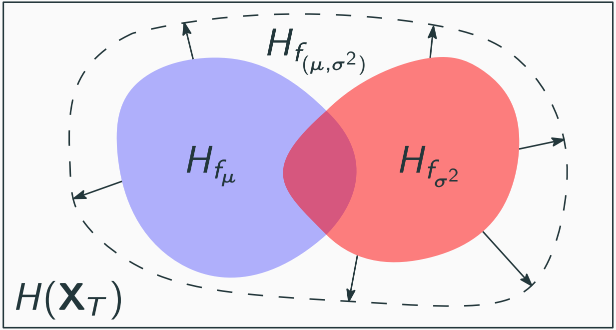

Surprising observations: simultaneous release

Variance and mean release

More information is revealed from the **joint release** $f_{(\mu,\sigma^2)}$ than from the **individual** function outputs f_μ and f_{σ^2}.

$$H_{f_\mu} + H_{f_{\sigma^2}} \leq H_{(f_\mu, f_{\sigma^2})}$$

- The quantity $H_{f_\mu} + H_{f_{\sigma^2}}$ itself isn't practically significant
- **Gap** between the curves implies A can learn **more** information about the target
- Does there exist some **nontrivial** function(s) $f_?$ that leaks the target's information entirely?

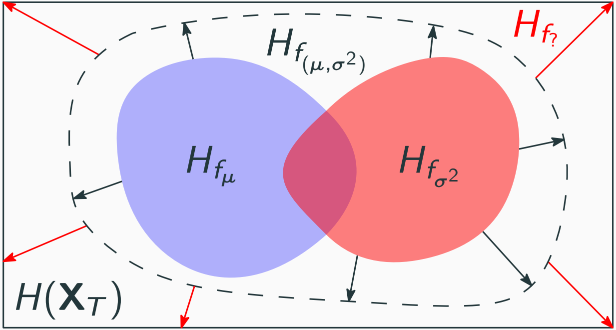

Ongoing work

- Theoretical framework from our comprehensive analysis of the average
- *Much* to learn for complex functions

Ongoing work

- Theoretical framework from our comprehensive analysis of the average
- *Much* to learn for complex functions

Further analysis of complex functions

- Derive analytical expressions the entropy
- Estimators suffer from the "curse of dimensionality"
  - Can project high-dimensional data into lower-dimensional space

Ongoing work

- Theoretical framework from our comprehensive analysis of the average
- *Much* to learn for complex functions

Further analysis of complex functions

- Derive analytical expressions the entropy
- Estimators suffer from the "curse of dimensionality"
  - Can project high-dimensional data into lower-dimensional space

Mitigation strategies

- Adding noise (DP)
- Synthetic inputs
- Modifying the function

Ongoing work

- Theoretical framework from our comprehensive analysis of the average
- *Much* to learn for complex functions

Further analysis of complex functions

- Derive analytical expressions the entropy
- Estimators suffer from the "curse of dimensionality"
  - Can project high-dimensional data into lower-dimensional space

Mitigation strategies

- Adding noise (DP)
- Synthetic inputs
- Modifying the function

Alternate metrics

- (min-, g-, cross) entropies

Conclusions

- Developed an approach for quantifying information disclosure from secure computation
- Comprehensively analyzed a practically significant function (the average)
- Applied our analysis to complex statistical measures through data-driven techniques

Conclusions

- Developed an approach for quantifying information disclosure from secure computation
- Comprehensively analyzed a practically significant function (the average)
- Applied our analysis to complex statistical measures through data-driven techniques

Other research interests

- MPC techniques and applications [Bac24, Part I], [BBY23]
- MPC compiler development
- Threshold FHE

Multi-Party Replicated Secret Sharing over a Ring with Applications to Privacy-Preserving Machine Learning

Authors: Alessandro Baccarini (University at Buffalo (SUNY)), Marina Blanton (University at Buffalo (SUNY)), Chen Yuan (University at Buffalo (SUNY))

Volume: 2023
Issue: 1
Pages: 608–626
DOI: https://doi.org/10.56553/popets-2023-0035

Privacy Enhancing Technologies Symposium (PETS), 2023

- Developed an approach for quantifying information disclosure from secure computation
- Comprehensively analyzed a practically significant function (the average)
- Applied our analysis to complex statistical measures through data-driven techniques

Other research interests

- MPC techniques and applications [Bac24, Part I], [BBY23]
- MPC compiler development
- Threshold FHE

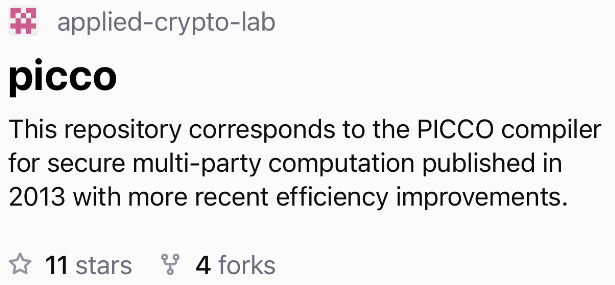

https://github.com/applied-crypto-lab/picco

Thank you!

Questions?

References

[AH17] P. Ah-Fat and M. Huth. "Secure Multi-party Computation: Information Flow of Outputs and Game Theory". In: *International Conference on Principles of Security and Trust*. 2017, pp. 71–92.

[Bac24] A. Baccarini. "New Directions in Secure Multi-Party Computation: Techniques and Information Disclosure Analysis". PhD Thesis. University at Buffalo, 2024.

[BBY23] A. Baccarini, M. Blanton, and C. Yuan. "Multi-Party Replicated Secret Sharing over a Ring with Applications to Privacy-Preserving Machine Learning". In: *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2023.1 (2023), pp. 608–626.

[BBZ24a] A. Baccarini, M. Blanton, and S. Zou. "Understanding Information Disclosure from Secure Computation Output: A Comprehensive Study of Average Salary Computation". In: *ACM Transactions on Privacy and Security (TOPS)* 28.1 (2024), pp. 1–36.

[BBZ24b] A. Baccarini, M. Blanton, and S. Zou. "Understanding Information Disclosure from Secure Computation Output: A Study of Average Salary Computation". In: *ACM CODASPY*. 2024, pp. 187–198.

[Cao+22] L. Cao, T. Tong, D. Trafimow, T. Wang, and X. Chen. "The A Priori Procedure for estimating the mean in both log-normal and gamma populations and robustness for assumption violations". In: *Methodology* 18.1 (2022), pp. 24–43.

[Gao+17] W. Gao, S. Kannan, S. Oh, and P. Viswanath. "Estimating mutual information for discrete-continuous mixtures". In: *Proceedings on Advances in Neural Information Processing Systems (NeurIPS)* 30 (2017), pp. 5988–5999.

[GOV18] W. Gao, S. Oh, and P. Viswanath. "Demystifying fixed k-nearest neighbor information estimators". In: *IEEE Transactions on Information Theory* 64.8 (2018), pp. 5629–5661.

[Pan03] L. Paninski. "Estimation of entropy and mutual information". In: *Neural Computation* 15.6 (2003), pp. 1191–1253.