Beyond Differential Privacy: Composition Theorems and Relational Logic for f-divergences between Probabilistic Programs

Gilles Barthe
Federico Olmedo

IMDEA Software Institute, Madrid, Spain

40^{th} International Colloquium on Automata, Languages and Programming
2013.09.07

f-divergences are everywhere

Cryptography

Pattern Recognition

Information Theory

Image Processing

f-divergences

Data Mining

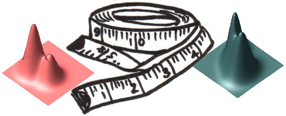

- Improving security bounds for Key-Alternating Cipher via Hellinger Distance [Steinberger:2012].

Crux of his proof: bounding the f-divergence between two probabilistic computations.

$$\Delta_f\left(c_1, c_2\right) \leq \delta$$

In this Work

Goal

Lay the foundations for reasoning about f-divergences between probabilistic programs.

- Observe that the notion of distance used to characterize differential privacy (DP) belongs to the family of f-divergences.
- Extend techniques from the DP literature to reason about arbitrary f-divergences.

Differential Privacy Primer

General Scenario

Contributor **privacy** **VS** Data mining **utility**

We want to release statistical information about a sensitive dataset without comprising the privacy of individual respondents.

Dwork's Solution [ICALP '06]

The output of the mining process should be indistinguishable when run with two databases d_1 and d_2 differing in a single record.

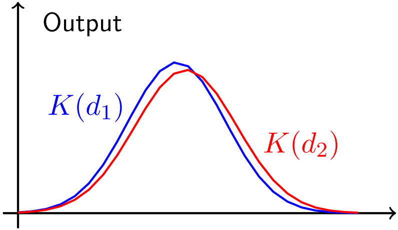

Dwork's Solution [ICALP '06]

The output of the mining process should be indistinguishable when run with two databases d_1 and d_2 differing in a single record.

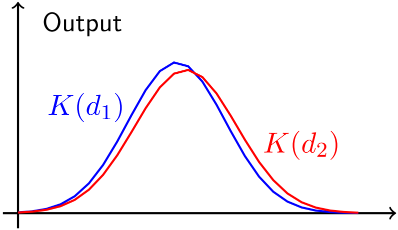

A randomized mechanism K is (ϵ, δ)-*differentially private* iff

$$\forall d_1, d_2 \bullet \Delta(d_1, d_2) \leq 1 \implies \Delta_\alpha (K(d_1), K(d_2)) \leq \delta$$

where $\alpha = \exp(\epsilon)$.

f-divergences - Definition

The f-*divergence* between two distributions μ_1 and μ_2 over a set A is defined as

$$\Delta_f(\mu_1, \mu_2) \triangleq \sum_{a \in A} \mu_2(a) f\left(\frac{\mu_1(a)}{\mu_2(a)}\right)$$

where $f : \mathbb{R}^{\geq 0} \to \mathbb{R}$ is a continuous convex function s.t. $f(1) = 0$.

Some examples

- Statistical distance (Δ_{SD}) $\quad f(t) = \frac{1}{2}|t-1|$
- Kullback-Leibler (Δ_{KL}) $\quad f(t) = t\ln(t)$
- Hellinger distance (Δ_{HD}) $\quad f(t) = \frac{1}{2}(\sqrt{t}-1)^2$

f-divergences - Definition

The f-*divergence* between two distributions μ_1 and μ_2 over a set A is defined as

$$\Delta_f(\mu_1, \mu_2) \triangleq \sum_{a \in A} \mu_2(a) f\left(\frac{\mu_1(a)}{\mu_2(a)}\right)$$

where $f : \mathbb{R}^{\geq 0} \to \mathbb{R}$ is a continuous convex function s.t. $f(1) = 0$.

Some examples

- Statistical distance (Δ_{SD}) $\quad f(t) = \frac{1}{2}|t-1|$
- Kullback-Leibler (Δ_{KL}) $\quad f(t) = t\ln(t)$
- Hellinger distance (Δ_{HD}) $\quad f(t) = \frac{1}{2}(\sqrt{t}-1)^2$
- α-distance (Δ_α) $\quad f(t) = \max\{t-\alpha, 0\}$

Sequential Composition Theorem of DP

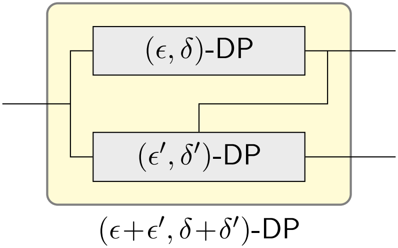

Sequential Composition Theorem of α-distance

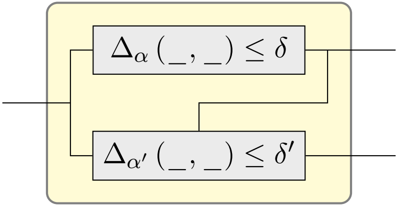

$$\Delta_{\alpha\alpha'}(\_, \_) \leq \delta + \delta'$$

Sequential Composition Theorem of f-divergences

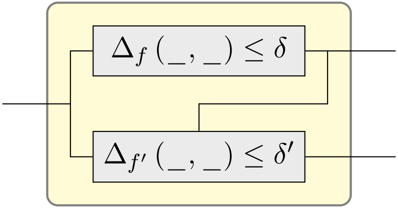

$$\Delta_{f''}\left(\_,\_\right) \leq \delta + \delta'$$

Sequential Composition Theorem of f-divergences

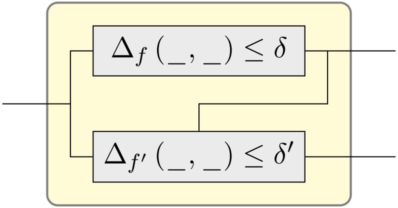

$$\Delta_{f''}(\_, \_) \leq \delta + \delta'$$

We extend the sequential composition theorem of DP by

- Introducing the notion of f-divergence composability.

  (f, f') is f''-*composable*

- Showing that Δ_{SD}, Δ_{KL} and Δ_{HD} are self-composable.

Probabilistic Relational Reasoning for DP [Barthe:2012a].

They propose an approximate relational Hoare logic

$$c_1 \sim_{\alpha,\delta} c_2 : \Psi \Rightarrow \Phi$$

A program c is (ϵ, δ)-DP iff

$$c \sim_{\exp(\epsilon),\delta} c : \Psi \Rightarrow \equiv$$

database adjacency

equality on program states

Relational Hoare Logic for f-divergences

Judgments have the form

$$c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi$$

Such a judgment is *valid* iff for all memories m_1 and m_2

$$m_1 \, \Psi \, m_2 \implies (\llbracket c_1 \rrbracket \, m_1) \, \mathcal{L}_f^{\delta}(\Phi) \, (\llbracket c_2 \rrbracket \, m_2)$$

Relational Hoare Logic for f-divergences

Judgments have the form

$$c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi$$

Such a judgment is *valid* iff for all memories m_1 and m_2

$$m_1 \, \Psi \, m_2 \implies (\llbracket c_1 \rrbracket \, m_1) \, \mathcal{L}_f^\delta(\Phi) \, (\llbracket c_2 \rrbracket \, m_2)$$

Lifting of Φ to a relation over **distributions** on program states

(f, δ)-lifting of Relations

$$\mathcal{L}_f^{\delta}(\cdot) : \mathcal{P}\left(A \times B\right) \to \mathcal{P}\left(\mathcal{D}(A) \times \mathcal{D}(B)\right)$$

- Generalizes previous lifting operator for the exact setting (*ie* $\delta = 0$).
- More or less involved definition for arbitrary relations, but admits simpler characterization for equivalence relations.
- In the case of equality we have

$$\mu_1 \, \mathcal{L}_f^{\delta}(\equiv) \, \mu_2 \iff \Delta_f\left(\mu_1, \mu_2\right) \leq \delta$$

Relational Hoare Logic for f-divergences - Applications

- Bound the f-divergence between programs

$$\Delta_f\left(\llbracket c_1 \rrbracket\, m_1, \llbracket c_2 \rrbracket\, m_2\right) \leq \delta$$

- Relate the probability of individual events

$$\Pr\left[c_2(m_2) : E_2\right]\ f\left(\frac{\Pr\left[c_1(m_1) : E_1\right]}{\Pr\left[c_2(m_2) : E_2\right]}\right) \leq \delta$$

- Model other quantitative notions such as such as continuity or approximate non-interference.

Selected Rules

Weakening

$$\frac{\models c_1 \sim_{f',\delta'} c_2 : \Psi' \Rightarrow \Phi' \qquad \Psi \Rightarrow \Psi' \qquad \Phi' \Rightarrow \Phi \qquad f \leq f' \qquad \delta' \leq \delta}{\models c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi}$$

Sequential composition

$$\frac{(f_1, f_2) \text{ is } f_3\text{-composable} \qquad \models c_1 \sim_{f_1,\delta_1} c_2 : \Psi \Rightarrow \Phi' \qquad \models c_1' \sim_{f_2,\delta_2} c_2' : \Phi' \Rightarrow \Phi}{\models c_1; c_1' \sim_{f_3,\delta_1+\delta_2} c_2; c_2' : \Psi \Rightarrow \Phi}$$

Contributions

- We unveil a connection between differential privacy and f-divergences.
- We generalize the sequential composition theorem of DP to some well-known f-divergences.
- We introduce a program logic for upper-bounding the f-divergences between probabilistic programs.

Thanks for your attention!

References I

Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella-Béguelin.
Probabilistic relational reasoning for differential privacy.
In *39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012*, pages 97–110, New York, 2012. ACM.

John Steinberger.
Improved security bounds for key-alternating ciphers via hellinger distance.
Cryptology ePrint Archive, Report 2012/481, 2012.
`http://eprint.iacr.org/`.

- Improving security bounds for Key-Alternating Cipher via Hellinger Distance [Steinberger:2012].

Improving security bounds for **Key-Alternating Cipher** via Hellinger Distance [Steinberger:2012].

$$E_P(k,\cdot) : \{0{,}1\}^n \to \{0{,}1\}^n$$

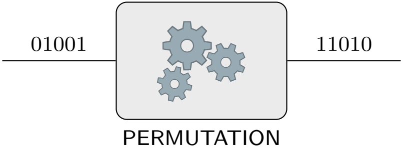

f-divergences in Crypto

Improving **security** bounds for Key-Alternating Cipher via Hellinger Distance [Steinberger:2012].

Hard to distinguish $E_P(k,\cdot)$ *from a true random permutation* Q

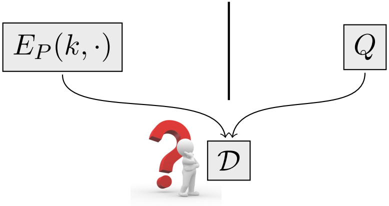

Formally stated as an upper bound of

$$\Delta_{\mathsf{SD}}\left(\mathcal{D}^{E_P(k,\cdot)}, \mathcal{D}^{Q}\right)$$

Improved security guarantees by bounding instead the f-divergence

$$\Delta_{\mathsf{HD}}\left(\mathcal{D}^{E_P(k,\cdot)}, \mathcal{D}^{Q}\right)$$

Key-Alternating Ciphers

$$E_P(k, m) = m'$$

$P=(P_i)_{i=1}^t$ $\quad$ $k=k_0\|\cdots\|k_t$

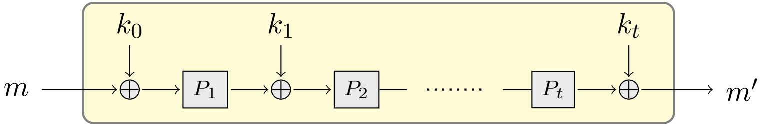

Generalized Data Processing Theorem

For any distribution transformer $h : \mathcal{D}(A) \to \mathcal{D}(B)$

$$\Delta_f \left(h(\mu_1), h(\mu_2)\right) \leq \Delta_f \left(\mu_1, \mu_2\right)$$

Generalized Data Processing Theorem

For any distribution transformer $h : \mathcal{D}(A) \to \mathcal{D}(B)$

$$\Delta_f\left(h(\mu_1), h(\mu_2)\right) \leq \Delta_f\left(\mu_1, \mu_2\right)$$

As a corollary,

$$\Delta_f\left(\llbracket c_1 \rrbracket m_1, \llbracket c_2 \rrbracket m_2\right) \leq \delta \implies \Delta_f\left(\pi_S(\llbracket c_1 \rrbracket m_1), \pi_S(\llbracket c_2 \rrbracket m_2)\right) \leq \delta$$

$$
\begin{array}{lcll}
\mathcal{C} & ::= & \mathsf{skip} & \text{nop} \\
& | & \mathcal{C};\ \mathcal{C} & \text{sequence} \\
& | & \mathcal{V} \leftarrow \mathcal{E} & \text{assignment} \\
& | & \mathcal{V} \xleftarrow{\$} \mathcal{D} & \text{random sampling} \\
& | & \mathsf{if}\ \mathcal{E}\ \mathsf{then}\ \mathcal{C}\ \mathsf{else}\ \mathcal{C} & \text{conditional} \\
& | & \mathsf{while}\ \mathcal{E}\ \mathsf{do}\ \mathcal{C} & \text{while loop} \\
& | & \mathcal{V} \leftarrow \mathcal{P}(\mathcal{E}, \ldots, \mathcal{E}) & \text{procedure call}
\end{array}
$$

$$\frac{\forall m_1, m_2 \bullet\ m_1\ \Psi\ m_2 \implies (m_1\,\{\llbracket e_1 \rrbracket\ m_1/x_1\})\ \Phi\ (m_2\,\{\llbracket e_2 \rrbracket\ m_2/x_2\})}{\vdash x_1 \leftarrow e_1 \sim_{f,0} x_2 \leftarrow e_2 : \Psi \Rightarrow \Phi}\,[\mathsf{assn}]$$

$$\frac{\forall m_1, m_2 \bullet\ m_1\ \Psi\ m_2 \implies \Delta_f\,(\llbracket \mu_1 \rrbracket\ m_1, \llbracket \mu_2 \rrbracket\ m_2) \le \delta}{\vdash x_1 \xleftarrow{\$} \mu_1 \sim_{f,\delta} x_2 \xleftarrow{\$} \mu_2 : \Psi \Rightarrow x_1\langle 1\rangle = x_2\langle 2\rangle}\,[\mathsf{rand}]$$

$$\frac{\begin{array}{c}\Psi \implies b\langle 1\rangle \equiv b'\langle 2\rangle \\ \vdash c_1 \sim_{f,\delta} c_1' : \Psi \wedge b\langle 1\rangle \Rightarrow \Phi \qquad \vdash c_2 \sim_{f,\delta} c_2' : \Psi \wedge \neg b\langle 1\rangle \Rightarrow \Phi\end{array}}{\vdash \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \sim_{f,\delta} \mathsf{if}\ b'\ \mathsf{then}\ c_1'\ \mathsf{else}\ c_2' : \Psi \Rightarrow \Phi}\,[\mathsf{cond}]$$

$$\frac{\begin{array}{c}(f_1, \ldots, f_n)\ \text{composable and monotonic} \\ \Theta \triangleq b\langle 1\rangle \equiv b'\langle 2\rangle \qquad \Psi \wedge e\langle 1\rangle \le 0 \implies \neg b\langle 1\rangle \\ \vdash c \sim_{f_1,\delta} c' : \Psi \wedge b\langle 1\rangle \wedge b'\langle 2\rangle \wedge e\langle 1\rangle = k \Rightarrow \Psi \wedge \Theta \wedge e\langle 1\rangle < k\end{array}}{\vdash \mathsf{while}\ b\ \mathsf{do}\ c \sim_{f_n,n\delta} \mathsf{while}\ b'\ \mathsf{do}\ c' : \Psi \wedge \Theta \wedge e\langle 1\rangle \le n \Rightarrow \Psi \wedge \neg b\langle 1\rangle \wedge \neg b'\langle 2\rangle}\,[\mathsf{while}]$$

$$\frac{}{\vdash \mathsf{skip} \sim_{f,0} \mathsf{skip} : \Psi \Rightarrow \Psi}\,[\mathsf{skip}] \qquad \frac{\begin{array}{c}(f_1, f_2)\ \text{is}\ f_3\text{-composable} \\ \vdash c_1 \sim_{f_1,\delta_1} c_2 : \Psi \Rightarrow \Phi' \quad \vdash c_1' \sim_{f_2,\delta_2} c_2' : \Phi' \Rightarrow \Phi\end{array}}{\vdash c_1; c_1' \sim_{f_3,\delta_1+\delta_2} c_2; c_2' : \Psi \Rightarrow \Phi}\,[\mathsf{seq}]$$

$$\frac{\begin{array}{c}\vdash c_1 \sim_{f,\delta} c_2 : \Psi \wedge \Theta \Rightarrow \Phi \\ \vdash c_1 \sim_{f,\delta} c_2 : \Psi \wedge \neg\Theta \Rightarrow \Phi\end{array}}{\vdash c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi}\,[\mathsf{case}] \qquad \frac{\begin{array}{c}\vdash c_1 \sim_{f',\delta'} c_2 : \Psi' \Rightarrow \Phi' \\ \Psi \Rightarrow \Psi' \quad \Phi' \Rightarrow \Phi \quad f \le f' \quad \delta' \le \delta\end{array}}{\vdash c_1 \sim_{f,\delta} c_2 : \Psi \Rightarrow \Phi}\,[\mathsf{weak}]$$

(f, δ)-lifting of Relations

$$\mathcal{L}_f^{\delta}(\cdot) : \mathcal{P}\left(A \times B\right) \rightarrow \mathcal{P}\left(\mathcal{D}(A) \times \mathcal{D}(B)\right)$$

$$\mu_1 \, \mathcal{L}_f^{\delta}(R) \, \mu_2 \triangleq \exists \mu_L, \mu_R \bullet \begin{cases} \mathsf{supp}\left(\mu_L\right) \subseteq R \; \wedge \; \mathsf{supp}\left(\mu_R\right) \subseteq R \\ \pi_1(\mu_L) = \mu_1 \; \wedge \; \pi_2(\mu_R) = \mu_2 \\ \Delta_f\left(\mu_L, \mu_R\right) \leq \delta \end{cases}$$

The α-*distance* $\Delta_\alpha(\mu_1, \mu_2)$ between distributions μ_1 and μ_2 is defined as

$$\Delta_\alpha(\mu_1, \mu_2) \triangleq \max_S \Pr[\mu_1 \in S] - \alpha \Pr[\mu_2 \in S]$$