Reasoning about Divergences via Span-liftings

(Extended Abstract)

Tetsuya Sato
Department of Computer Science and Engineerings
University at Buffalo, SUNY
tetsuyas@buffalo.edu

1 Introduction

Differential privacy (DP) [3] is a definition of data privacy that guarantee strong privacy against database attacks using background-knowledge. This privacy has attracted the attention of several academic and people in the industry. Differential privacy restricts the range of *privacy loss* random variables for any two "adjacent" datasets differing in at most one data record. Rényi differential privacy (RDP) [8] and zero-concentrated differential privacy (zCDP) [2] are relaxed notion of differential privacy constraining *the moment* of the privacy loss random variable. These relaxations can be good definitions of data privacy of machine learning mechanisms such as privacy-preserving mechanisms for Bayesian inference [4].

This work is motivated to verify such all privacy-preserving mechanisms. Since distribution of datasets may be continuous, we also want to verify continuous probabilistic programs.

We give semantic models for reasoning about RDP and zCDP by extending fpRHL [1] to suppport more *general statistical divergences* of *both discrete and continuous* distributions. Furthermore, our extended semantic models can be used not only for reasoning about differential privacy but also developing more general approximate logical relations reasoning about the probabilistic behavior of continuous probabilistic programs.

2 Problems

To extend the semantic model of fpRHL, we face the following two technical difficulties: first, we need a framework that supports more general divergences than f-divergences, although there is a framework for reasoning about f-divergences [1]. RDP and zCDP can be defined by the α-Rényi divergence [9], which is the logarithm of a f-divergence. Strictly, it is not formulated as a f-divergences. In particular, when we characterize zCDP by statistical divergences, it can return negative values.

Second, we also aim to give a semantic model for continuous programming languages. In the previous work [10], we have a "witness-free lifting" for approximate DP supporting the continuous case, but the approach in the previous work [10] does not work well for RDP and zCDP. Our previous approach in [10] is based on a method to give categorical monad lifting, named codensity lifting [5]. Roughly speaking, codensity lifting is defined as a *large intersection* indexed by all relation-preserving maps from a given relation to fixed relations. Fortunately, in the case of approximate differential privacy, we simplify such large intersections [10]. However, we could not simplify such large intersections for other statistical divergences.

Summarizing the above, we have the following two technical difficulties on semantics framework:

1. We need semantics models which support more general statistical divergences beyond f-divergences.
2. We need it to support continuous distributions, but our previous approach of witness-free lifting in [10] does not work well.

3 Solutions

To solve the first technical difficulty, we first relax the notion of divergences to sub-probability distributions (subdistributions). Actually, we begin with just functions of the form

$$\Delta_X \colon \mathrm{Dist}(X) \times \mathrm{Dist}(X) \to \mathbb{R} \cup \{-\infty, +\infty\}$$

where $\mathrm{Dist}(X)$ is the set of subdistributions on X. Then, we *axiomatize* some *basic properties* of divergences inspired from the composability, additivity, and continuity of f-divergences discussed in [1, 7].

To solve the second technical difficulty, we extend the notion of "2-witness lifting" introduced in [1] to a novel notion of *span-lifting*. It is difficult to extend 2-witness lifting to the continuous case directly (in the previous work [10], the author took a different way).

Technically, 2-witness lifting extends a binary relation $R \subseteq X \times Y$ to a binary relation $R^{\sharp\mathrm{DP}(\varepsilon,\delta)} \subseteq \mathrm{Dist}(X) \times \mathrm{Dist}(Y)$ of subdistributions:

$$\begin{aligned} \mu_1 R^{\sharp\mathrm{DP}(\varepsilon,\delta)} \mu_2 \iff \exists \mu_L, \mu_R \in \mathrm{Dist}(R).&\\ \mu_1 = \pi_1(\mu_L), \pi_2(\mu_R) = \mu_2&\\ \Delta^{\mathrm{DP}(\varepsilon)}(\mu_L, \mu_R) \leq \delta& \end{aligned}$$

PPS'18, January 9, 2018, Los Angeles, CA, USA
2018.

where $\Delta^{\mathrm{DP}(\varepsilon)}$ is an f-divergence describing approximate DP [1] (we can replace it to other divergences), and $\pi_i(\mu)$ is the i-th marginal of μ.

In the relation $R^{\sharp\mathrm{DP}(\varepsilon,\delta)}$, two subdistributions μ_1, μ_2 are related by the existence of witness μ_L, μ_R. It is a problematic to extend to the continuous case because the recovering μ_L, μ_R from the membership $(\mu_1, \mu_2) \in R^{\sharp\mathrm{DP}(\varepsilon,\delta)}$ is *restricted* in the continuous case.

For example, we consider a relation-preserving map $(f, g) \colon S \to R^{\sharp\mathrm{DP}(\varepsilon,\delta)}$, that is, two functions f, g such that $(f(x), g(y)) \in R^{\sharp\mathrm{DP}(\varepsilon,\delta)}$ whenever $(x, y) \in S$. In the discrete case, it is no problem to take a mapping $(x, y) \mapsto (\mu_L, \mu_R)$ by the axiom of choice. However, in the continuous case, it is problematic that such mapping needs to be a *measurable* function while the axiom of choice does not guarantee measurability.

This problem is hard to solve, but easy to avoid. It suffices to enrich the structure of 2-witness liftings to make precise witness distributions. In short, we consider the following 4-ary relations instead of binary relations:

$$
\begin{aligned}
&(\mu_1, \mu_2, \mu_L, \mu_R) \in R^{\sharp\mathrm{DP}(\varepsilon,\delta)} \\
&\iff \mu_1 = \pi_1(\mu_L), \pi_2(\mu_R) = \mu_2, \Delta^{\mathrm{DP}(\varepsilon)}(\mu_L, \mu_R) \le \delta.
\end{aligned}
$$

This modification is not problematic to give a semantic model of probabilistic language. First, in many practical cases, measurable functions $(x, y) \mapsto (\mu_L, \mu_R)$ are almost obviously given. Second, thanks to the axiom of choice, this modification covers all the discrete case discussed in [1].

3.1 construction

Based on the above ideas, we extend the notion of divergences, and introduce a novel notion of *span-liftings for general divergences*. Instead of 4-ary relations we use spans $X \xleftarrow{h} \Phi \xrightarrow{k} Y$ in the category **Meas** of measurable spaces and measurable functions. We then relate basic properties of divergences to semantical properties of span-liftings for the divergences. The span-liftings for divergences form a *graded monad* [6], which gives main structures for formal verifications when the divergences satisfy some basic properties. Finally, we check basic properties of divergences for RDP, zCDP, and approximate DP, and apply them to our framework.

To sum up, in this study, we extend the semantic model of fpRHL to support general divergences in both the discrete and continuous case, and instantiate it for RDP, zCDP and approximate DP as follows:

1. We introduce general notions of divergences, and axiomatize the basic properties as in [1, 7].
2. We introduce a novel notion of span-lifting for general divergences to support both discrete and continuous case. Then, we relate basic properties of divergences and span-liftings for the divergences.
3. We instantiate this framework for RDP, zCDP, and approximate DP by checking basic properties of divergences.

Acknowledgment

This extended abstract is based on a joint work with Gilles Barthe, Marco Gaboardi, Justin Hsu, and Shin-ya Katsumata. The author thanks the collaborators for fruitful discussions and stimulating suggestions. The author also thanks the anonymous reviewers for advices improving this extended abstract.

References

[1] Gilles Barthe and Federico Olmedo. 2013. Beyond Differential Privacy: Composition Theorems and Relational Logic for f-divergences between Probabilistic Programs. In *Automata, Languages, and Programming: 40th International Colloquium, ICALP 2013, Proceedings, Part II*. Springer Berlin Heidelberg, Berlin, Heidelberg, 49–60. https://doi.org/10.1007/978-3-642-39212-2_8

[2] Mark Bun and Thomas Steinke. 2016. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B*. 635–658. https://doi.org/10.1007/978-3-662-53641-4_24

[3] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). LNCS, Vol. 3876. Springer Berlin Heidelberg, 265–284. https://doi.org/10.1007/11681878_14

[4] Joseph Geumlek, Shuang Song, and Kamalika Chaudhuri. 2017. Rényi Differential Privacy Mechanisms for Posterior Sampling. *CoRR* abs/1710.00892 (2017).

[5] Shin-ya Katsumata and Tetsuya Sato. 2015. Codensity Liftings of Monads. In *Conference on Algebra and Coalgebra in Computer Science (CALCO 2015) (Leibniz Intern. Proc. in Informatics (LIPIcs))*, Vol. 35. Schloss Dagstuhl, 156–170. https://doi.org/10.4230/LIPIcs.CALCO.2015.156

[6] Shin-ya Katsumata. 2014. Parametric Effect Monads and Semantics of Effect Systems. In *ACM Symposium on Principles of Programming Languages (POPL '14)*. ACM, New York, NY, USA, 633–645. https://doi.org/10.1145/2535838.2535846

[7] Friedrich Liese and Igor Vajda. 2006. On Divergences and Informations in Statistics and Information Theory. *IEEE Transactions on Information Theory* 52, 10 (Oct 2006), 4394–4412. https://doi.org/10.1109/TIT.2006.881731

[8] Ilya Mironov. 2017. Rényi Differential Privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. 263–275. https://doi.org/10.1109/CSF.2017.11

[9] Alfred Renyi. 1961. On Measures of Entropy and Information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. University of California Press, Berkeley, Calif., 547–561. http://projecteuclid.org:443/euclid.bsmsp/1200512181

[10] Tetsuya Sato. 2016. Approximate Relational Hoare Logic for Continuous Random Samplings. *ENTCS* 325 (2016), 277–298. https://doi.org/10.1016/j.entcs.2016.09.043 Mathematical Foundations of Programming Semantics (MFPS XXXII).