

A randomized computation M satisfies ϵ -differential privacy if for any two possible input data sets A and B , and any subset of possible outputs S ,

$$P(M(A) \in S) \leq P(M(B) \in S) \times \exp(\epsilon \times |A \ominus B|)$$


where $A \ominus B$ is the set of records in A or B , but not both.

1. No computational / informational assumptions
2. Agnostic to data type
3. Formal
- ...
- n. Exposes conditional probabilities:


$$P(\text{outcome} = z \mid \text{data} = X).$$

Noisy observations z of hidden data X induces marginal posterior over model parameters θ .

$$p(z|\theta) = \int dX p(z|X) p(X|\theta)$$




Directly observable data




Our setting

1. Integrate multiple observations;
2. Express/use confidence in answers;
3. Use prior knowledge about data;
4. Posteriors over unmasked questions.


Example: PCA



$\epsilon = 0.003$



$\epsilon = 0.01$



$\epsilon = 0.1$

$$p(x|\theta) = \mathcal{N}(0, \theta\theta^\top + \sigma^2 I).$$

Example: Logistic Regression

	SYNTH	CM2	ADULT
<i>Heuristic</i>	37.40 \pm 15.75	9.32 \pm 1.18	43.15 \pm 7.85
<i>Inference</i>	29.14 \pm 5.54	8.84 \pm 0.79	36.07 \pm 6.32
<i>Benchmark</i>	16.40	5.40	26.09
<i>NIPS 08</i>		19.03 \pm 11.05	

Cumulative density functions for classification error rates:

