A Gentle Introduction to Differential Privacy with Use Cases

José Cabrero-Holgueras – jose.cabrero.holgueras@cern.ch

DOI: 10.5281/zenodo.5095316

Differential Privacy

- Differential Privacy (DP) is a property that gives privacy guarantees to user of a population. It ensures that belonging to a dataset will not release private information of himself, but only from the whole population.
- The formal definition of DP grants that with a high probability computing a statistic ($\mathcal{M}$) on a dataset (x) and the same dataset but without an individual (y) would yield the same result, with a really small difference parametrized by ε and δ.
- The formal definition is: $Pr[\mathcal{M}(x) \in S] \leq e^{\varepsilon} * Pr[\mathcal{M}(y) \in S] + \delta$.
- It also means that any individual can be discarded from the dataset without affecting the result of the algorithm.
  - The individual helps building the population, but there is no need of him to create the dataset.
- These techniques can act before (Local DP) or after (Global DP) the creation of the dataset.

DOI: 10.5281/zenodo.5095316

[1] C. Dwork and A. Roth, The Algorithmic Foundations of Differential Privacy, 2014.

Local Differential Privacy: Randomized Response

- We want to understand the percentage of population affected by some sensitive disease (COVID19, AIDS,…).
- Participants may not want to answer honestly because of the implications it may have.
- This randomized response algorithm adds plausible deniability. That is, the participant cannot be hold accountable for his answer, because it is influenced by the randomness of the coins.
- The algorithm is the is run by tossing two coins. Participants answer based on their affection and the results of the coins:
  - If the first coin lands on heads, the person will tell the truth (whether he is affected or not).
  - If the first coin lands on tails, the answer will be determined by the second coin and not by the user reality (the result is defined by the coin, not the participant).
- With 75% probability we obtain correct answer AND with sufficiently large populations, we obtain accurate statistics, while the privacy of participants is never individually infringed.

DOI: 10.5281/zenodo.5095316

Global Differential Privacy: Towards Privacy-Preserving Dataset Release

- A hospital wants to release a dataset of faces:
  - We want to understand statistics about faces.
  - We cannot allow re-identification of a person.
- We can introduce modifications to specific sections of the picture with high sensitivity.
- Sensitivity is the amount of information carried out by a specific piece of data. Sections with more sensitivity are obfuscated harder to avoid reidentification.
- The figure shows the sensitivity of a dataset of human faces.
  - Yellow areas (eyes, hair, earings/necklaces) are more identificative (i.e., carry more information).
  - Blue areas (corners or neck) do not deliver a lot of information (i.e., are less predominant for reidentification).

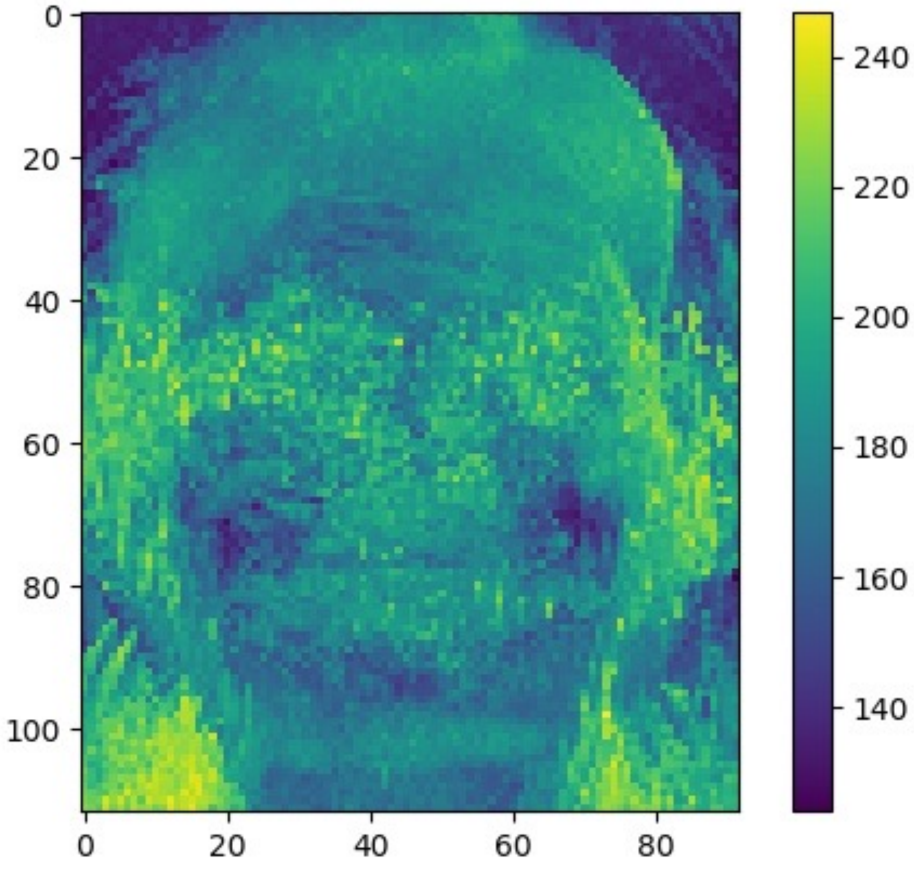

DOI: 10.5281/zenodo.5095316

Global Differential Privacy: The Gaussian Mechanism

- The Gaussian Mechanism is a differentially private algorithm, that determines the amount of noise to be added to guarantee DP based on the sensitivity.

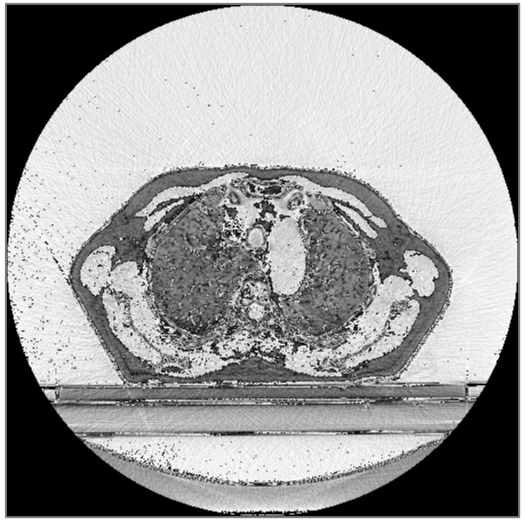

Original Image

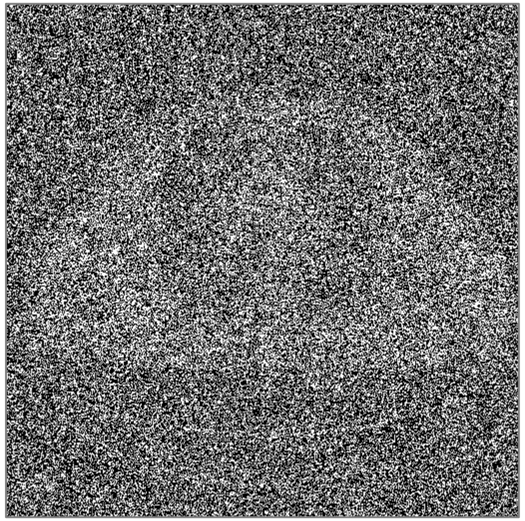

Differentially Private Image

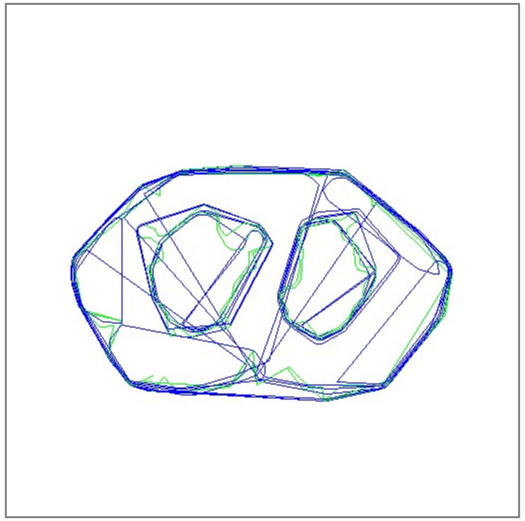

Noise resilient segmentation of DP Image.

DOI: 10.5281/zenodo.5095316

[2] NSCLC-Radiomics – The Cancer Imaging Archive (TCIA) Public Access

Other interesting sources:

- I want to know more about DP without getting in detail:
  - Why differential privacy is awesome - Damien Desfontaines
  - Understanding differential privacy and why it matters for digital rights

- Why do we need perfect randomness for Differential Privacy?
  - Dodis, Y., López-Alt, A., Mironov, I., & Vadhan, S. (2012, August). Differential privacy with imperfect randomness. In *Annual Cryptology Conference* (pp. 497-516). Springer, Berlin, Heidelberg

- The original work on Differential Privacy:
  - Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, *9*(3-4), 211-407.

DOI: 10.5281/zenodo.5095316