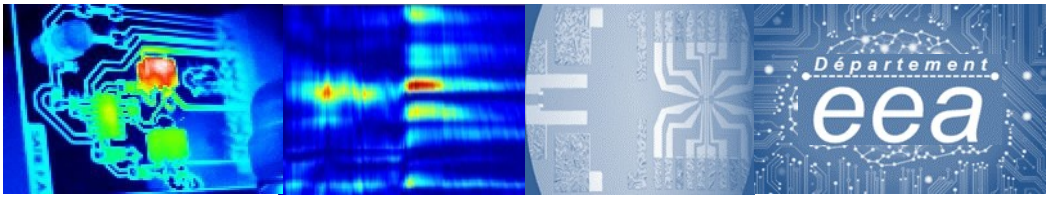

école normale supérieure paris—saclay

Differential Privacy applied to Crowd Density Estimation

Solal Nathan,
SATIE, ENS Paris-Saclay

Abstract: Taking into account the importance of crowd density estimation in multiple areas of research and society, I tried to adress the raising concern of privacy in this domain of computer vision.

SATIE

Work:

Goal: Be able to release public information about density estimation coming from video surveillance footage to ensure better **safety** and more **openly accesible data** for researchers without the fear of compromising user personal information. Getting the **density map** sufficently "blurry" to fool the adversarial algorithm, but retaining enough clarity for safety decision making inference and research.

Methodology: Work with toy models (1D or 2D case) to develop post-processing and **noises** that can be applied to the density map to fool the **adversarial algorithm**.

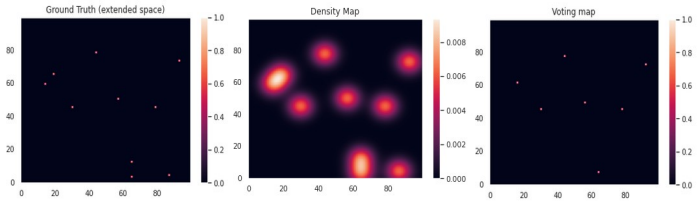

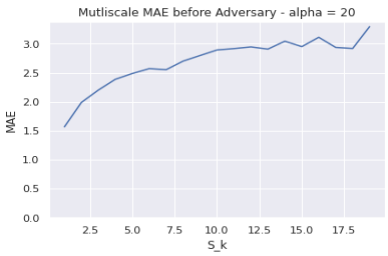

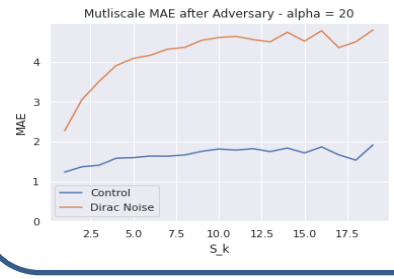

We can observe the **Multiscale MAE** (Mean Average Error) before and after the adversarial algorithm processes the data. The goal is to have a small MAE for a large scale (low S_k) and a large MAE after a certain scale S_k. The idea is to retain density information on a large scale but withdraw any possible information that could target individuals on a small scale.

A simple **Dirac noise** added to the original data already has good results. However other means of post-processing like the **density box method** (giving only the density for a given square of size S_k) or the **random swaps method** (randomly swapping bits of information within a given S_k to maintain the local density) could still be tested and obtain better results.

Conclusion: The results are promising on simple toy models. Training a real crowd density estimation model (tiny heads/LFE) helped on the generation of data but was not used as input. More post-processings should be tested, using real input data and adding the temporal dimension. Differential Privacy is very formal and hard to apply in computer vision and a relaxed version has to be used.

Main references:

- Dwork, et al. (2006). Calibrating noise to sensitivity in private data analysis.
- Zhan, B., et al. (2008). Crowd analysis: A survey.
- Lecuyer, M., et al. (2019). Certified robustness to adversarial examples with differential privacy.