INTERNATIONAL INSTITUTE OF INFORMATION TECHNOLOGY
HYDERABAD

R&D SHOWCASE 2021

Differential Privacy in Artificial Intelligence

Privacy Preserving Artificial Intelligence

Is it possible to answer questions using data we cannot see?

Privacy Preserving AI Techniques

- Federated Learning
- Differential Privacy
- Remote Execution
- Encryption
- Multi-Party Computation

These techniques enable us to build AI solutions for sensitive problems like cancer, dementia, depression, covid etc. Applying these techniques to various algorithms in AI and ML has become an active area of research. They also help provide optimizations in other aspects of AI problems. A need for new techniques such as DP arises due to infeasibility of older encryption like methods.

Differentially Private Constraint Optimization

Goal

Designing a scalable DCOP algorithm that preserves constraint privacy from related participants through differential privacy techniques

Challenges

- Protecting information leak during information exchange
- From the final assignment
- Ensuring Scalability

Constraint Privacy loss due to,
1. Sampling
2. Difference in utility

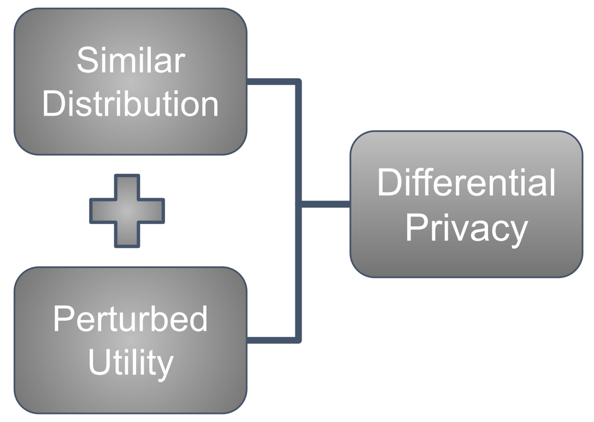

Sample:
1. Gibbs Distribution
2. Agent maximizes expected locall utility

Root:
1. Track's solutions
2. Broadcasts the best solution, so far

Value:
1. Difference in utility
2. Sum is global difference

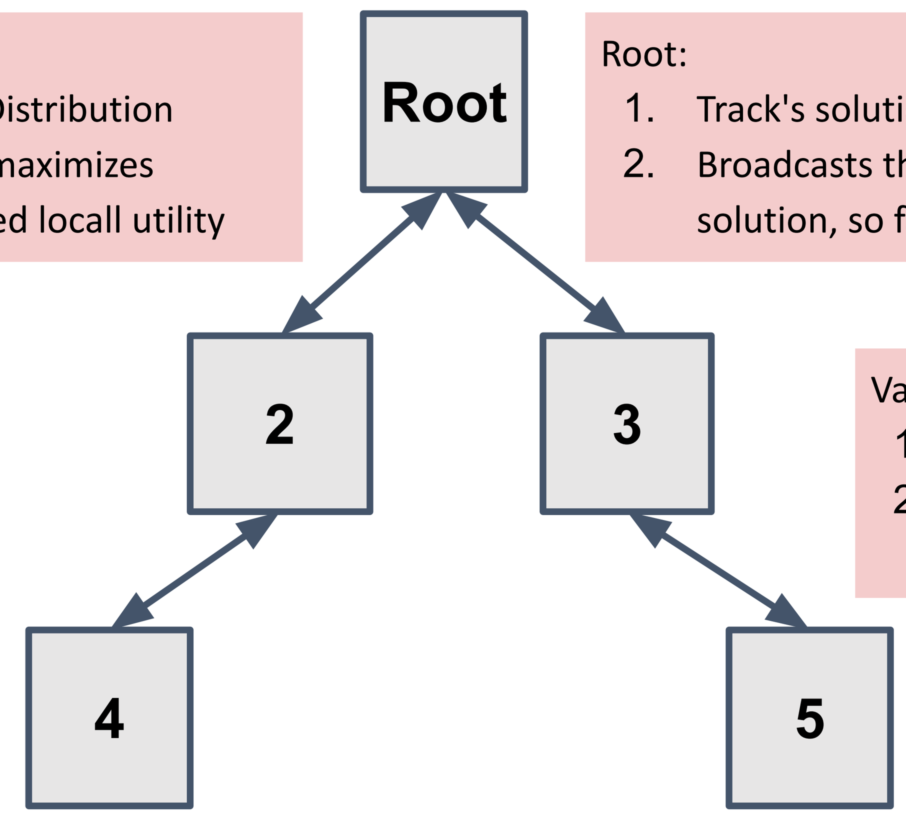

SD-Gibbs Algorithm Framework

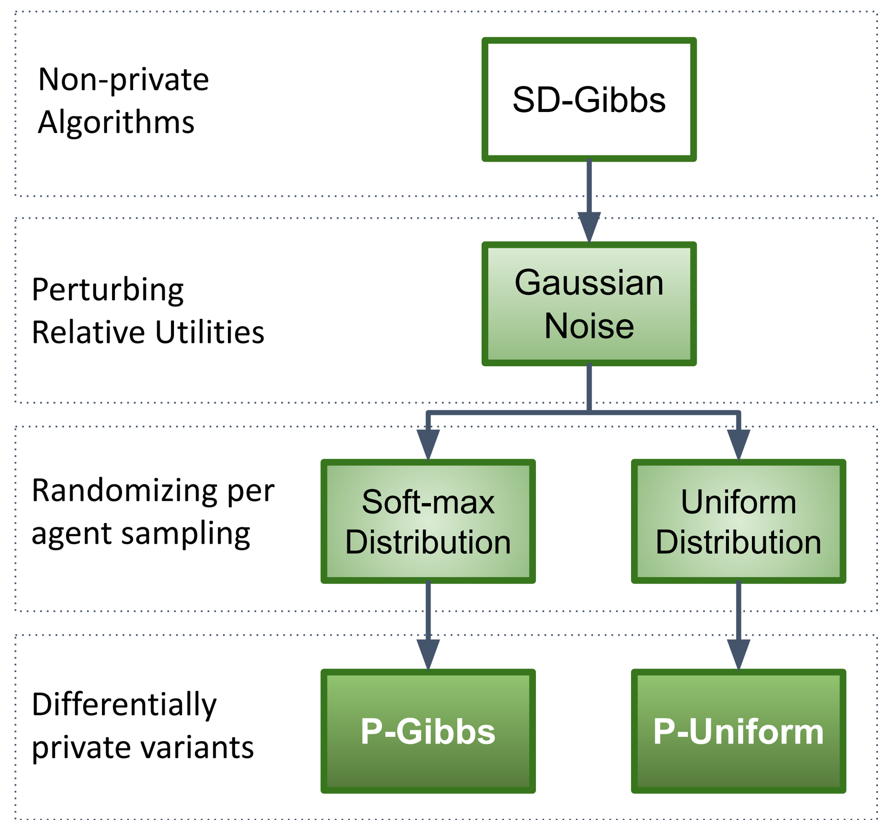

Introduction to Differential Privacy

Differential Privacy (DP) is a system for publicly sharing information about a dataset that masks individual contributions while retaining the big picture, via data randomization. It allows us to quantify the privacy loss.

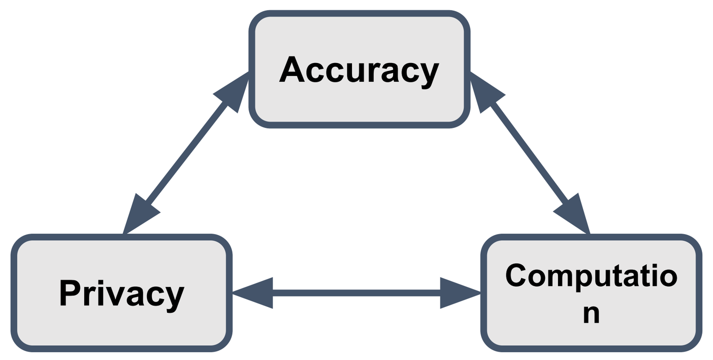

(ε,δ)-DP. A randomized algorithm M gives a privacy guarantee of (ε,δ)-DP if for all pairs of adjacent datasets d, d', and all outputs S with $\varepsilon,\delta > 0$, we have

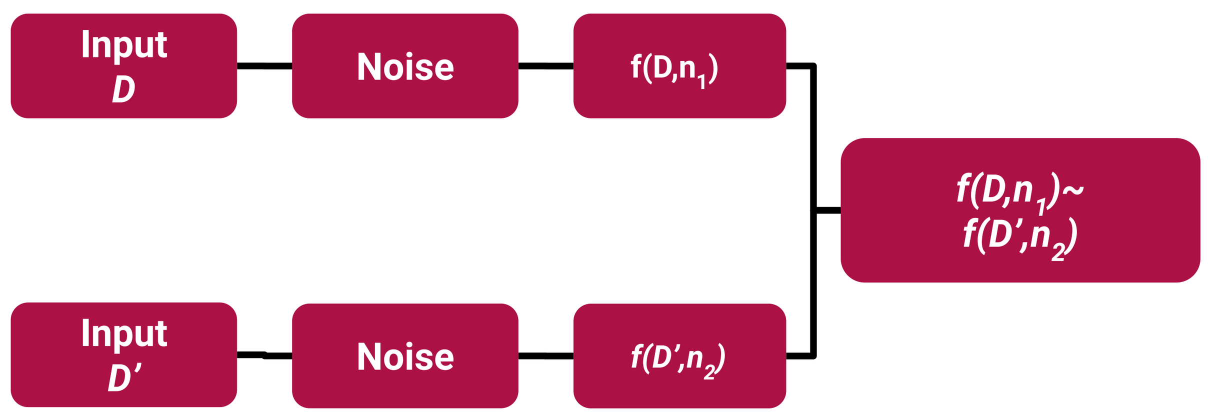

$$\Pr[M(d) = S] \le \exp(\varepsilon) \cdot \Pr[M(d') = S] + \delta$$

Smaller εs, δs => **better** privacy

Differentially Private Deep Reinforcement Learning

RLis a sub-field of ML where we train an agent to learn a policy to perform a task in an environment by offering a reward to it for every action it takes. Deep RL uses Deep NNs for training with policy gradients.

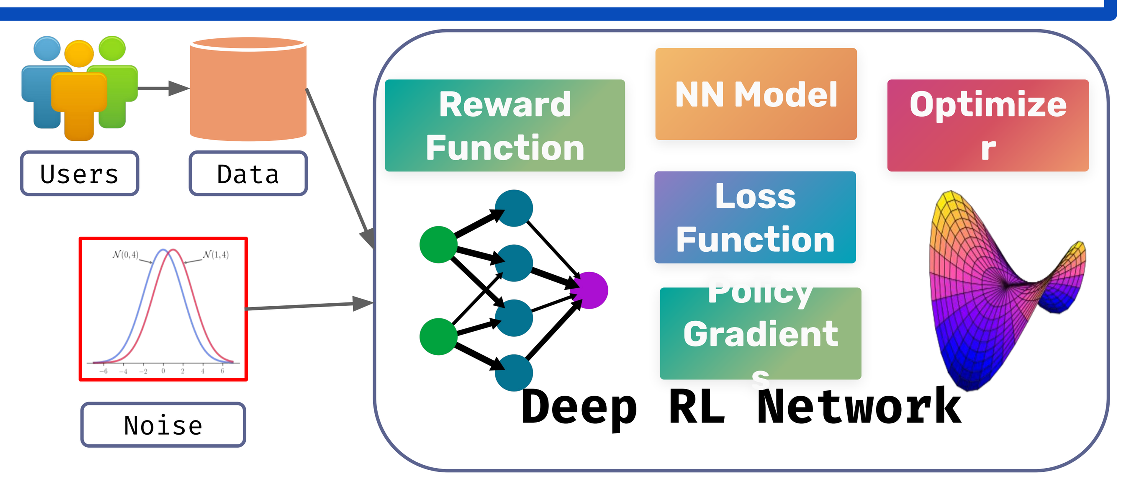

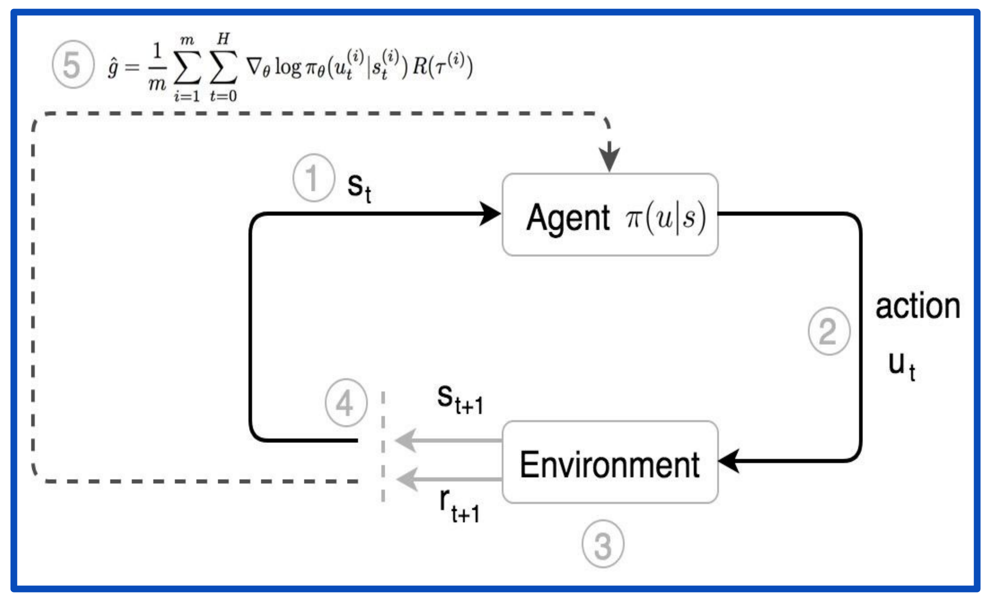

We need to protect the private reward function being used to train the agent by introducing DP to Deep RL by smartly adding noise to the learning process. We also investigate the relationship between privacy budget and generalization ability of the learned policy.

Publications: Sankarshan Damle, Aleksei Triastcyn, Boi Faltings, Sujit Gujar. "Differentially Private Multi-Agent Constraint Optimization". In Second AAAI Workshop on Privacy-Preserving Artificial Intelligence (PPAI-21)

Authors: Dr. Praveen Paruchuri, Dr. Sujit Gujar, Kritika Prakash, Sankarshan Damle