Differentially Private Distributed Convex Optimization via Functional Perturbation

Erfan Nozari

Department of Mechanical and Aerospace Engineering
University of California, San Diego
http://carmenere.ucsd.edu/erfan

July 6, 2016

Joint work with **Pavankumar Tallapragada** and **Jorge Cortés**

UC San Diego
Jacobs School of Engineering

Distributed Coordination

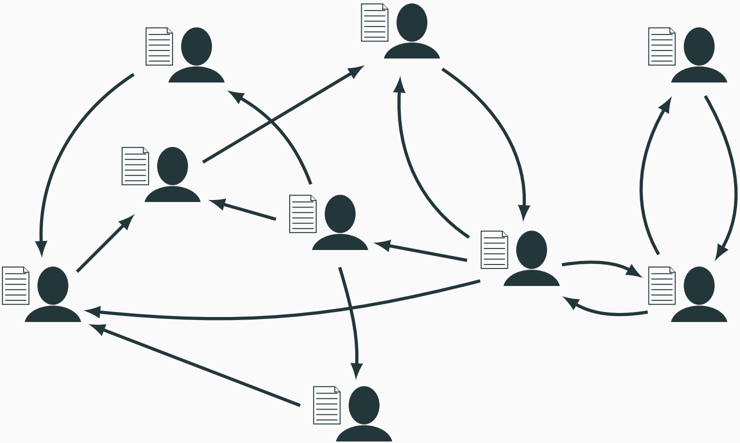

Distributed Coordination

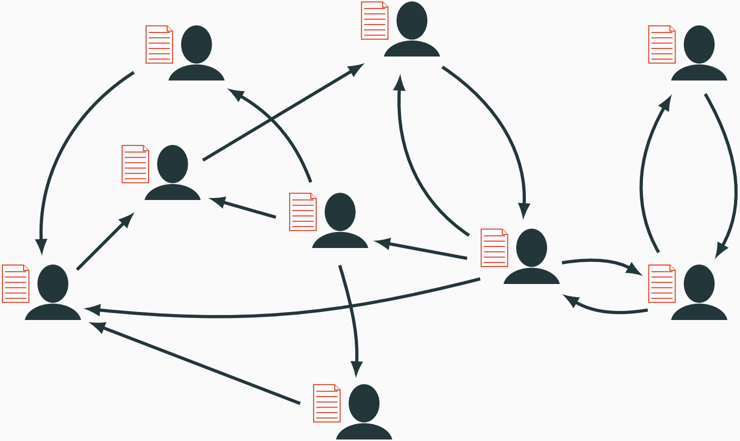

What if local information is sensitive?

Motivating Scenario: Optimal EV Charging
[Han *et. al.*, 2014]

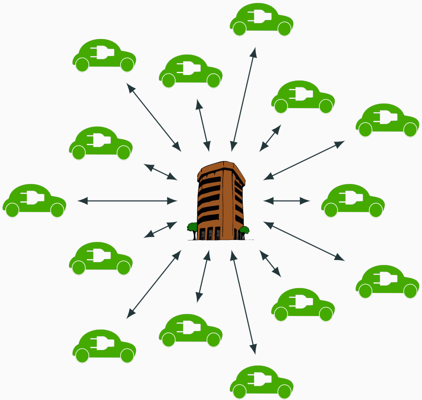

Motivating Scenario: Optimal EV Charging
[Han *et. al.*, 2014]

Central aggregator solves:

$$\underset{r_1,\ldots,r_n}{\text{minimize}} \quad U\big(\textstyle\sum_{i=1}^{n} r_i\big)$$
$$\text{subject to} \quad r_i \in \mathcal{C}_i \quad i \in \{1,\ldots,n\}$$

- U = energy cost function
- $r_i = r_i(t)$ = charging rate
- $\mathcal{C}_i$ = local constraints

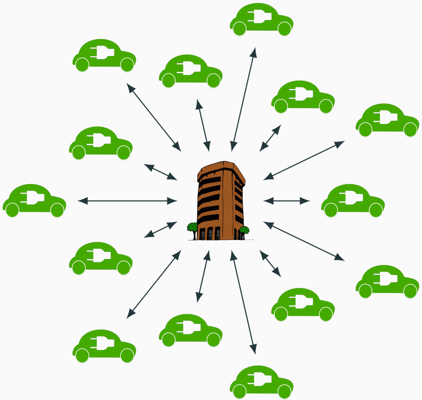

Motivating Scenario: Optimal EV Charging
[Han *et. al.*, 2014]

Central aggregator solves:

$$\begin{aligned} &\underset{r_1,\dots,r_n}{\text{minimize}} \quad U\Big(\sum_{i=1}^{n} r_i\Big) \\ &\text{subject to} \quad r_i \in \mathcal{C}_i \quad i \in \{1,\dots,n\} \end{aligned}$$

- U = energy cost function
- $r_i = r_i(t)$ = charging rate
- $\mathcal{C}_i$ = local constraints

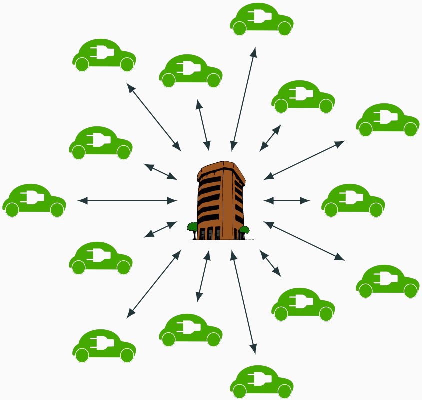

Myth: Aggregation Preserves Privacy

Myth: Aggregation Preserves Privacy

- Fact: NOT in the presence of **side-information**

Myth: Aggregation Preserves Privacy

- Fact: NOT in the presence of **side-information**
- Toy example:

Database

1	100
2	120
⋮	
n	90

→ Average = 110

Myth: Aggregation Preserves Privacy

- Fact: NOT in the presence of **side-information**
- Toy example:

Database

1	100
2	120
⋮	
n	90

→ Average = 110

Side Information

2	120
⋮	
n	90

$\Rightarrow d_1 = 100$

Myth: Aggregation Preserves Privacy

- Fact: NOT in the presence of **side-information**

- Toy example:

Database

1	100
2	120
⋮	
n	90

→ Average = 110

Side Information

2	120
⋮	
n	90

$\Rightarrow d_1 = 100$

- Real example: A. Narayanan and V. Shmatikov successfully **de-anonymized** Netflix Prize dataset (2007)
  Side information: IMDB databases!

Outline

1. DP Distributed Optimization
   - Problem Formulation
   - Impossibility Result
2. Functional Perturbation
   - Perturbation Design
3. DP Distributed Optimization via Functional Perturbation
   - Regularization
   - Algorithm Design and Analysis

Outline

1. DP Distributed Optimization
   - Problem Formulation
   - Impossibility Result
2. Functional Perturbation
   - Perturbation Design
3. DP Distributed Optimization via Functional Perturbation
   - Regularization
   - Algorithm Design and Analysis

Standard additive convex optimization problem:

$$\begin{aligned} \underset{x \in D}{\text{minimize}} \quad & f(x) \triangleq \sum_{i=1}^{n} f_i(x) \\ \text{subject to} \quad & G(x) \leq 0 \\ & Ax = b \end{aligned}$$

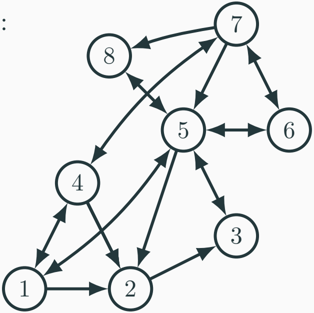

Assumption:

- D is compact
- f_i's are strongly convex and C^2

Problem Formulation

Optimization

Standard additive convex optimization problem:

$$\begin{aligned} \underset{x \in D}{\text{minimize}} \quad & f(x) \triangleq \sum_{i=1}^{n} f_i(x) \\ \text{subject to} \quad & G(x) \leq 0 \\ & Ax = b \end{aligned}$$

$$\underset{x \in X}{\text{minimize}} \quad f(x) \triangleq \sum_{i=1}^{n} f_i(x)$$

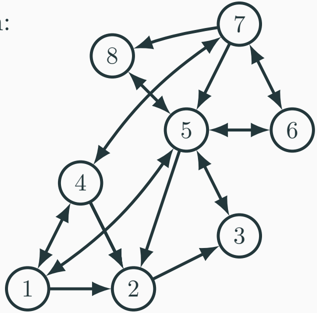

Assumption:

- D is compact
- f_i's are strongly convex and C^2

Standard additive convex optimization problem:

$$\underset{x \in X}{\text{minimize}} \quad f(x) \triangleq \sum_{i=1}^{n} f_i(x)$$

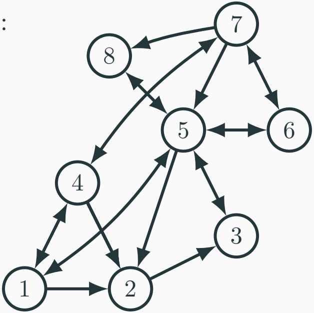

Assumption:

- D is compact
- f_i's are strongly convex and C^2

Standard additive convex optimization problem:

$$\underset{x \in X}{\text{minimize}} \quad f(x) \triangleq \sum_{i=1}^{n} f_i(x)$$

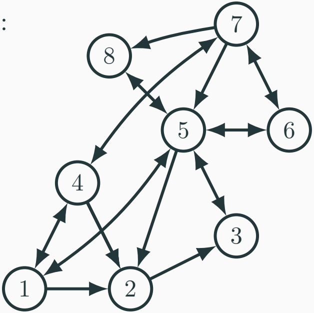

- A **non-private** solution [Nedic *et. al.*, 2010]:

$$x_i(k+1) = \text{proj}_X(z_i(k) - \alpha_k \nabla f_i(z_i(k)))$$

$$z_i(k) = \sum_{j=1}^{n} w_{ij} x_j(k)$$

Assumption:

- D is compact
- f_i's are strongly convex and C^2

Standard additive convex optimization problem:

$$\underset{x \in X}{\text{minimize}} \quad f(x) \triangleq \sum_{i=1}^{n} f_i(x)$$

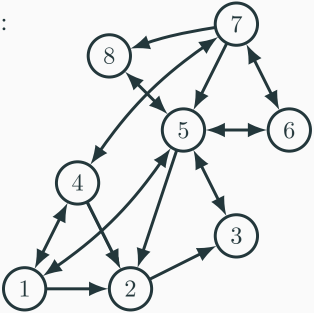

- A **non-private** solution [Nedic *et. al.*, 2010]:

$$x_i(k+1) = \text{proj}_X(z_i(k) - \alpha_k \nabla f_i(z_i(k)))$$

$$z_i(k) = \sum_{j=1}^{n} w_{ij} x_j(k)$$

$$\begin{cases} \sum \alpha_k = \infty \\ \sum \alpha_k^2 < \infty \end{cases}$$

Assumption:

- D is compact
- f_i's are strongly convex and C^2

- "Information": $F = (f_i)_{i=1}^n \in \mathcal{F}^n$

- "Information": $F = (f_i)_{i=1}^n \in \mathcal{F}^n$
- Given $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$ with $\mathcal{V} \subseteq \mathcal{F}$,

Adjacency

$F, F' \in \mathcal{F}^n$ are **$\mathcal{V}$-adjacent** if there exists $i_0 \in \{1, \ldots, n\}$ such that

$$f_i = f_i' \text{ for } i \neq i_0 \quad \text{and} \quad f_{i_0} - f_{i_0}' \in \mathcal{V}$$

- "Information": $F = (f_i)_{i=1}^n \in \mathcal{F}^n$
- Given $(\mathcal{V}, \|\cdot\|_{\mathcal{V}})$ with $\mathcal{V} \subseteq \mathcal{F}$,

Adjacency

$F, F' \in \mathcal{F}^n$ are **$\mathcal{V}$-adjacent** if there exists $i_0 \in \{1, \ldots, n\}$ such that

$$f_i = f'_i \text{ for } i \neq i_0 \quad \text{and} \quad f_{i_0} - f'_{i_0} \in \mathcal{V}$$

- For a **random map** $\mathcal{M} : \mathcal{F}^n \times \Omega \to \mathcal{X}$ and $\epsilon \in \mathbb{R}^n_{>0}$

Differential Privacy (DP)

$\mathcal{M}$ is **ϵ-DP** if

$\forall$ $\mathcal{V}$-adjacent $F, F' \in \mathcal{F}^n \quad \forall \mathcal{O} \subseteq X$

$$\mathbb{P}\{\mathcal{M}(F', \omega) \in \mathcal{O}\} \leq e^{\epsilon_{i_0} \|f_{i_0} - f'_{i_0}\|_{\mathcal{V}}} \mathbb{P}\{\mathcal{M}(F, \omega) \in \mathcal{O}\}$$

- Training records: $\{(a_j, b_j)\}_{j=1}^N$ where $a_j \in [0,1]^2$ and $b_j \in \{-1, 1\}$
- Goal: find the best separating hyperplane $x^T a$

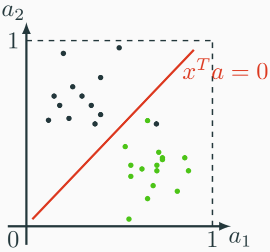

Case Study
Linear Classification with Logistic Loss Function

- Training records: $\{(a_j, b_j)\}_{j=1}^N$ where $a_j \in [0,1]^2$ and $b_j \in \{-1, 1\}$
- Goal: find the best separating hyperplane $x^T a$

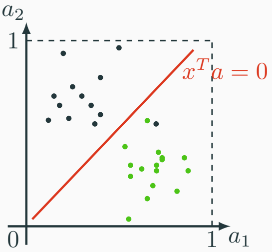

Convex Optimization Problem

$$x^* = \underset{x \in X}{\operatorname{argmin}} \quad \sum_{j=1}^N \left(\ell(x; a_j, b_j) + \frac{\lambda}{2} |x|^2 \right)$$

- Logistic loss: $\ell(x; a, b) = \ln(1 + e^{-b a^T x})$

- Training records: $\{(a_j, b_j)\}_{j=1}^N$ where $a_j \in [0,1]^2$ and $b_j \in \{-1, 1\}$
- Goal: find the best separating hyperplane $x^T a$

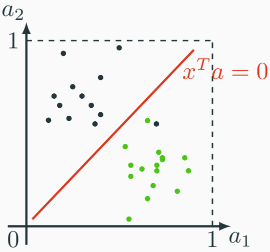

Convex Optimization Problem

$$x^* = \operatorname*{argmin}_{x \in X} \sum_{i=1}^{n} \sum_{j=1}^{N_i} \left(\ell(x; a_{i,j}, b_{i,j}) + \frac{\lambda}{2} |x|^2 \right)$$

- Logistic loss: $\ell(x; a, b) = \ln(1 + e^{-b a^T x})$

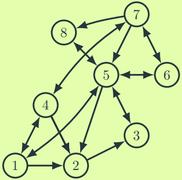

Message Perturbation vs. Objective Perturbation

A generic distributed optimization algorithm:

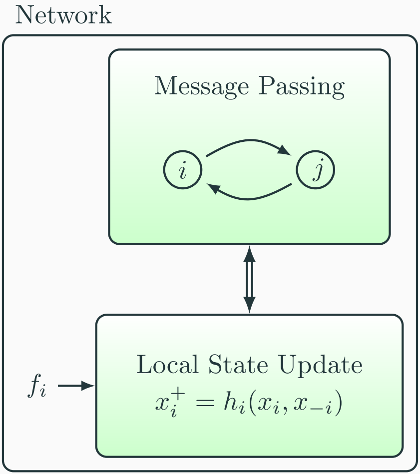

Message Perturbation vs. Objective Perturbation

Message Perturbation:

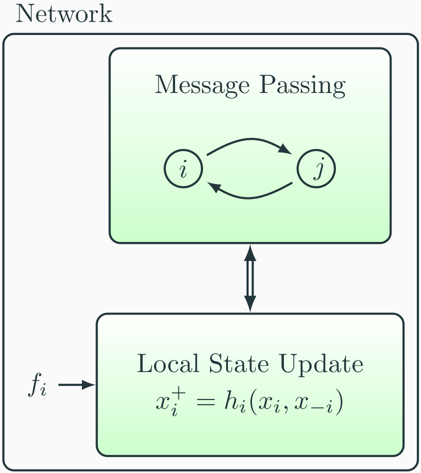

Objective Perturbation:

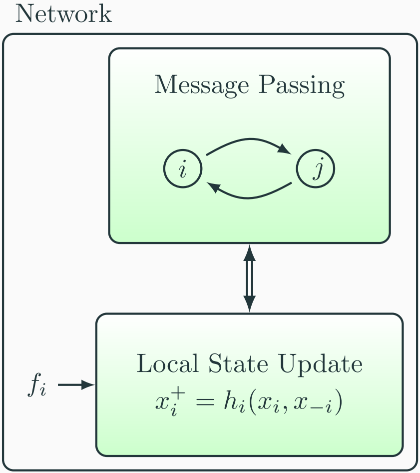

Message Perturbation vs. Objective Perturbation

Message Perturbation:

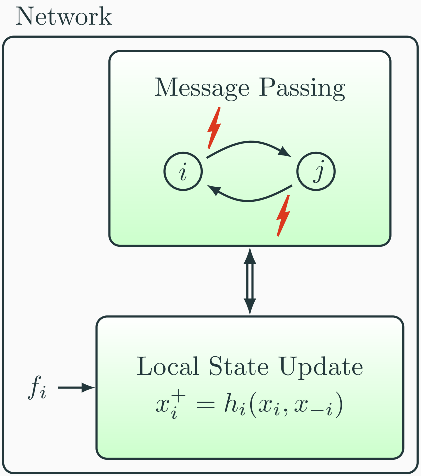

Objective Perturbation:

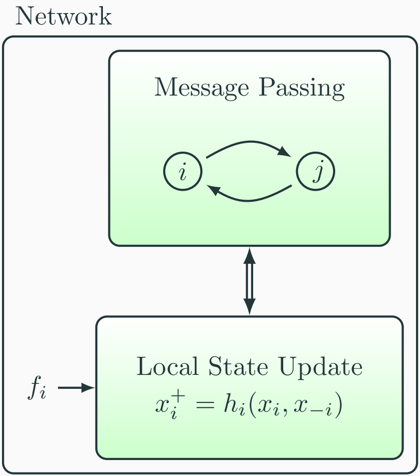

Message Perturbation vs. Objective Perturbation

Message Perturbation:

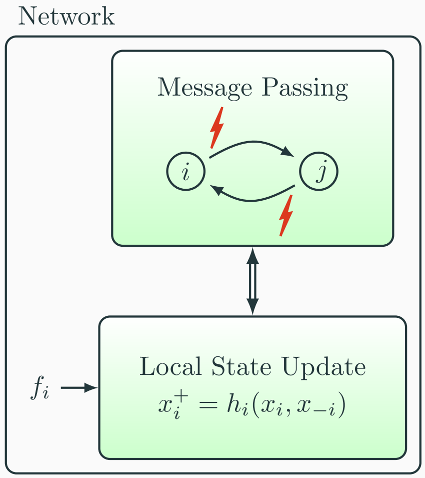

Objective Perturbation:

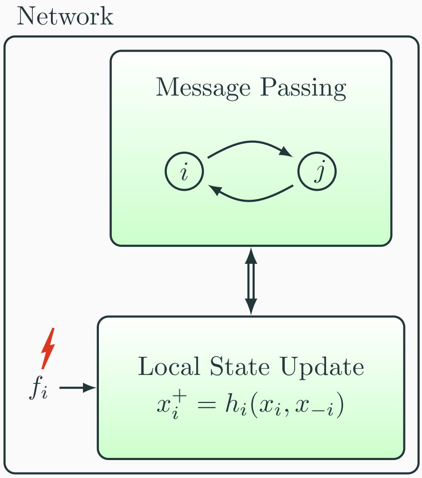

Impossibility Result

Generic message-perturbing algorithm:

$$x(k+1) = a_{\mathcal{I}}(x(k), \xi(k))$$
$$\xi(k) = x(k) + \eta(k)$$

Impossibility Result

Generic message-perturbing algorithm:

$$x(k+1) = a_{\mathcal{I}}(x(k), \xi(k))$$
$$\xi(k) = x(k) + \eta(k)$$

Theorem

If

- The $\eta \to x$ dynamics is **0-LAS**
- $\eta_i(k) \sim \text{Lap}(b_i(k))$ or $\eta_i(k) \sim \mathcal{N}(0, b_i(k))$
- $b_i(k)$ is $O(\frac{1}{k^p})$ for some $p > 0$

Then **no ϵ-DP** of the information set $\mathcal{I}$ for any $\epsilon > 0$

Impossibility Result: An Example

Algorithm proposed in [Huang *et. al.*, 2015]:

$$x_i(k+1) = \mathrm{proj}_X(z_i(k) - \alpha_k \nabla f_i(z_i(k)))$$

$$z_i(k) = \sum_{j=1}^{n} w_{ij} \xi_j(k)$$

$$\xi_j(k) = x_j(k) + \eta_j(k)$$

Impossibility Result: An Example

Algorithm proposed in [Huang *et. al.*, 2015]:

$$x_i(k+1) = \mathrm{proj}_X(z_i(k) - \alpha_k \nabla f_i(z_i(k)))$$

$$z_i(k) = \sum_{j=1}^{n} w_{ij} \xi_j(k)$$

$$\xi_j(k) = x_j(k) + \eta_j(k)$$

Impossibility Result: An Example

Algorithm proposed in [Huang *et. al.*, 2015]:

$$x_i(k+1) = \mathrm{proj}_X(z_i(k) - \alpha_k \nabla f_i(z_i(k)))$$

$$z_i(k) = \sum_{j=1}^{n} w_{ij} \xi_j(k)$$

$$\xi_j(k) = x_j(k) + \eta_j(k)$$

- $\eta_j(k) \sim \mathrm{Lap}(\propto p^k)$
- $\alpha_k \propto q^k$

$0 < q < p < 1$

Impossibility Result: An Example

Algorithm proposed in [Huang *et. al.*, 2015]:

$$x_i(k+1) = \mathrm{proj}_X(z_i(k) - \alpha_k \nabla f_i(z_i(k)))$$

$$z_i(k) = \sum_{j=1}^{n} w_{ij} \xi_j(k)$$

$$\xi_j(k) = x_j(k) + \eta_j(k)$$

- $\eta_j(k) \sim \mathrm{Lap}(\propto p^k)$
- $\alpha_k \propto q^k$ → Finite sum

$0 < q < p < 1$

Impossibility Result: An Example

Algorithm proposed in [Huang *et. al.*, 2015]:

- Simulation results for a linear classification problem:

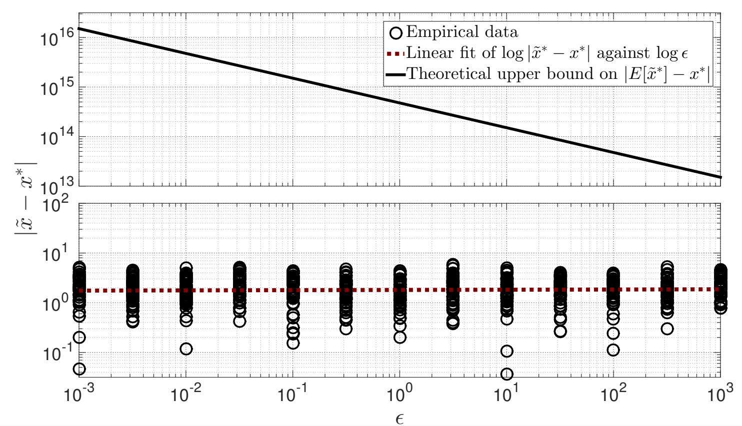

Outline

State of the Art

- [Chaudhuri *et. al.*, 2011]
    - First proposed "objective perturbation" by adding linear random functions
    - Extended by [Kifer *et. al.*, 2012] to constrained and non-differentiable problems
    - Preserves DP of objective function parameters

State of the Art

- [Chaudhuri *et. al.*, 2011]
  - First proposed “objective perturbation” by adding linear random functions
  - Extended by [Kifer *et. al.*, 2012] to constrained and non-differentiable problems
  - Preserves DP of objective function parameters

- [Zhang *et. al.*, 2012]
  - Proposed objective perturbation by adding sample path of Gaussian stochastic process
  - Preserves DP of objective function parameters

State of the Art

- [Chaudhuri *et. al.*, 2011]
  - First proposed "objective perturbation" by adding linear random functions
  - Extended by [Kifer *et. al.*, 2012] to constrained and non-differentiable problems
  - Preserves DP of objective function parameters

- [Zhang *et. al.*, 2012]
  - Proposed objective perturbation by adding sample path of Gaussian stochastic process
  - Preserves DP of objective function parameters

- [Hall *et. al.*, 2013]
  - Proposed objective perturbation by adding quadratic random functions
  - Preserves DP of objective function parameters

Prelim: Hillbert Spaces

- Hilbert space $\mathcal{H}$ = complete inner-product space
- Orthonormal basis $\{e_k\}_{k \in I} \subset \mathcal{H}$
- If $\mathcal{H}$ is separable:

$$h = \sum_{k=1}^{\infty} \overbrace{\langle h, e_k \rangle}^{\delta_k} e_k$$

Prelim: Hillbert Spaces

- Hilbert space $\mathcal{H}$ = complete inner-product space
- Orthonormal basis $\{e_k\}_{k \in I} \subset \mathcal{H}$
- If $\mathcal{H}$ is separable:

$$h = \sum_{k=1}^{\infty} \overbrace{\langle h, e_k \rangle}^{\delta_k} e_k$$

- For $D \subseteq \mathbb{R}^d$, $L_2(D)$ **is a separable Hilbert space** $\Rightarrow \mathcal{F} = L_2(D)$

Functional Perturbation via Laplace Noise

- Φ : coefficient sequence $\boldsymbol{\delta} \to$ function $h = \sum_{k=1}^{\infty} \delta_k e_k$
- Adjacency space:

$$\mathcal{V}_q = \left\{\Phi(\boldsymbol{\delta}) \mid \sum_{k=1}^{\infty} (k^q \delta_k)^2 < \infty\right\}$$

Functional Perturbation via Laplace Noise

- Φ : coefficient sequence $\boldsymbol{\delta} \to$ function $h = \sum_{k=1}^{\infty} \delta_k e_k$
- Adjacency space:

$$\mathcal{V}_q = \left\{\Phi(\boldsymbol{\delta}) \mid \sum_{k=1}^{\infty} (k^q \delta_k)^2 < \infty\right\}$$

- Random map:

$$\boxed{\mathcal{M}(f, \boldsymbol{\eta}) = \Phi\left(\Phi^{-1}(f) + \boldsymbol{\eta}\right) = f + \Phi(\boldsymbol{\eta})}$$

Functional Perturbation

Functional Perturbation via Laplace Noise

- Φ : coefficient sequence $\boldsymbol{\delta} \to$ function $h = \sum_{k=1}^{\infty} \delta_k e_k$
- Adjacency space:

$$\mathcal{V}_q = \left\{ \Phi(\boldsymbol{\delta}) \mid \sum_{k=1}^{\infty} (k^q \delta_k)^2 < \infty \right\}$$

- Random map:

$$\boxed{\mathcal{M}(f, \boldsymbol{\eta}) = \Phi\left(\Phi^{-1}(f) + \boldsymbol{\eta}\right) = f + \Phi(\boldsymbol{\eta})}$$

Functional Perturbation

Theorem

For $\eta_k \sim \text{Lap}(\frac{\gamma}{k^p})$, $q > 1$, and $p \in \left(\frac{1}{2}, q - \frac{1}{2}\right)$, $\mathcal{M}$ guarantees ϵ-DP with

$$\epsilon = \frac{1}{\gamma}\sqrt{\zeta(2(q-p))}$$

Outline

Resilience to Post-processing

Algorithm sketch:

1. Each agent **perturbs its own** objective function (offline)
2. Agents **participate in an arbitrary** distributed optimization algorithm with perturbed functions (online)

Resilience to Post-processing

Algorithm sketch:

1. Each agent **perturbs its own** objective function (offline)
2. Agents **participate in an arbitrary** distributed optimization algorithm with perturbed functions (online)

⬆

- $\mathcal{M} : L_2(D)^n \times \Omega \to L_2(D)^n$
- $\mathcal{F} : L_2(D)^n \to \mathcal{X}$, where $(\mathcal{X}, \Sigma_{\mathcal{X}})$ is an arbitrary measurable space

Corollary (special case of [Ny & Pappas 2014, Theorem 1])

If $\mathcal{M}$ is ϵ-DP, then $\mathcal{F} \circ \mathcal{M} : L_2(D)^n \times \Omega \to \mathcal{X}$ is ϵ-DP.

Ensuring Regularity of Perturbed Functions

- $\hat{f}_i = \mathcal{M}(f_i, \boldsymbol{\eta_i})$ may be discontinuous/non-convex/...

Ensuring Regularity of Perturbed Functions

- $\hat{f}_i = \mathcal{M}(f_i, \boldsymbol{\eta_i})$ may be discontinuous/non-convex/...

- $\mathcal{S} = \{\text{Regular functions}\} \subset C^2(D) \subset L_2(D)$

Ensuring Regularity of Perturbed Functions

- $\hat{f}_i = \mathcal{M}(f_i, \boldsymbol{\eta_i})$ may be discontinuous/non-convex/...
- $\mathcal{S} = \{\text{Regular functions}\} \subset C^2(D) \subset L_2(D)$
- **Ensuring Smoothness:** $C^2(D)$ is dense in $L_2(D)$ so

$$\forall \varepsilon_i > 0 \;\; \text{pick } \hat{f}_i^s \in C^2(D) \quad \text{such that} \quad \|\hat{f}_i - \hat{f}_i^s\| < \varepsilon_i$$

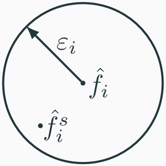

Ensuring Regularity of Perturbed Functions

- $\hat{f}_i = \mathcal{M}(f_i, \boldsymbol{\eta_i})$ may be discontinuous/non-convex/...
- $\mathcal{S} = \{\text{Regular functions}\} \subset C^2(D) \subset L_2(D)$
- **Ensuring Smoothness:** $C^2(D)$ is dense in $L_2(D)$ so

$$\forall \varepsilon_i > 0 \;\; \text{pick } \hat{f}_i^s \in C^2(D) \quad \text{such that} \quad \|\hat{f}_i - \hat{f}_i^s\| < \varepsilon_i$$

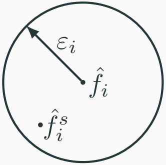

- **Ensuring Regularity:**

$$\tilde{f}_i = \text{proj}_{\mathcal{S}}(\hat{f}_i^s)$$

Proposition

$\mathcal{S}$ is convex and closed relative to $C^2(D)$

Algorithm

1. Each agent **perturbs** its function:

$$\hat{f}_i = \mathcal{M}(f_i, \boldsymbol{\eta}_i) = f_i + \Phi(\boldsymbol{\eta}_i), \quad \eta_{i,k} \sim \mathrm{Lap}(b_{i,k}), \quad b_{i,k} = \frac{\gamma_i}{k^{p_i}}$$

2. Each agent **selects** $\hat{f}_i^s \in \mathcal{S}_0$ such that

$$\|\hat{f}_i - \hat{f}_i^s\| < \varepsilon_i$$

3. Each agent **projects** $\hat{f}_i^s$ onto $\mathcal{S}$:

$$\tilde{f}_i = \mathrm{proj}_{\mathcal{S}}(\hat{f}_i^s)$$

4. Agents **participate** in *any* distributed optimization algorithm with $(\tilde{f}_i)_{i=1}^n$

Algorithm

Offline:

1. Each agent **perturbs** its function:

$$\hat{f}_i = \mathcal{M}(f_i, \boldsymbol{\eta}_i) = f_i + \Phi(\boldsymbol{\eta}_i), \quad \eta_{i,k} \sim \mathrm{Lap}(b_{i,k}), \quad b_{i,k} = \frac{\gamma_i}{k^{p_i}}$$

2. Each agent **selects** $\hat{f}_i^s \in \mathcal{S}_0$ such that

$$\|\hat{f}_i - \hat{f}_i^s\| < \varepsilon_i$$

3. Each agent **projects** $\hat{f}_i^s$ onto $\mathcal{S}$:

$$\tilde{f}_i = \mathrm{proj}_{\mathcal{S}}(\hat{f}_i^s)$$

4. Agents **participate** in *any* distributed optimization algorithm with $(\tilde{f}_i)_{i=1}^n$

Accuracy Analysis

- Set of "regular" functions:

$$\mathcal{S} = \{h \in C^2(D) \mid \alpha I_d \leq \nabla^2 h(x) \leq \beta I_d \text{ and } |\nabla h(x)| \leq \overline{u}\}$$

Lemma ($\mathcal{K}$-Lipschitzness of argmin)

For $f, g \in \mathcal{S}$,

$$\Big| \operatorname*{argmin}_{x \in X} f - \operatorname*{argmin}_{x \in X} g \Big| \leq \kappa_{\alpha,\beta}(\|f - g\|)$$

- Set of "regular" functions:

$$\mathcal{S} = \{h \in C^2(D) \mid \alpha I_d \leq \nabla^2 h(x) \leq \beta I_d \text{ and } |\nabla h(x)| \leq \overline{u}\}$$

Lemma ($\mathcal{K}$-Lipschitzness of argmin)

For $f, g \in \mathcal{S}$,

$$\Big| \operatorname*{argmin}_{x \in X} f - \operatorname*{argmin}_{x \in X} g \Big| \leq \kappa_{\alpha,\beta}(\|f - g\|)$$

- Define

$$\tilde{x}^* = \operatorname*{argmin}_{x \in X} \sum_{i=1}^n \tilde{f}_i \quad \text{and} \quad x^* = \operatorname*{argmin}_{x \in X} \sum_{i=1}^n f_i,$$

Theorem (Accuracy)

$$\mathbb{E}\,|\tilde{x}^* - x^*| \leq \sum_{i=1}^n \kappa_n\left(\frac{\zeta(q_i)}{\epsilon_i}\right) + \kappa_n(\varepsilon_i)$$

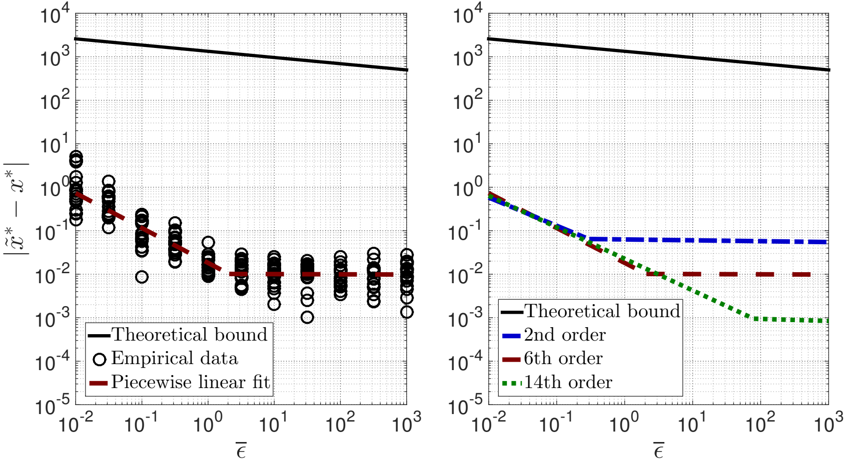

$|\tilde{x}^* - x^*|$
$\overline{\epsilon}$
Theoretical bound
Empirical data
Piecewise linear fit
Theoretical bound
2nd order
6th order
14th order

Conclusions and Future Work

In this talk, we

- Proposed a definition of DP for functions
- Illustrated a fundamental limitation of message-perturbing strategies
- Proposed the method of functional perturbation
- Discussed how functional perturbation can be applied to distributed convex optimization

Conclusions and Future Work

In this talk, we

- Proposed a definition of DP for functions
- Illustrated a fundamental limitation of message-perturbing strategies
- Proposed the method of functional perturbation
- Discussed how functional perturbation can be applied to distributed convex optimization

Future work includes

- relaxation of the smoothness, convexity, and compactness assumptions
- comparing the numerical efficiency of different bases for L_2
- characterizing the expected sub-optimality gap of the algorithm and the optimal privacy-accuracy trade-off curve
- further understanding the appropriate scales of privacy parameters for particular applications

Questions and Comments

Full results of this talk available in:

E. Nozari, P. Tallapragada, J. Cortés, "Differentially Private Distributed Convex Optimization via Functional Perturbation," *IEEE Trans. on Control of Net. Sys.*, provisionally accepted, `http://arxiv.org/abs/1512.00369`

Formal Definition

in original context [Dwork *et. al.*, 2006]

Context:

- $D \in \mathcal{D}$: A database of records
- Adjacency: $D_1, D_2 \in \mathcal{D}$ are adjacent if they differ by at most 1 record
- $(\Omega, \Sigma, \mathbb{P})$: Probability space
- $q : \mathcal{D} \to X$: (Honest) query function
- $\mathcal{M} : \mathcal{D} \times \Omega \to X$: Randomized/sanitized query function
- $\epsilon > 0$: Level of privacy

Definition

$\mathcal{M}$ is ϵ-DP if

$$\forall \text{ adjacent } D_1, D_2 \in \mathcal{D} \quad \forall \mathcal{O} \subseteq X \qquad \mathbb{P}\{\mathcal{M}(D_1) \in \mathcal{O}\} \le e^{\epsilon} \mathbb{P}\{\mathcal{M}(D_2) \in \mathcal{O}\}$$

- Adjacency is symmetric $\Rightarrow \begin{cases} \mathbb{P}\{\mathcal{M}(D_1) \in \mathcal{O}\} \le e^{\epsilon} \mathbb{P}\{\mathcal{M}(D_2) \in \mathcal{O}\} \\ \mathbb{P}\{\mathcal{M}(D_2) \in \mathcal{O}\} \le e^{\epsilon} \mathbb{P}\{\mathcal{M}(D_1) \in \mathcal{O}\} \end{cases}$

Formal Definition: Geometric Interpretation
in original context

Definition

$\mathcal{M}$ is ϵ-DP if

$$\forall \text{ adjacent } D_1, D_2 \in \mathcal{D} \quad \forall \mathcal{O} \subseteq X \qquad \mathbb{P}\{\mathcal{M}(D_1) \in \mathcal{O}\} \leq e^{\epsilon} \mathbb{P}\{\mathcal{M}(D_2) \in \mathcal{O}\}$$

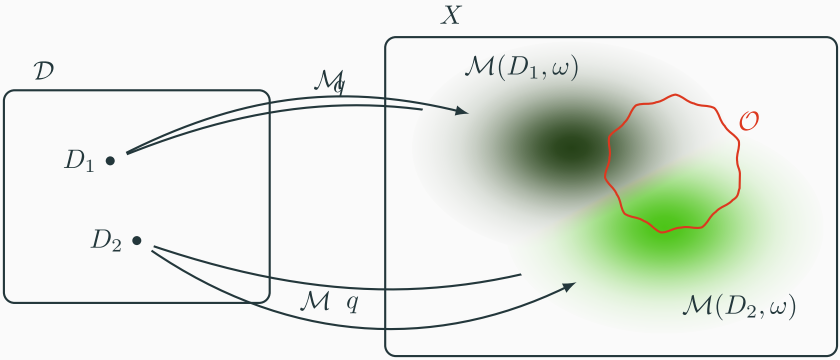

Operational Meaning of DP

A binary decision example [Geng&Pramod, 2013]

- Adversary's decision $= \begin{cases} \text{TRUE} & \text{if} \quad \mathcal{M}(D,\omega) \in \mathcal{O} \\ \text{FALSE} & \text{if} \quad \mathcal{M}(D,\omega) \in \mathcal{O}^c \end{cases}$

- MD $= \{\mathcal{M}(D_1,\omega) \in \mathcal{O}^c\}$
- FA $= \{\mathcal{M}(D_2,\omega) \in \mathcal{O}\}$

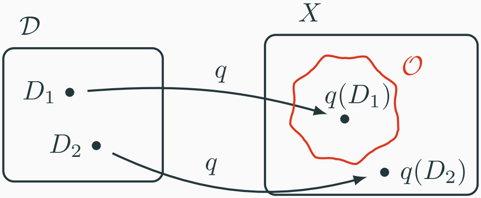

- If $\mathcal{M}$ is ϵ-DP then

$$\begin{cases} \mathbb{P}\{\mathcal{M}(D_1,\omega) \in \mathcal{O}\} \le e^\epsilon \mathbb{P}\{\mathcal{M}(D_2,\omega) \in \mathcal{O}\} \\ \mathbb{P}\{\mathcal{M}(D_2,\omega) \in \mathcal{O}^c\} \le e^\epsilon \mathbb{P}\{\mathcal{M}(D_1,\omega) \in \mathcal{O}^c\} \end{cases} \Rightarrow \begin{cases} 1 - p_{\mathrm{MD}} \le e^\epsilon p_{\mathrm{FA}} \\ 1 - p_{\mathrm{FA}} \le e^\epsilon p_{\mathrm{MD}} \end{cases}$$

$$\Rightarrow p_{\mathrm{MD}}, p_{\mathrm{FA}} \ge \frac{e^\epsilon - 1}{e^{2\epsilon} - 1}$$

Generalizing the Definition: Using Metrics
[Chatzikokolakis *et. al.*, 2013]

- If D_1, D_2 differ in N elements then

$$\mathbb{P}\{\mathcal{M}(D_1,\omega) \in \mathcal{O}\} \leq e^{N\epsilon}\mathbb{P}\{\mathcal{M}(D_2,\omega) \in \mathcal{O}\}$$

- $d : \mathcal{D} \times \mathcal{D} \to [0,\infty)$ metric on $\mathcal{D}$

Definition –revisited

$\mathcal{M}$ gives/preserves ϵ-differential privacy if

$\forall D_1, D_2 \in \mathcal{D} \quad \forall \mathcal{O} \subseteq X$ we have

$$\mathbb{P}\{\mathcal{M}(D_1,\omega) \in \mathcal{O}\} \leq e^{\epsilon d(D_1, D_2)}\mathbb{P}\{\mathcal{M}(D_2,\omega) \in \mathcal{O}\}$$