When Worst-Case Sensitivity is Atypical

Reminder

Definition: the Global Sensitivity of a function $f : D^n \to \mathbb{R}^d$

$$GS_f = \max_{x,y:d(x,y)=1} \|f(x) - f(y)\|_1$$

with $d(x, y) = |\{i : x_i \neq y_i\}|$

An Example: the Median

Let $f_{med}(x) = \text{median}(x_1, \ldots, x_n)$

x_i are real numbers from a bounded interval $D = [0, \Lambda]$

Suppose for simplicity that x_i are ordered and n is odd.

$x_1 \leq \cdots \leq x_n$. Let $m = \frac{n+1}{2}$ the rank of the median.

The *global sensitivity* $GS_{f_{med}}$ is the size of the range.

$GS_{f_{med}} = \Lambda$

Take: $x_1 = \cdots = x_m = 0$ and $x_{m+1} = \cdots = x_n = \Lambda$

then $f_{med}(x_1, \ldots, x_n) = 0$

but $f_{med}(x_1, \ldots, x_{m-1}, \Lambda, x_{m+1}, \ldots, x_n) = \Lambda$

However, on typical inputs, f_{med} is not very sensitive

Definition: Local Sensitivity

For $f : D^n \to \mathbb{R}^d$ *and* $x \in D^n$, *the local sensitivity of* f *at* x

$$LS_f(x) = \max_{y:d(x,y)=1} \|f(x) - f(y)\|_1$$

Notice that $GS_f = \max_x LS_f(x)$

Back to the Median Example

$$LS_{f_{med}}(x) = \max\left(x_m - x_{m-1}, x_{m+1} - x_m\right)$$

(Explain on board)

We Want to Use Noise Proportional to $LS_f(x)$

But, as $LS_f(x)$ a function of x, its magnitude can reveal information

Example

Consider the two neighbouring datasets:

$$x_1 = \cdots = x_{m+1} = 0, x_{m+2} = \cdots = x_n = \Lambda$$

$$x_1 = \cdots = x_m = 0, x_{m+1} = \cdots = x_n = \Lambda$$

booth have a zero median,

but the local sensitivity is zero for the first, and Λ for the second

Definition: A smooth Bound on *LS*

For $\beta > 0$, a function $S : D^n \to \mathbb{R}^+$ is a β-smooth upper bound on the local sensitivity of f if it satisfies the following requirements:

$$\forall x \in D^n : \quad S(x) \geq LS_f(x) ; \tag{1}$$

$$\forall x, y \in D^n, d(x, y) = 1 : \quad S(x) \leq e^{\beta} \cdot S(y) . \tag{2}$$

- e.g. GS_f is zero-smooth
- As β increases the bound gets tighter, and is allowed to fluctuate more

Definition: S^* The minimal β-smooth bound

For $\beta > 0$, the β-smooth sensitivity of f is

$$S^*_{f,\beta}(x) = \max_{y \in D^n} \left(LS_f(y) \cdot e^{-\beta d(x,y)} \right)$$

Lemma

$S^*_{f,\beta}$ is a β-smooth upper bound on LS_f

$S^*_{f,\beta}(x) \leq S(x)$ for all $x \in D^n$ for every β-smooth upper bound S on LS_f.

Proof $S^*_{f,\beta}$ *is a* β*-smooth upper bound on* LS_f

$S^*_{f,\beta}(x) \geq LS_f(x)$ this is obvious from the definition of smooth sensitivity

next we show it is β-smooth:

Take x,y neighbor datasets, set x' such that $S^*_{f,\beta}(x) = LS_f(x') \cdot e^{-\beta d(x,x')}$

It holds that $d(y,x') \leq d(y,x) + d(x,x') = d(x,x') + 1$

$$
\begin{aligned}
S^*_{f,\beta}(y) \geq LS_f(x') \cdot e^{-\beta d(y,x')} &\geq LS_f(x') \cdot e^{-\beta(d(x,x')+1)} \\
&= e^{-\beta} \cdot LS_f(x') \cdot e^{-\beta d(x,x')} \\
&= e^{-\beta} \cdot S^*_{f,\beta}(x)
\end{aligned}
$$

Proof the minimality of $S^*_{f,\beta}(x)$

Let S be a β-smooth bound

it is enough to establish that $S(x) \geq LS_f(y) \cdot e^{-\beta d(x,y)}$ for all $x, y \in D^n$.

induction on $d(x, y)$

base case, $S(x) \geq LS_f(x)$ (by the definition of a β-smooth bound)

suppose $S(x') \geq LS_f(y) \cdot e^{-\beta d(x',y)}$ for all x', y at distance k

Consider x, y at distance $k + 1$. There exists x': $d(x, x') = 1, d(x', y) = k$

by the definition of a β-smooth bound $S(x) \geq S(x') \cdot e^{-\beta}$

IH $\Rightarrow$ $S(x) \geq LS_f(y) \cdot e^{-\beta(d(x',y)+1)} = LS_f(y) \cdot e^{-\beta d(x,y)}$

Calibrating Noise to Smooth Upper Bounds on LS_f

Main result (1D Case)

Let $f : D_n \to \mathbb{R}$ *be any real-valued function*

$S : D^n \to \mathbb{R}$ *be a* β*-smooth upper bound on the local sensitivity of* f

1. *If* $\beta \leq \frac{\epsilon}{2(\gamma+1)}$ *and* $\gamma > 1$, *the algorithm* $x \mapsto f(x) + \frac{2(\gamma+1)S(x)}{\epsilon} \cdot \eta$, *where* η *is sampled from distribution with density* $h(z) \propto \frac{1}{1+|z|^\gamma}$, *is* ϵ*-differentially private.*
2. *If* $\beta \leq \frac{\epsilon}{2\ln(\frac{2}{\delta})}$ *and* $\delta \in (0,1)$, *the algorithm* $x \mapsto f(x) + \frac{2S(x)}{\epsilon} \cdot \eta$, *where* $\eta \sim \mathrm{Lap}(1)$, *is* $\left(\epsilon, \frac{\delta}{2}(e^{\frac{\epsilon}{2}} + 1)\right)$ -differentially private

For the proof, we need more definitions
but before that - we will see an example for computing S*(x)

Computing Smooth Sensitivity for the Median

Another definition: *The sensitivity of f at distance k is*

$$A^{(k)}(x) = \max_{y \in D^n:\ d(x,y) \le k} LS_f(y)$$

Now notice that it can be used to express the *smooth sensitivity*:

$$S^*_{f,\beta}(x) = \max_{y \in D^n} \left(LS_f(y) \cdot e^{-\beta d(x,y)} \right) = \max_{k=0,1,\ldots,n} e^{-\beta k} \left(\max_{y:\ d(x,y)=k} LS_f(y) \right)$$

$$= \max_{k=0,1,\ldots,n} e^{-\beta k} A^{(k)}(x)$$

Smooth Sensitivity for the Median

Definition: the *smooth sensitivity* of the *median* is

$$S^*_{f_{med},\beta}(x) = \max_{k=0,\ldots,n} \left(e^{-k\beta} \cdot \underbrace{\max_{t=0,\ldots,k+1} \left(x_{m+t} - x_{m+t-k-1} \right)}_{} \right)$$

By prev slide suffice to show this is $A^{(k)}(x)$

Compute naively in $O(n^2)$

$$S^*_{f_{med},\beta}(x) = \max_{k=0,\dots,n} \left(e^{-k\beta} \cdot \max_{t=0,\dots,k+1} (x_{m+t} - x_{m+t-k-1}) \right)$$

Example

Consider the dataset where points are bounded and evenly spread in the interval [0,1].

that is $\hat{x}_i = \frac{i}{n}$ for $i = 1, \dots, n$

According to the definition above:

$$S^*(\hat{x}) = \max_k \left(e^{-k\beta} \cdot \frac{k+1}{n} \right) = e^{\beta - 1} \frac{1}{\beta n} \geq \frac{1}{n}$$

Completing the proof for $S^*_{f_{med},\beta}(x)$

Need to show that $A^{(k)}(x) = \max_{y:\ d(x,y)\le k} LS(y) = \max_{0\le t\le k+1}(x_{m+t} - x_{m+t-k-1})$

Observation#1: after changing k points in the data, the median must be in the interval $[x_{m-k}, x_{m+k}]$

Observation#2: The local sensitivity at distance k is maximized when the new median is an endpoint of a large empty interval

- This is achieved by taking the largest interval between two points that are *k*-apart.
- Then “pushing” the *k* point in the interval left and right

Calibrating Noise to Smooth Upper Bounds on LS_f

Back to Main result (1D Case)

Let $f : D_n \to \mathbb{R}$ *be any real-valued function*

$S : D^n \to \mathbb{R}$ *be a* β*-smooth upper bound on the local sensitivity of* f

1. *If* $\beta \leq \frac{\epsilon}{2(\gamma+1)}$ *and* $\gamma > 1$, *the algorithm* $x \mapsto f(x) + \frac{2(\gamma+1)S(x)}{\epsilon} \cdot \eta$, *where* η *is sampled from distribution with density* $h(z) \propto \frac{1}{1+|z|^\gamma}$, *is* ϵ*-differentially private.*
2. *If* $\beta \leq \frac{\epsilon}{2\ln(\frac{2}{\delta})}$ *and* $\delta \in (0,1)$, *the algorithm* $x \mapsto f(x) + \frac{2S(x)}{\epsilon} \cdot \eta$, *where* $\eta \sim \text{Lap}(1)$, *is* $\left(\epsilon, \frac{\delta}{2}(e^{\frac{\epsilon}{2}} + 1)\right)$ -differentially private

Admissible Noise Distributions

Notation: For $\mathcal{S} \subset \mathbb{R}^d$ $\quad \mathcal{S} + \Delta := \{z + \Delta \mid z \in \mathcal{S}\}$

$$e^{\lambda}\mathcal{S} := \{e^{\lambda} \cdot z \mid z \in \mathcal{S}\}$$

Definition: Given ϵ, δ A probability distribution on $\mathbb{R}^d$ (with pdf h) is

$(\alpha, \beta) - admissible$ where α, β are a function of ϵ, δ if:

for all $\Delta \in \mathbb{R}^d$ *and* $\lambda \in \mathbb{R}$ *satisfying* $\|\Delta\|_1 \leq \alpha$ *and* $|\lambda| \leq \beta$,

for all measurable subsets $\mathcal{S} \subseteq \mathbb{R}^d$:

Sliding Property: $\quad \Pr_{Z \sim h}\left[Z \in \mathcal{S}\right] \leq e^{\frac{\epsilon}{2}} \cdot \Pr_{Z \sim h}\left[Z \in \mathcal{S} + \Delta\right] + \frac{\delta}{2}$

Dilation Property: $\quad \Pr_{Z \sim h}\left[Z \in \mathcal{S}\right] \leq e^{\frac{\epsilon}{2}} \cdot \Pr_{Z \sim h}\left[Z \in e^{\lambda} \cdot \mathcal{S}\right] + \frac{\delta}{2}$

Notice that for every event S: $\frac{\int_S h(z)dz}{\int_{S\cdot e^\lambda} h(z)dz} = \frac{\int_S h(z)dz}{e^\lambda \int_S h(e^\lambda y)dy}$ hence we can illustrate the two conditions as follow:

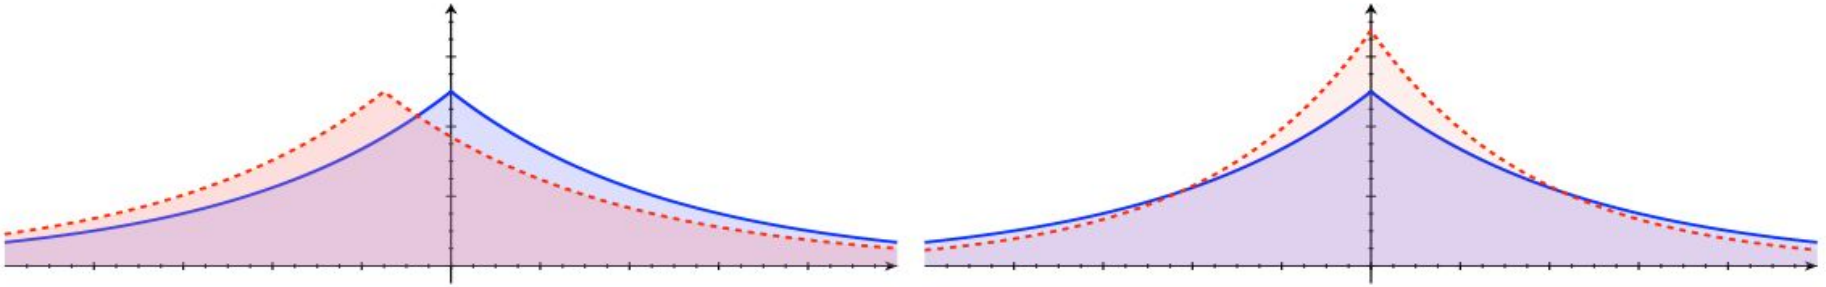

Figure 1: Sliding and dilation for the Laplace distribution with p.d.f. $h(z) = \frac{1}{2}e^{-|z|}$, plotted as a solid line. The dotted lines plot the densities $h(z + 0.3)$ (left) and $e^{0.3}h(e^{0.3}z)$ (right).

Lemma

Let h be an (α, β)-admissible noise probability density function, $Z \sim h$

For a function $f : D^n \to \mathbb{R}^d$, let $S : D^n \to \mathbb{R}$ be a β-smooth upper bound

Then algorithm $\mathcal{A}(x) = f(x) + \frac{S(x)}{\alpha} \cdot Z$ is $\left(\epsilon, \frac{\delta}{2}(e^{\frac{\epsilon}{2}} + 1)\right)$ -differentially private

Proof

For all neighboring $x, y \in D^n$ and all sets $\mathcal{S}$, we need to show that

$$\Pr[\mathcal{A}(x) \in \mathcal{S}] \leq e^{\epsilon} \cdot \Pr[\mathcal{A}(y) \in \mathcal{S}] + \frac{\delta}{2}\left(e^{\frac{\epsilon}{2}} + 1\right)$$

Denote: $N(x) = \frac{S(x)}{\alpha}$, $\mathcal{S}_1 = \frac{\mathcal{S} - f(x)}{N(x)}$, $\mathcal{S}_2 = \mathcal{S}_1 + \frac{f(y) - f(x)}{N(x)} = \frac{\mathcal{S} - f(y)}{N(x)}$, $\mathcal{S}_3 = \mathcal{S}_2 \cdot \frac{N(x)}{N(y)} = \frac{\mathcal{S} - f(y)}{N(y)}$

Observe that $\mathcal{A}(x) \in \mathcal{S}$ if and only if $Z \in \mathcal{S}_1$

$$\begin{aligned}
\Pr[\mathcal{A}(x) \in \mathcal{S}] &= \Pr_{z \sim h}[z \in \mathcal{S}_1] \\
&\leq \Pr_{z \sim h}[z \in \mathcal{S}_2] \cdot e^{\epsilon/2} + \frac{\delta}{2} \\
&\leq \Pr_{z \sim h}[z \in \mathcal{S}_3] \cdot e^{\epsilon} + \frac{\delta}{2} \cdot e^{\epsilon/2} + \frac{\delta}{2} \\
&= Pr\left[\mathcal{A}(y) \in \mathcal{S}\right] \cdot e^{\epsilon} + \frac{\delta}{2}(e^{\frac{\epsilon}{2}} + 1)
\end{aligned}$$

Lemma we just proved

Let h be an (α, β)-admissible noise probability density function, $Z \sim h$

For a function $f : D^n \to \mathbb{R}^d$, let $S : D^n \to \mathbb{R}$ be a β-smooth upper bound

Then algorithm $\mathcal{A}(x) = f(x) + \frac{S(x)}{\alpha} \cdot Z$ is $\left(\epsilon, \frac{\delta}{2}(e^{\frac{\epsilon}{2}} + 1)\right)$ -differentially private

Now we can implement this lemma for a given distribution. For instance as was stated:

1. *If $\beta \leq \frac{\epsilon}{2(\gamma+1)}$ and $\gamma > 1$, the algorithm* $x \mapsto f(x) + \frac{2(\gamma+1)S(x)}{\epsilon} \cdot \eta$, *where η is sampled from distribution with density $h(z) \propto \frac{1}{1+|z|^\gamma}$, is ϵ-differentially private.*

We just need to show that the distribution is admissible for some α, β

Distribution with pdf $h(z) \propto \frac{1}{1+|z|^\gamma}$ is Admissible

Lemma:

For any $\gamma > 1$, the distribution with density $h(z) \propto \frac{1}{1+|z|^\gamma}$ is $(\frac{\epsilon}{2(\gamma+1)}, \frac{\epsilon}{2(\gamma+1)})$-admissible

Notice the parameters are independent of δ , hence it can be set to 0.

Dilation Property: $\Pr_{Z\sim h}\left[Z\in\mathcal{S}\right]\le e^{\frac{\epsilon}{2}}\cdot\Pr_{Z\sim h}\left[Z\in e^{\lambda}\cdot\mathcal{S}\right]+\frac{\delta}{2}$

Proof

Dilation by e^{λ}: assume $|\lambda|\le\frac{\epsilon}{2(\gamma+1)}$ we need to show: $\frac{\int_S h(z)dz}{\int_{e^{\lambda}\cdot S}h(z)dz}\le e^{\frac{\epsilon}{2}}$

$$\frac{\int_S h(z)dz}{\int_{S\cdot e^{\lambda}}h(z)dz}=\frac{\int_S h(z)dz}{e^{\lambda}\int_S h(e^{\lambda}y)dy}=\frac{\int_S\frac{1}{1+|z|^{\gamma}}dz}{e^{\lambda}\int_S\frac{1}{1+(e^{\lambda}|y|)^{\gamma}}dy}$$

It is sufficient to show: $\frac{1+(e^{\lambda}|z|)^{\gamma}}{e^{\lambda}(1+|z|^{\gamma})}\le e^{\frac{\epsilon}{2}}$

If $\lambda\ge 0$ then $\frac{1+(e^{\lambda}|z|)^{\gamma}}{1+|z|^{\gamma}}\le\frac{(e^{\lambda}|z|)^{\gamma}}{|z|^{\gamma}}=e^{\lambda\gamma}\Longrightarrow\frac{1+(e^{\lambda}|z|)^{\gamma}}{e^{\lambda}(1+|z|^{\gamma})}\le e^{\lambda(\gamma-1)}$

which is $\le e^{\frac{\epsilon}{2}}$ by the assumption on λ

If $\lambda<0$ then $\frac{1+(e^{\lambda}|z|)^{\gamma}}{1+|z|^{\gamma}}\le 1\Longrightarrow\frac{1+(e^{\lambda}|z|)^{\gamma}}{e^{\lambda}(1+|z|^{\gamma})}\le e^{-\lambda}\le e^{\frac{\epsilon}{2(\gamma+1)}}\le e^{\frac{\epsilon}{2}}$

Sliding Property: $\Pr_{Z\sim h}\left[Z \in \mathcal{S}\right] \leq e^{\frac{\epsilon}{2}} \cdot \Pr_{Z\sim h}\left[Z \in \mathcal{S}+\Delta\right] + \frac{\delta}{2}$

Proof

Sliding by Δ : assume $|\Delta| \leq \frac{\epsilon}{2(\gamma+1)}$ define: $\phi(z) = \ln(1+z^{\gamma})$

we get: $\ln\left(\frac{h(z)}{h(z)+\Delta}\right) = \phi(|z+\Delta|) - \phi(|z|)$

By Lagrange's mean value theorem $\exists \zeta > 0. |\phi(|z+\Delta|) - \phi(|z|)| = |\Delta\phi'(\zeta)|$

$|\Delta\phi'(\zeta)|$ is bounded by $|\Delta|\gamma$ since $\phi'(\zeta) = \frac{\gamma\zeta^{\gamma-1}}{1+\zeta^{\gamma}} = \frac{\gamma}{\zeta+\zeta^{-(\gamma-1)}}$ (as $\gamma > 1$)

By the assumption on Δ : $|\Delta|\gamma \leq \frac{\epsilon\gamma}{2(\gamma+1)} \leq \frac{\epsilon(\gamma+1)}{2(\gamma+1)} = \frac{\epsilon}{2}$