Recovery of sparse heat source signals from locally differentially private sensor data via constrained ℓ_1-minimisation

Anna Gilbert and Audra McMillan
Department of Mathematics
University of Michigan
Ann Arbor, Michigan
Email: amcm@umich.edu

***Abstract*—In recent years, wireless technology has allowed the power of lightweight (thermal, light, motion, etc.) sensors to be explored. This data offers important benefits to society. For example, thermal sensor data now plays an important role in controlling HVAC systems and minimising energy consumption in smart buildings. Simultaneously, we have begun to understand the extent to which our privacy is compromised by allowing this increased level of data collection. In this work we explore the utility of locally differentially private thermal sensor data. Our work indicates that it is possible to produce locally private sensor measurements that both keep the exact locations of the heat sources private and permit recovery of the "general vicinity" of the sources.**

We consider the discretisation of the 1-dimensional heat equation on an unbounded domain. Let $u(x,t) : \mathbb{R} \times [0,\infty) \to \mathbb{R}$ be the temperature at location x at time t, $f(x)$ be the initial temperature (with bounded support) and μ be the diffusion constant. Let $T = \mu t$, $n > 0$ and suppose the support of f is contained in the discrete set $\{\frac{1}{n}, \cdots, 1\}$ so f can be represented by a vector $f_0 \in \mathbb{R}^n$ where $f(x) = \sum_{i=1}^{n} (f_0)_i \delta(\frac{i}{n} - x)$. Let $m > 0$ and suppose we take m measurements at locations $\frac{1}{m}, \cdots, 1$. Then our thermal measurements $y \in \mathbb{R}^m$ are given by

$$y = Af_0 \quad \text{where} \quad A_{ij} = g\left(\frac{i}{n} - \frac{j}{m}, t\right),$$

where $g(x,t) = \frac{1}{\sqrt{4\pi\mu t}} e^{\frac{-|x|^2}{4\mu t}}$. The heat source recovery problem is recover f_0 from a possibly noisy version of y.

We say two source vectors f_0 and f_0' are *neighbours* if they differ in location of a single source by one unit. For example, the following two source vectors are neighbours (where $n = 5$).

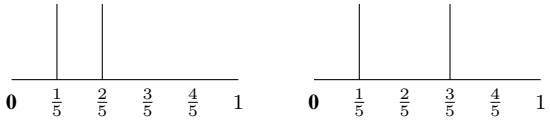

Our privacy goal is to output a perturbed version, $\tilde{y}$ of y such that an adversary can not determine with good confidence whether the source vector was f_0 or it's neighbour f_0'. That is, we design an algorithm that is locally differentially private, where "neighbouring databases" are replaced with "neighbouring source vectors". The "local" part refers to the fact that the measurements are made private before they are sent to the data analyst. This is desirable for sensor measurements since the data analyst is often the person the consumer would like to be protected against. Local differential privacy is also attractive from an implementation perspective. It is often the case with wireless sensors that the data must be communicated via some untrusted channel. Usually this step would involve encrypting the data, incurring significant computational and communication overhead. However, if the data is made private prior to being sent, then an argument can then be made that it no longer needs to be encrypted.

In this paper, we design algorithms $\mathcal{A}$ and R that satisfy the following diagram and criterion:

$$f_0 \xrightarrow{A} \left\{ \begin{array}{ccc} y_1 & \xrightarrow{\mathcal{A}} & \tilde{y_1} \\ \vdots & & \vdots \\ y_m & \xrightarrow{\mathcal{A}} & \tilde{y_m} \end{array} \right\} \xrightarrow{R} \hat{f}$$

1) **(Privacy) For all neighbouring source vectors f_0 and f_0', sensor locations s and Borel measurable sets E we have**

$$\mathbb{P}(\mathcal{A}(u_{f_0}(s,t)) \in E) \leq e^{\epsilon} \mathbb{P}(\mathcal{A}(u_{f_0'}(s,t)) \in E) + \delta.$$

2) **(Utility) The true source vector, f_0, and the recovered estimate $\hat{f}$ are close in the Earth Mover Distance (EMD).**

In order to prove that we have good recovery, we assume that f_0 is a sparse vector, the sources are well-separated and the measurements are taken within the right time interval. A picture of the set-up for 2 sources is provided below:

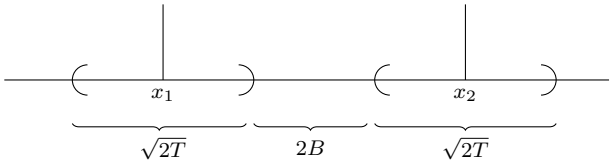

The algorithm $\mathcal{A}$ is derived from the Gaussian mechanism from the differential privacy literature [1]. We prove that adding independent Gaussian noise to each coordinate of y with standard deviation $\sigma = O\left(\frac{2\log(1.25/\delta)\sqrt{m}}{\epsilon n T^{1.5}}\right)$ is sufficient to satisfy our privacy constraint. The

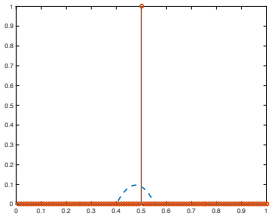

Fig. 1: The red line represents f_0 and the blue (dashed) line represents $\hat{f} = R(\tilde{y})$ (both normalised to have unit ℓ_1 norm) where $n = 100$, $m = 50$, $T = 0.5$ and $\sigma = 0.5$.

TABLE I: Asymptotic upper bounds for noisy recovery

VARIABLE	$\text{EMD}\left(\frac{f_0}{\|f_0\|_1}, \frac{\hat{f}}{\|\hat{f}\|_1}\right)$
n	$O(1)$
m	$O(1)$
T	$O(\sqrt{T} + T^2)$
σ	$O(\sigma)$

recovery algorithm R is constrained ℓ_1-minimisation, introduced by Li et al. [2]. The upper bound we prove for noisy recovery is a generalisation of a result of Li et al. [2] to source vectors with more than one source. Figure 1 demonstrates the typical behaviour of the algorithm R. As can be seen in the figure, this algorithm returns an estimate $\hat{f}$ that is indeed close to f_0 in the EMD but not necessarily close in more traditional norms like the ℓ_1 and ℓ_2 norms.

We first provide an upper bound for the performance of our recovery algorithm R on a noisy version $\tilde{y} = y + N(0, \sigma^2)$. Table I describes the asymptotics of our theoretical upper bound if all variables except the variable denoted in the left column are held constant. Table II then shows the asymptotics of our theoretical upper bound on EMD error of the private recovery algorithm. Figure 2 shows the results of simulations of our locally differentially private recovery algorithm. These results indicate that the locally private data can indeed be used to recover an estimate, $\hat{f}$, that is close to the true source locations, f_0, in the EMD.

We also provide a lower bound for noisy recovery of f_0:

$$\inf_{\hat{f}} \sup_{f_0} \mathbb{E}[\textbf{EMD}(f_0, \hat{f})] = \Omega\left(\min\left\{\frac{1}{2}, \frac{T^{1.5}\sigma}{\sqrt{m}}\right\}\right).$$

where $\inf_{\hat{f}}$ is the infimum over all estimators $\hat{f} : \mathbb{R}^m \to [0,1]^n$, $\sup_{f_0}$ is the supremum over all source vectors in $[0,1]^n$ and $\tilde{y}$ is sampled from $y + N(0, \sigma^2 I_m)$. Note that this lower bound matches our upper bound asymptotically in σ and is slightly loose in T. It varies by a factor of $\sqrt{m}$ from our theoretical upper bound but matches our experimental results in m (experimental results for noisy recovery can be found in the full version of this paper).

TABLE II: Asymptotic upper bounds for private recovery

VARIABLE	$\text{EMD}\left(\frac{f_0}{\|f_0\|_1}, \frac{\hat{f}}{\|\hat{f}\|_1}\right)$
n	$O(1 + \frac{1}{\sqrt{n}})$
m	$O(\sqrt{m})$
t	$O(\frac{1}{t} + \sqrt{t} + t^2 e^{\frac{-B^2}{4\mu t}})$

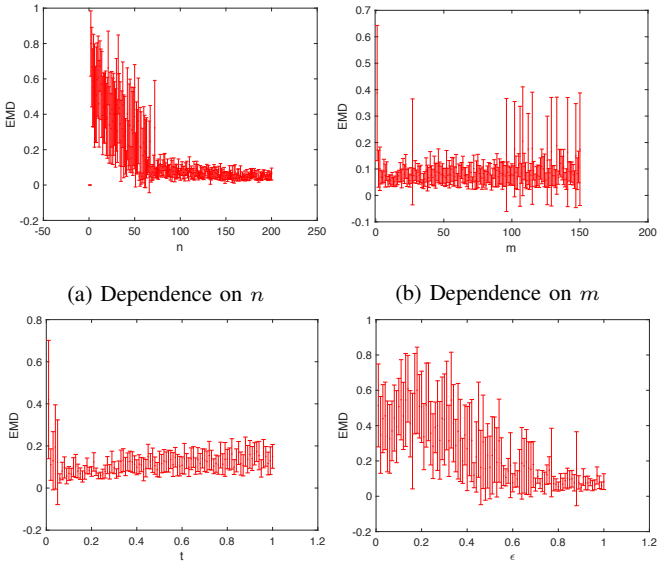

(a) Dependence on n (b) Dependence on m

(c) Dependence on t (d) Dependence on ϵ

Fig. 2: Empirical results for the EMD error of experiments running the recovery algorithm R on the locally differentially private thermal measurements. Unless specified otherwise, $n = 100$, $m = 50$, $t = 0.1$, $\delta = 0.1$ and $\epsilon = 1$.

Finally, we explore the relationship between the condition number of a matrix, M, and the amount of noise we need to add to the measurements to maintain privacy. That is, if we have a vector $x \in \mathbb{R}^n$ and measurement vector $y = Mx$, how much noise do we need to add to y to keep x private? We find that the amount of noise is greater than $1/(\sqrt{n}\kappa_2(M))$, where $\kappa_2(M)$ is the condition number. That is, if a problem is well-conditioned then we necessarily need to add a significant amount of noise to maintain privacy. The converse however is not generally true. It is possible to have a ill-conditioned matrix such that we still need to add a considerable amount of noise to maintain privacy.

REFERENCES

[1] C. DWORK AND A. ROTH, *The algorithmic foundations of differential privacy*, Found. Trends Theor. Comput. Sci., 9 (2014), pp. 211–407.

[2] Y. LI, S. OSHER, AND R. TSAI, *Heat source identification based on L1 constrained minimization*, Inverse Problems and Imaging, 1 (2014).