Adversarial Robustness of Streaming Algorithms through Importance Sampling

Vladimir Braverman (Johns Hopkins University, Google)
Avinatan Hassidim (Google)
Yossi Matias (Google)
Mariano Schain (Google)
Sandeep Silwal (MIT)
Samson Zhou (Carnegie Mellon University)

AdvML @ ICML 2021
Workshop at ICML 2021
July 24, 2021

Model

❖ Input: Elements of an underlying data set S, which arrives sequentially and *adversarially.* Adversary can choose future inputs after seeing previous outputs by honest algorithm
❖ Output: Evaluation (or approximation) of a given function
❖ Goal: Use space *sublinear* in the size m of the input S

❖ Surprising separation between "classic" streaming model where the stream input is fixed but the order of the updates may be given adversarially
❖ Hardt and Woodruff [HW13] showed that Linear sketches are NOT robust to adversarial attacks, must use $\Omega(n)$ space by giving an attack on AMS F_2 algorithm

Applications / Motivations

❖ Adversarial machine learning: ML problems where the input is chosen by an adversary
❖ Database queries: For multiple queries to a database, each query may depend on the responses to the previous queries
❖ Transparency of Algorithms: Internal state of honest algorithms may be entirely revealed or otherwise compromised

Coresets

❖ Coreset: Returns an ϵ-approximation on a query space

❖ Merge and reduce framework: Each $C_{1,j}$ is an $\frac{\epsilon}{\log n}$ coreset of the corresponding partition of the substream

❖ Applications: k-means clustering, k-median clustering, projective clustering, principal component analysis, Bayesian logistic regression, generative adversarial networks, k-line center, M-estimators

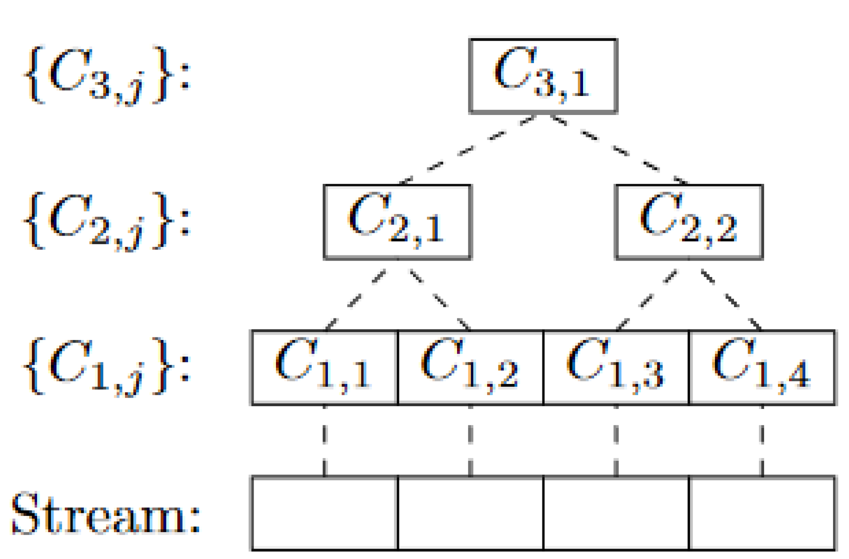

Row Sampling Algorithms for Linear Algebra

❖ Row-arrival model: $M_1, \ldots, M_n \in R^d$ rows of a matrix M
❖ Sample each row based on its "importance" to obtain $(1+\epsilon)$-approximate solutions to each problem

❖ Linear Regression: Output $x \in R^d$ to minimize $\|Mx-b\|_2$
❖ Spectral Sparsification / Subspace Embedding: Output $A \in R^{m\times d}$ so that $\|Mx\|_2 \approx \|Ax\|_2$ for all $x \in R^d$ and $m \ll n$
❖ Low-Rank Approximation: Output $A \in R^{m\times d}$ so that $\|M - MP\|_F \approx \|A - AP\|_F$ for all rank k projection matrices P
❖ L_1 Subspace Embedding: Output $A \in R^{m\times d}$ so that $\|Mx\|_1 \approx \|Ax\|_1$ for all $x \in R^d$ and $m \ll n$

Edge Sampling Algorithms for Graphs

❖ Edge-arrival model: $e_1, \ldots, e_m$ edges of a graph G
❖ Sample each edge based on its "importance" to obtain $(1+\epsilon)$-approximate solutions to each problem

❖ Graph Sparsification: Output H so that $Cut_H(S, V\setminus S) \approx Cut_G(S, V\setminus S)$ for any $S \subset V$

Results and Related Work

❖ Our result: Importance sampling based algorithms are adversarially robust!
❖ Intuition: Importance is a robust metric and sampling based algorithms use public randomness that is independent of previous randomness

❖ Corollary: Merge-and-reduce is adversarially robust
❖ Corollary: Row sampling algorithms are adversarially robust
❖ Corollary: Edge sampling algorithm is adversarially robust

Empirical Evaluations

❖ Streaming k-means clustering: a series of point batches where all points except the last batch are randomly sampled from a two-dimensional standard normal distribution. Points in the last batch sampled but around a distant center

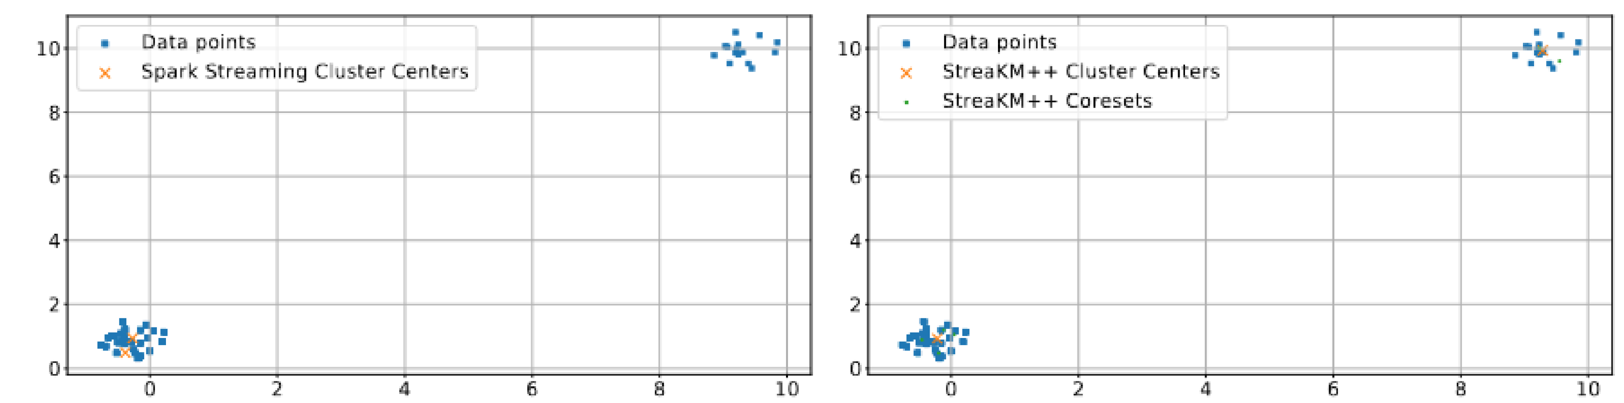

❖ Streaming linear regression: all batches except the last one are sampled around a constellation of four points in the plane such that the optimal regression line is of -1 slope through the origin. The last batch is at (L, L), far from the origin so the resulting optimal regression line has slope 1 through the origin

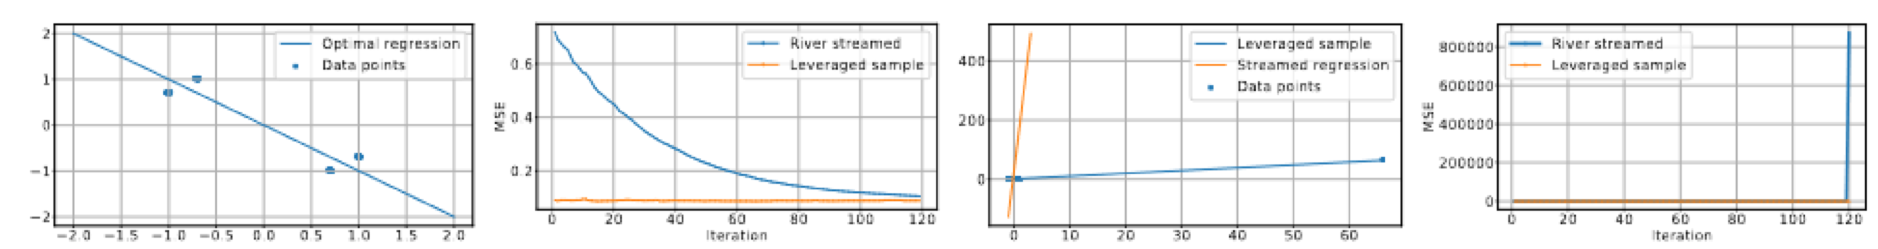

❖ Sampling vs. sketching: For a random unit sketching matrix S (each of its elements is sampled from $\{-1,1\}$ with equal probability), we create an adversarial data stream M such that its columns are in the nullspace of S for linear regression

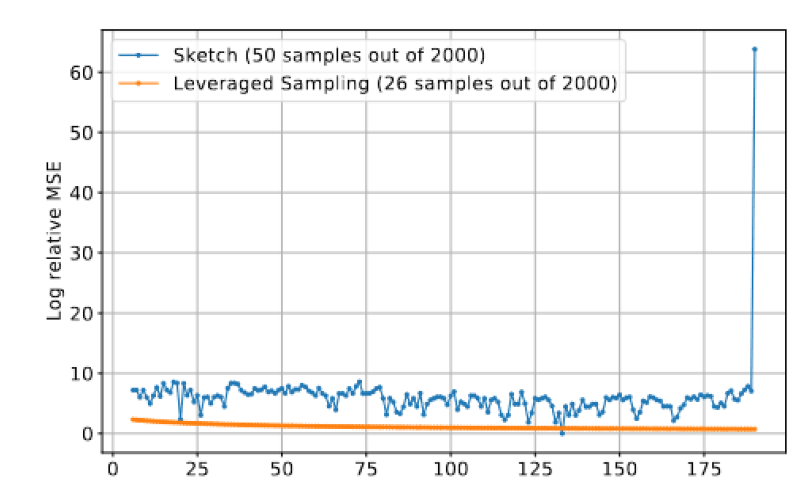

References

❖ [HW13] Moritz Hardt, David P. Woodruff: How robust are linear sketches to adaptive inputs? STOC 2013: 121-130
❖ [KMNS21] Haim Kaplan, Yishay Mansour, Kobbi Nissim, Uri Stemmer: Separating Adaptive Streaming from Oblivious Streaming. CRYPTO 2021 (to appear)