PopArt: Efficient Sparse Regression and Experimental Design for Optimal Sparse Linear Bandits

Chicheng Zhang

University of Arizona

Joint work with Kyoungseok Jang (NYU) and Kwang-Sung Jun (UArizona)

Problem: sparse linear bandits

- For time step $t = 1,2,..T$:
  - Take an action a_t from action space $\mathcal{A} \subseteq [-1,+1]^d$
  - Receive reward $r_t = \langle \theta_*, a_t \rangle + \eta_t$

    zero mean, 1-subgaussian noise

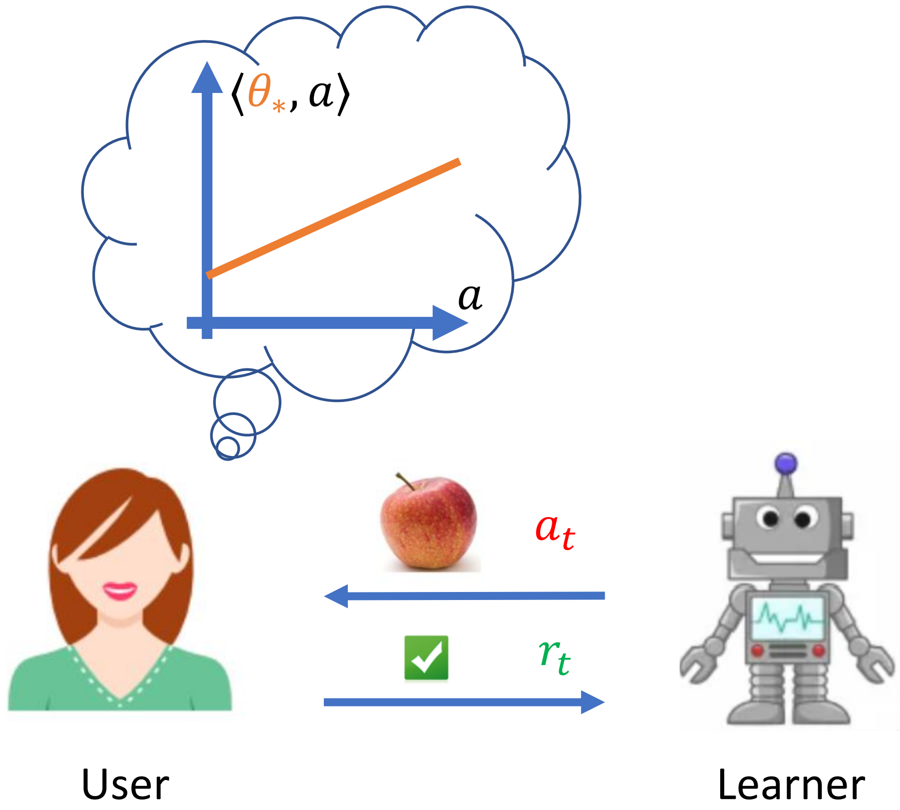

- Goal: minimize cumulative regret $\sum_{t=1}^{T} \max_{a \in \mathcal{A}} \langle \theta_*, a \rangle - \langle \theta_*, a_t \rangle$
- Tradeoff between exploration and exploitation
- θ_* is *sparse*

Problem: experimental design for sparse recovery

- For time step $t = 1,2, .. T$:
  - Make a measurement a_t from $\mathcal{A} \subseteq [-1, +1]^d$
  - Receive noisy response $r_t = \langle \theta_*, a_t \rangle + \eta_t$

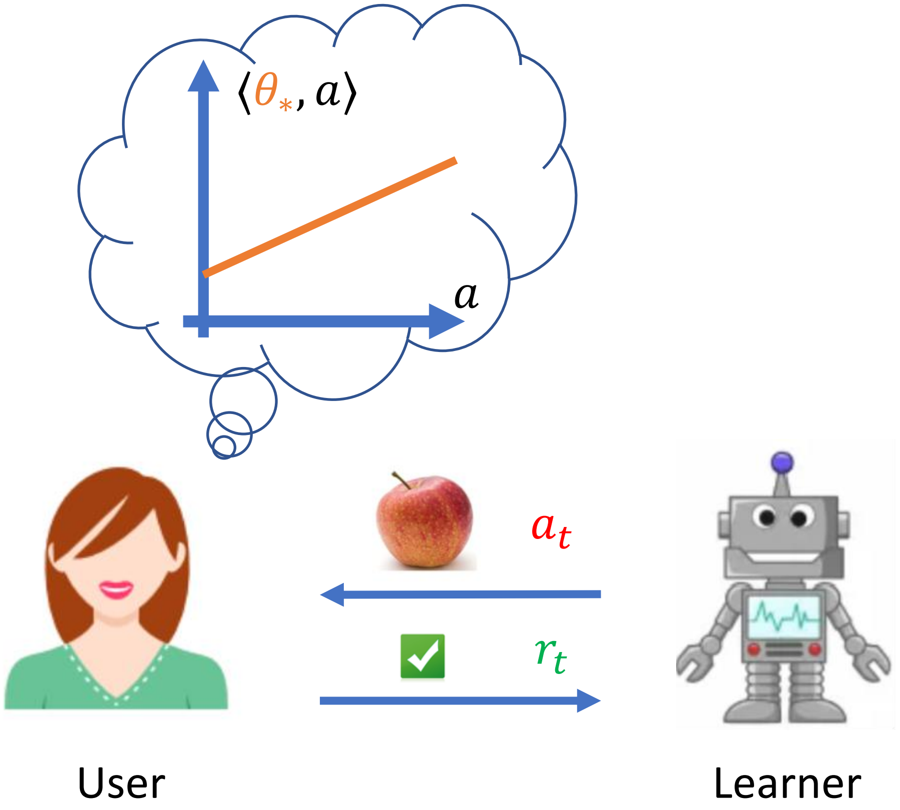

- Goal: output $\hat{\theta}$, such that $\left\| \hat{\theta} - \theta_* \right\|_1$ is small
- Assumption: θ_* is *sparse*
- No exploration / exploitation tradeoff

Main result: experimental design for sparse recovery

- Assumption: θ_* is s-sparse ($s \ll d$)
- ℓ_1-recovery error guarantee:

- $C_* = \max_{\pi \in \Delta(\mathcal{A})} \lambda_{\min}(Q(\pi)), \quad Q(\pi) = \mathbb{E}_{a\sim\pi}[aa^\top]$
- N_* is a new measurement set-dependent quantity
  - $N_* \geq C_*$ and can be $\gg C_*$

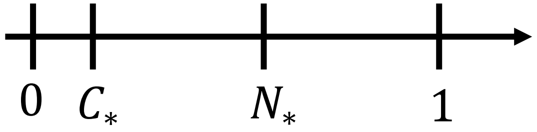

Main result: sparse linear bandits in data-poor regime

- Assumption: θ_* is s-sparse ($s \ll d$) and T is small ($T \ll d$)

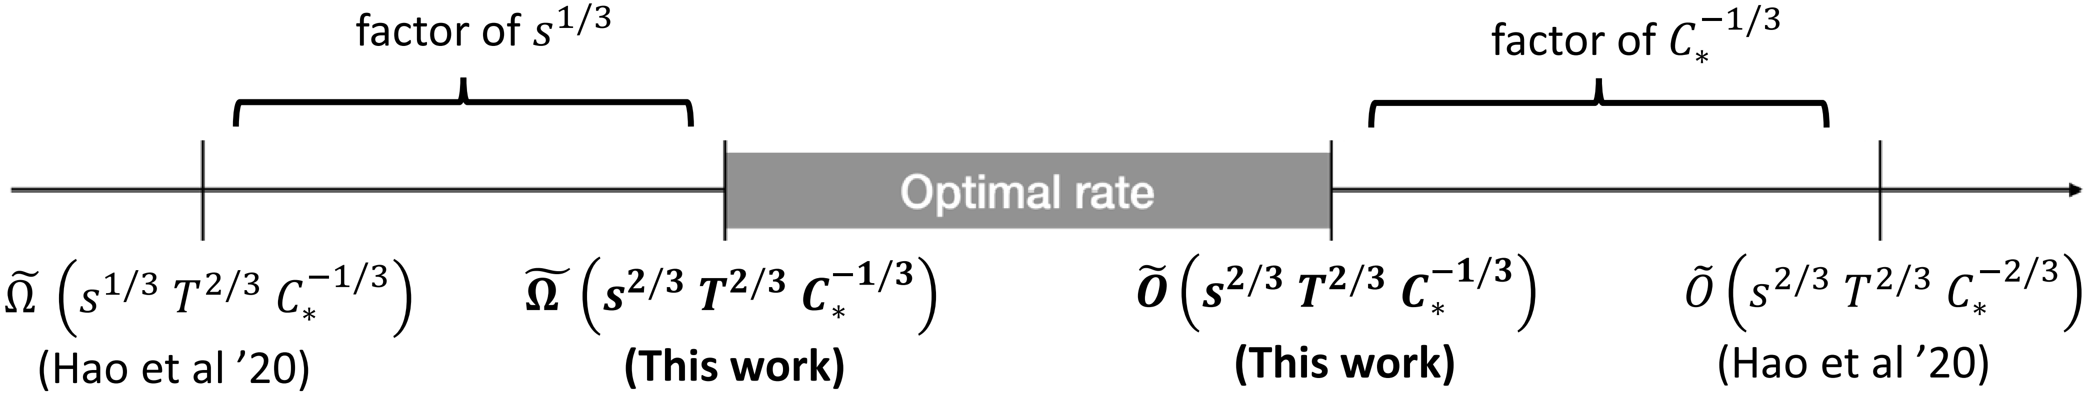

- Additional result:
  - $\widetilde{\Theta}\left(s^{2/3}\, T^{2/3} N_*^{-1/3}\right)$ regret upper and lower bounds

PopArt: new algorithm for sparse estimation

- Sparse estimation
  - Given: iid (a_t, r_t)'s such that $r_t = \langle \theta_*, a_t \rangle + \eta_t$, θ_* is s-sparse
  - Goal: recover θ_*

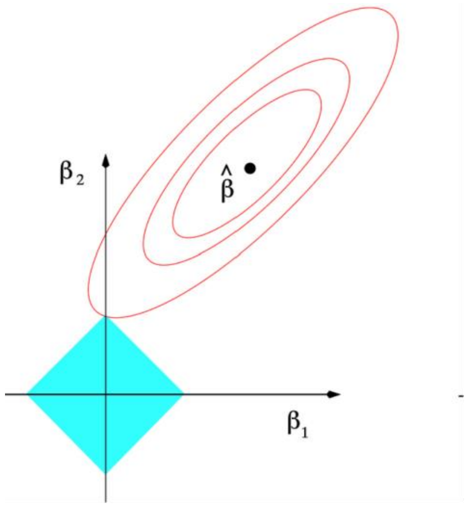

- Predominant approach: Lasso (Tibshirani'96)

$$\min_{\theta} \sum_t (r_t - \langle \theta, a_t \rangle)^2 + \lambda \|\theta\|_1$$

- Our approach: POPulation covariance regression with hARd Thresholding (PopArt)
  - Key insight: in many applications, the population covariance matrix, $Q = \mathbb{E}[aa^\top]$ is known and can be utilized

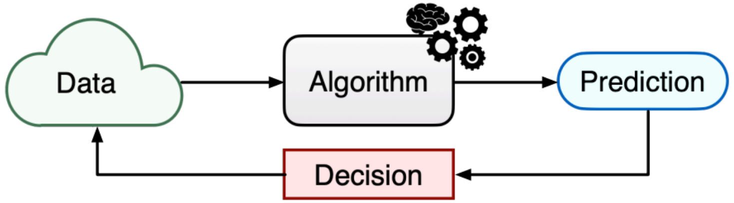

PopArt: the algorithm

- Key insight (Dani, Hayes, Kakade '07):

  $\tilde{\theta}_t = Q^{-1}(a_t r_t)$ is an unbiased estimator of θ_*

PopArt

Compute $\tilde{\theta}_t, t = 1, \ldots, T$

For $i = 1,2,..,d$:

  $\tilde{\theta}^i \leftarrow$ Robust-mean-estimator($\tilde{\theta}_1^i, \ldots, \tilde{\theta}_T^i$) (Catoni '12)

Return $H_s(\tilde{\theta})$.

- H_s: hard thresholding operator
  - Brings $\tilde{\theta}$ closer to θ^* in ℓ_1 distance, by exploiting sparsity of θ^*

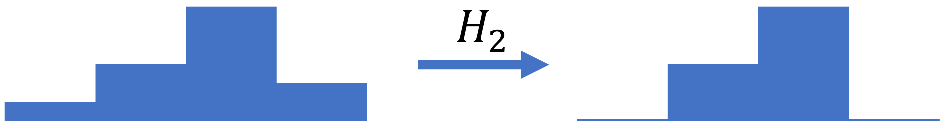

PopArt: recovery guarantee

- **Theorem:** (some variant of) PopArt outputs $\hat{\theta}$ that has ℓ_1 recovery error

$$\left\|\hat{\theta} - \theta_*\right\|_1 \leq \tilde{O}\left(s\sqrt{\max_i (Q^{-1})_{ii}}\sqrt{\frac{\ln d}{T}}\right)$$

  with high probability.

Experimental design for PopArt

$$\|\hat{\theta} - \theta_*\|_1 \leq \tilde{O}\left(s\sqrt{\max_i (Q^{-1})_{ii}}\sqrt{\frac{\ln d}{T}}\right)$$

- Given measurement set $\mathcal{A}$, minimize recovery bound by design sampling distribution

$$\mu^* = \operatorname*{argmin}_{\pi \in \Delta(\mathcal{A})} \max_i (Q(\pi)^{-1})_{ii}$$

- New experimental design criterion -> computationally tractable
- **Corollary:** PopArt with T samples from has ℓ_1 recovery error $\tilde{O}\left(\frac{s}{\sqrt{N_*}}\sqrt{\frac{\ln d}{T}}\right)$, where

$$N_* = \max_{\pi \in \Delta(\mathcal{A})} \min_i \frac{1}{(Q(\pi)^{-1})_{ii}}$$

Sparse linear bandits: algorithm & lower bounds

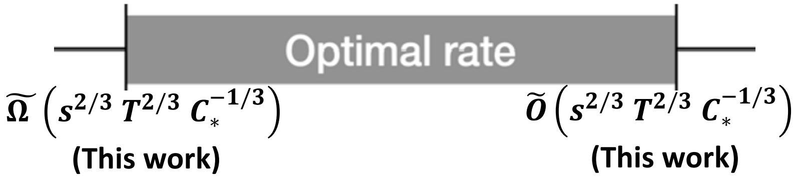

- New algorithm & regret upper bound: explore-then-commit
  - Use PopArt with μ^* for exploration
- New regret lower bound:
  - Main idea: reducing to lower bound sample complexity of estimating $\mathrm{support}(\theta^*)$
  - Prior work (Hao et al, 2020): use binary hypothesis testing
  - This work: using *symmetrization* to improve lower bound by a $s^{1/3}$ factor

Numerical simulations

- Sparse recovery algorithms evaluated:
  - PopArt with μ^* experimental design
  - Lasso with μ^* experimental design
  - Lasso with E-optimal design

- Experiment 1: $d = 10, s = 2$
  an $\mathcal{A}$ with $N_* \gg C_*$

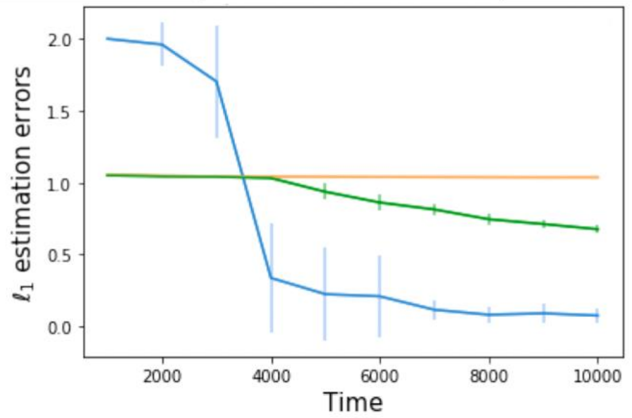

- Experiment 2: $d = 30, s = 2$
  $\mathcal{A}$ consists of 90 unit vectors ($N_* \approx C_*$)

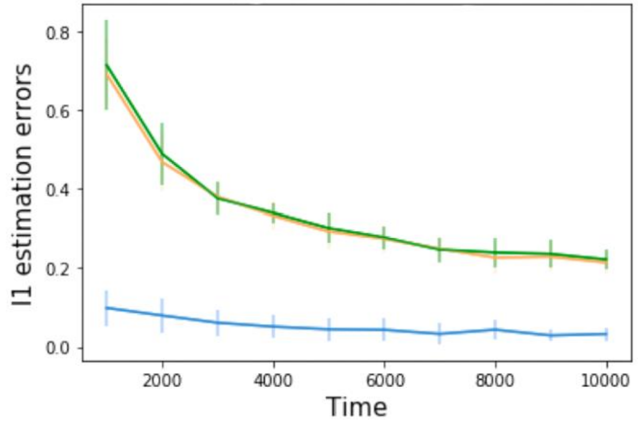

Conclusions & future work

- We propose PopArt, a new sparse estimation algorithm that can have better recovery guarantees than classical Lasso
- PopArt motivates a new experimental design criterion, and yields an optimal sparse linear bandit algorithm (in data-poor regime)

- Ongoing work: new estimator & experimental design for low-rank matrix regression
- Open questions:
  - (experimental design) Fixed measurement set $\mathcal{A}$ and budget T, can we design estimators with even lower recovery error?
  - (bandits) what's the minimax regret for sparse linear bandits in all regimes of parameters?

- Thank you! Paper: https://arxiv.org/abs/2210.15345

Backup

Main result: experimental design for sparse recovery

- Assumption: θ_* is s-sparse
- ℓ_1-recovery error guarantee:

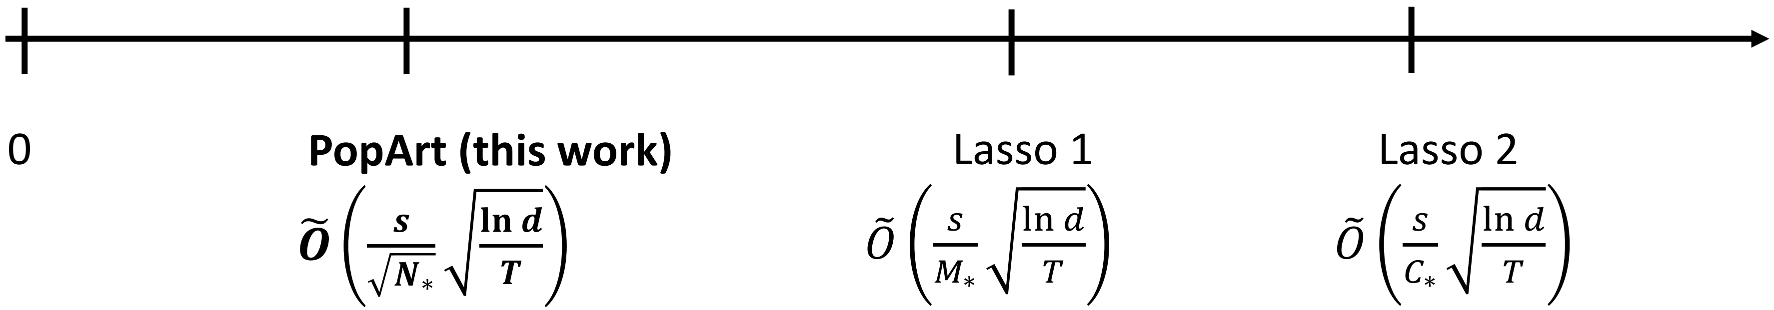

- $M_* = \max_{\pi \in \Delta(\mathcal{A})} \phi_0^2(Q(\pi), s)$ – restricted eigenvalue
  - Intractable Experimental design criterion (Bandeira et al, 2012)
  - Proposition: exists $\mathcal{A}$, such that $\frac{1}{\sqrt{N_*}} \ll \frac{1}{M_*}$

Definition of N_*

- $N_* = \max_{\pi \in \Delta(\mathcal{A})} \min_i \frac{1}{(Q(\pi)^{-1})_{ii}}$

- $\max_{\pi \in \Delta(\mathcal{A})} \min_i \frac{1}{(Q(\pi)^{-1})_{ii}}$

PopArt: recovery guarantee

- **Theorem:** (some variant of) PopArt outputs $\hat{\theta}$ that has ℓ_1 recovery error

$$\left\|\hat{\theta} - \theta_*\right\|_1 \leq \tilde{O}\left(s \sqrt{\max_i (Q^{-1})_{ii}} \sqrt{\frac{\ln d}{T}}\right)$$

with high probability.

- Key idea of analysis:
  - establish bounds on $\left|\tilde{\theta}^i - \theta_*^i\right|$ for all coordinate i
  - To this end, analyze variance of $\tilde{\theta}_t^i$

$\tilde{\theta}^i \leftarrow$ Robust-mean-estimator($\tilde{\theta}_1^i, \ldots, \tilde{\theta}_T^i$)

Sparse linear bandits: algorithm & lower bounds

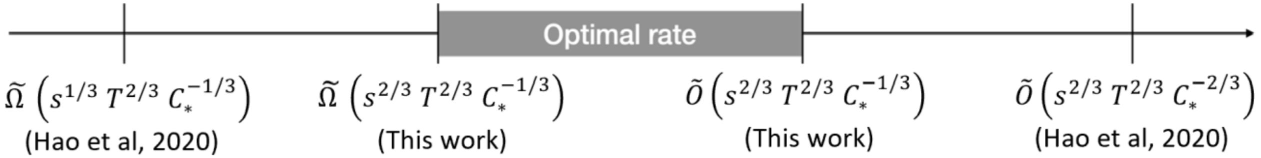

- New algorithm & regret upper bound: explore-then-commit

First T_1 rounds:
  take $a_t \sim \pi^*$, see r_t
use PopArt to compute $\hat{\theta}$

Remaining $T - T_1$ rounds:
  take $a_t = \mathrm{argmax}_{a \in \mathcal{A}} \langle \hat{\theta}, a \rangle$

- New regret lower bound:
  - Main idea: reducing to lower bound sample complexity of estimating $\mathrm{support}(\theta^*)$
  - Prior work (Hao et al, 2020): use binary hypothesis testing
  - This work: using *symmetrization* to improve lower bound by s factor

Sparse linear bandits: algorithm & lower bounds

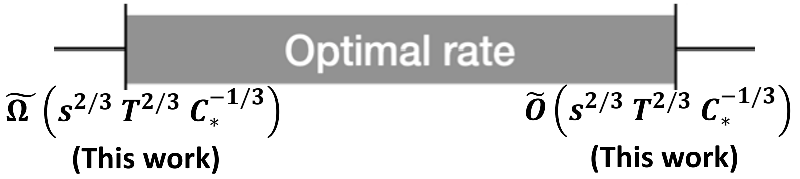

- New algorithm & regret upper bound: explore-then-commit

First T_1 rounds:
take $a_t \sim \mu^*$, see r_t
use PopArt to compute $\hat{\theta}$

Remaining $T - T_1$ rounds:
take $a_t = \mathrm{argmax}_{a \in \mathcal{A}} \langle \hat{\theta}, a \rangle$

- New regret lower bound:
  - Main idea: reducing to lower bound sample complexity of estimating $\mathrm{support}(\theta^*)$
  - Prior work (Hao et al, 2020): use binary hypothesis testing
  - This work: using *symmetrization* to improve lower bound by a $s^{1/3}$ factor