Active Learning and Best-Response Dynamics

Carnegie Mellon University

Maria-Florina Balcan
Carnegie Mellon

Christopher Berlind
Georgia Tech

Avrim Blum
Carnegie Mellon

Emma Cohen
Georgia Tech

Kaushik Patnaik
Georgia Tech

Le Song
Georgia Tech

Georgia Tech

Setting

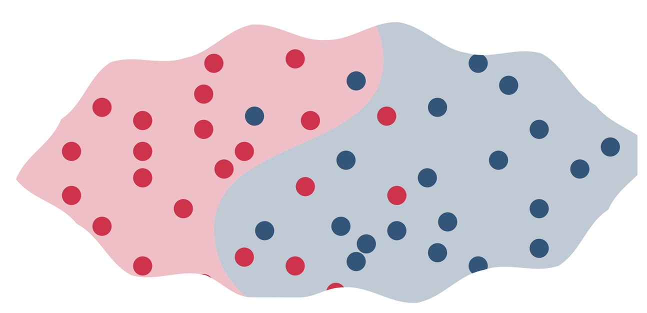

Many low-power distributed sensors

- Can only communicate locally
- Sensor readings are noisy

Single central agent

- Can make costly queries of any sensor
- Goal is to detect a spacial boundary

Example: Sensing chemical concentrations

Our Approach

Denoising

- Sensors play consensus game
- Payoff: correlation with neighbors
- Best response: update to neighbor majority
- Synchronous or asynchronous updates

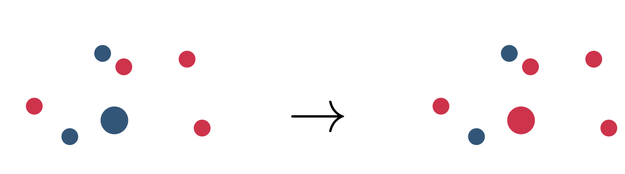

Active Learning

- Actively select intelligent sensor queries
- AL provably only effective in low noise
- Use noise-tolerant margin-based active learning algorithm

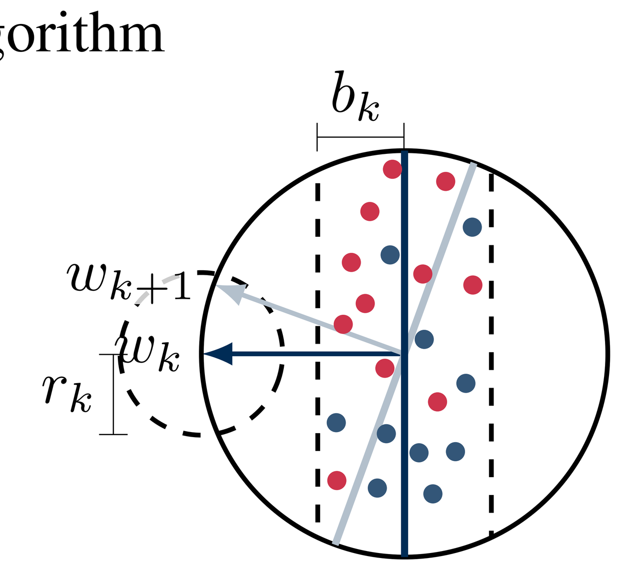

Contributions

Theory

- Positive results for denoising effectiveness, including synchronous and random updates
- Arbitrary asynchronous updates can fail

Experiments

- Denoising is very effective
- Denoising improves performance of AL

Notation and Setup

- N sensors with communication radius r
- Noisy sensor reading with probability η
- Sensors uniformly dist. in unit sphere in $\mathbb{R}^d$
- Boundary is homogeneous linear separator

Synchronous Updates

Theorem 1. *If*

$$N \geq \frac{2}{(r/2)^d(1/2-\eta)^2} \ln\left(\frac{2}{(r/2)^d(1/2-\eta)^2\delta}\right)$$

then, w.p. $\geq 1-\delta$, after one update every sensor at a distance $\geq r$ from the separator is correct.

Proof sketch:

- Consider a sensor x far from separator
- $\mathbb{E}[\text{nbrs}] = N \cdot \mu(B_r(x)) \geq \frac{1}{(1/2-\eta)^2}\ln(N/\delta)$
- $\mathbb{E}[\text{bad nbrs}]$ is η fraction
- Apply Bernstein and union bound

Random Order Asynchronous

Theorem 2. *If $r \leq O(\frac{1/2-\eta}{\sqrt{d}})$ and*

$$N \geq \frac{1}{(r/2)^d(1/2-\eta)^2} \ln\left(\frac{1}{r^d(1/2-\eta)^2\delta}\right)$$

then, w.p. $\geq 1-\delta$, all sensors at distance $\geq 2r$ from the separator will update correctly.

Proof sketch:

- Partition sensors based on distance from sep.
- Inside sensors have no guarantee
- Mid-distance sensors rarely mistaken
- Outside sensors never update incorrectly
- Apply Hoeffding bounds and Theorem 1

Arbitrary Order Asynchronous

Theorem 3. *For some c, if $\phi = \min(\eta, \frac{1}{2}-\eta)$,*

$$N \geq \frac{16}{(cr)^d\phi^2}\left(\ln\frac{8}{(cr)^d\phi^2} + \ln\frac{1}{\delta}\right),$$

and $r \leq 1/2$ then, w.p. $\geq 1-\delta$, there exists an ordering so that asynchronous updates in this order cause all points to have the same label.

Proof sketch:

- Wave of updates from left to right
- First half correctly turns negative (Thm. 1)
- Second half incorrectly turns negative

Denoising Results

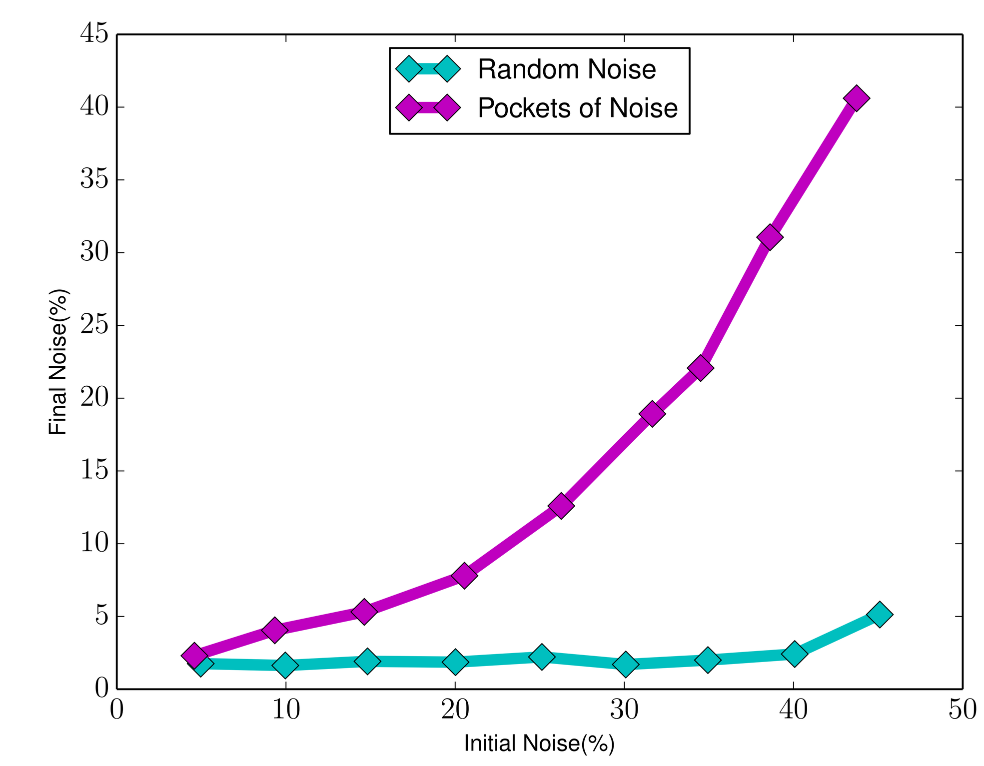

- Comparison of initial vs. final noise rates
- $N = 10,000$ and $r = 0.1$
- Synchronous (shown) and asynchronous perform comparably

Active Learning Results

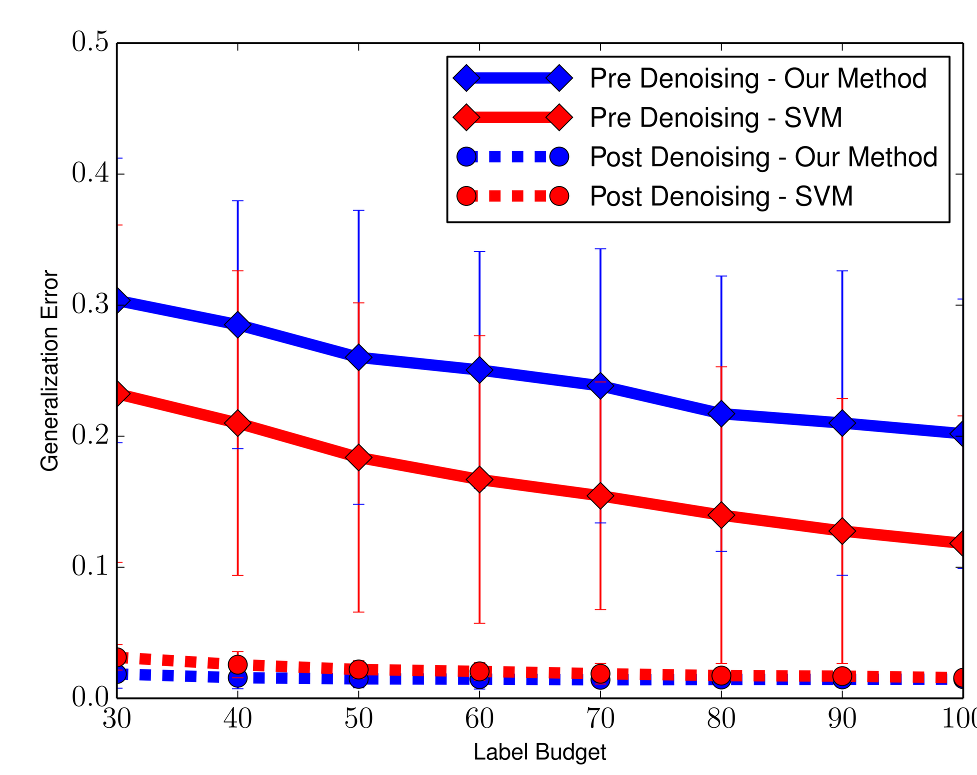

- Compared against passive SVM
- $N = 10,000$, $r = 0.1$, and $\eta = 0.35$
- Active outperforms passive after denoising

Discussion

- We seek specific ϵ-equilibrium configuration instead of complete consensus equilibrium
- Conservative best response: only update if confident on correct side of separator
- Combining denoising and kernelized AL algorithm can help detect nonlinear boundaries

Future Directions

- Guarantees for different boundary shapes
- Changing environments (moving boundary)
- Non-uniform sensor distributions
- Alternative denoising dynamics