Can reading half a bit help speed up your decision tree?

Renan Gross
Weizmann Institute

The problem

Given $f:\{-1,1\}^n \to \mathbb{R}$ and a uniform input $x \in \{-1,1\}^n$, compute $f(x)$ while querying as few bits as possible.

Decision trees: OUT! Fractional query algorithms: IN

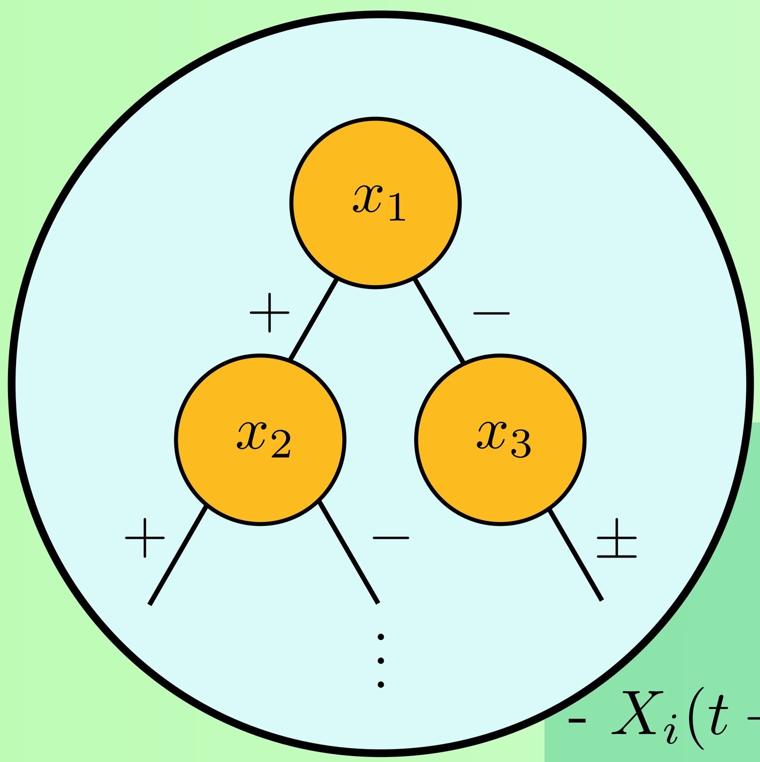

- Choose index i
- $X_i(t+1) = \begin{cases} +1 & \text{w.p. } \frac{1}{2} \\ -1 & \text{w.p. } \frac{1}{2} \end{cases}$
- Pay cost of 1

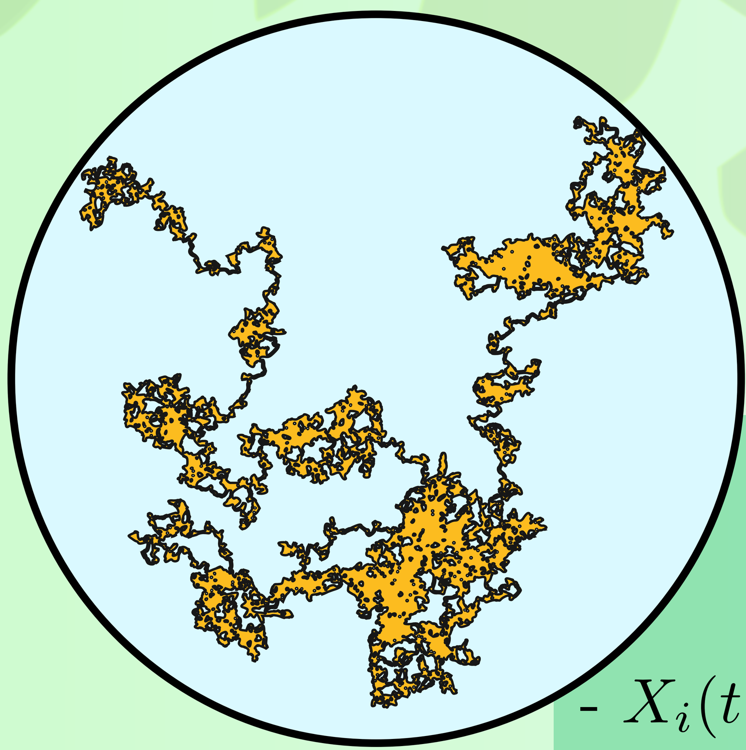

- Choose index i
- $X_i(t+1) = \begin{cases} X_i(t)+\varepsilon & \text{w.p. } \frac{1}{2} \\ X_i(t)-\varepsilon & \text{w.p. } \frac{1}{2} \end{cases}$
- Pay cost of ε^2

✔ $X(\infty)$ is uniform on $\{-1,1\}^n$.

✔ The process $X(t)$ is a martingale, and so is $f(X(t)) = \sum_{S\subseteq[n]} \hat{f}(S)\prod_{i\in S} X_i(t)$.

The Schramm-Steif theorem

For a fractional query algorithm $X(t)$, let τ be the time that $f(X(\infty))$'s value is known, and define $\delta := \max_{i\in[n]} \mathbb{E}\left[X_i(\tau)^2\right]$.

Theorem: For every $k \geq 1$, we have

$$\sum_{|S|=k} \hat{f}(S)^2 \leq \delta k \|f\|_2^2.$$

Noise sensitivity

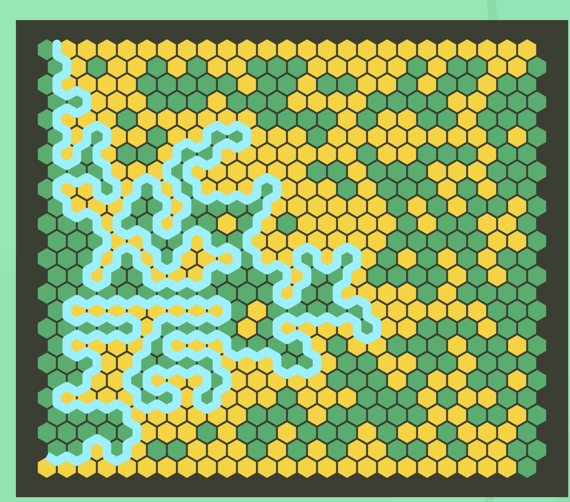

ε-noising

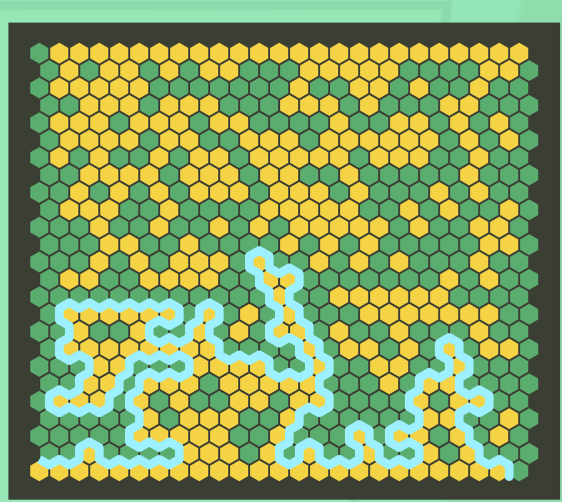

Lower bounds from partial differential equations

Let u_ε be the optimal $\sum_i \mathbb{E}\left[X_i(\tau)^2\right]$ when $X(0)=x$. Then

$$u_\varepsilon(x) = \min_i \frac{u_\varepsilon(x+\varepsilon e_i) + u_\varepsilon(x-\varepsilon e_i)}{2} + \varepsilon^2,$$

and when $\varepsilon \to 0$, we have

$$\min_i \frac{\partial^2 u}{\partial x_i^2} = -2.$$

Recursively solving the Dirichlet problem on the boundary gives complexity lower bounds.

Advantage???

Is there asymptotic advantage?

Example: n-bit OR

Always update the largest bit.
This reveals 1 bit as $n \to \infty$,
but decision trees need 2.