Recap Active learning

"Uncertainty"

Regression $\quad x_t \leftarrow \arg\max_x \sigma_{t-1}(x)$

↑ var
- closed form (linear generalized case)
- ensembles
- posterior var (Bayesian)

Classification $\quad f \in \mathcal{F}$

f^* - ground truth $\quad \bar{f}$ - best model in $\mathcal{F}$ $\quad \hat{f}$ - model in $\mathcal{F}$ trained using data

$$R(\hat{f}) - R(f^*) = \underbrace{R(\hat{f}) - R(\bar{f})}_{\text{est err}} + \underbrace{R(\bar{f}) - R(f^*)}_{\text{approx err}}$$

est err ≡ variance $\quad$ approx err ≡ bias^2

$$R(f) - R(g) = \int_{\{x:\, f(x) \neq g(x)\}} \left| P(Y=1|X) - \tfrac{1}{2} \right| dx$$

"Uncertainty" $\quad$ Probabilistic classifier $P(Y|X)$ eg. logistic regression

Binary class $\quad \arg\min_x \left| P(Y=\hat{y}|x) - \frac{1}{2} \right|$

Multi class
- $\arg\max_x \; 1 - P(Y=\hat{y}|X)$ $\quad$ least conf. ✓ ($\hat{y}$ ↑ most likely label)
- $\sum_y P(Y=y|X) \log \frac{1}{P(Y=y|x)}$ $\quad$ entropy
- $P(Y=1|X) - P(Y=2|X)$ $\quad$ margin ✓ (most likely, next most likely)

→ What @ non-probabilistic classifiers? linear, SVM, decision trees

Binary classes

Margin based active learning for linear classifiers $\quad$ 1-dim linear

binary bisection

(d-dim noisy)

$w_0 \sim$ unif $\quad \|w_0\| = 1 \quad \beta_0 = \pi$

for $e = 1 .. \log T$

Collect n samples $x_i \sim P_X$ (unif dist^n over unit ball)

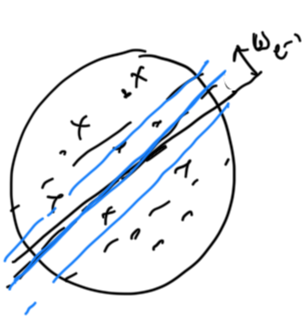

If $|\langle w_{e-1} \cdot x_i \rangle| > th = \frac{\beta_{e-1}}{\sqrt{d}} \sqrt{\log T}$

then collect y_i

Find $w_e \in \{w : \theta(w_{e-1}, w) \le \beta_{e-1}, \|w_2\|=1\}$

that minimizes empirical error

$$\sum_{i=1}^{n} 1_{y_i w \cdot x_i < 0}$$

Update $\beta_e \leftarrow c \cdot \beta_{e-1}$, $e \leftarrow e+1$

end.

$\frac{1}{n}$ vs. 2^{-n}

random/passive — active

(no noise)

$\frac{1}{2}$ $\frac{1}{2^2}$ ... $\frac{1}{2^n}$

c c^2 ... c^n - exp.

Decision Trees — Margin

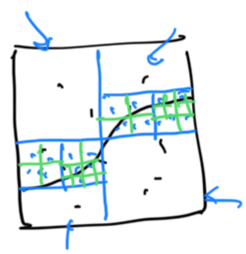

margin ≡ leaves @ the deepest level

		Passive	Active
Noiseless	1 dim	$\frac{1}{n}$	2^{-n}
	d-dim (linear)	$\frac{d}{n}$	$e^{-n/d}$
Noise	d dim	$\left(\frac{d}{n}\right)^{\frac{k}{2k-1}}$	$\left(\frac{d}{n}\right)^{\frac{k}{2k-2}}$

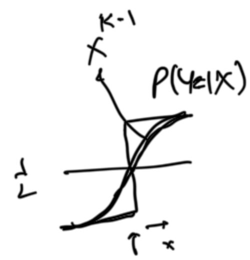

K-noise

→ Significant (exp) improvement is possible with active learning in noiseless setting

→ Gains decrease as there is more label noise.

Neural Networks

- $P(Y|X)$ - uncalibrated.

UnCalibrated $P(\hat{Y}=y \mid \hat{p}=p) \neq p$

- Margin

- Margin

- Ensemble methods (Query by committee, Disagreement based active learning)
  ↑ variance

Dropout

Coreset sampling $\left| \frac{1}{n}\sum_{i=1}^{n} l(x_i, y_i) - \frac{1}{k}\sum_{i=1}^{n} l(x_i, y_i) \right|$ ↙

↓ k-centers d

{ Representation sampling
{ Diversity sampling

Sequential Decision Making

Bandits Decision choices ≡ actions ≡ arms

reward

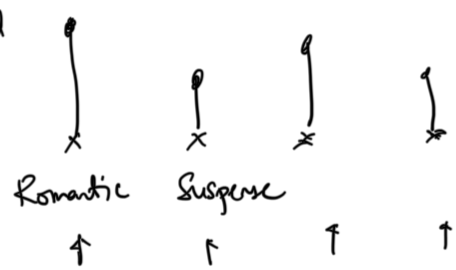

goal: max reward
not learn reward

Stochastic opt { Best arm identification $a^* = \arg\max_{a \in A} \mu(a)$ ← μ - mean reward.

{ Simple Regret $\min_{a_1..a_T} \mu(a^*) - \mu(a_T)$ ←

→ Cumulative Regret $\min_{a_1..a_T} \sum_{t=1}^{T} \mu(a^*) - \mu(a_t) = R(T)$ ← T, K ..

A - finite, finite action space $|A| = K$.

T - rounds

Non-adaptive exploration

$E[r(a)] = \mu(a)$

① Uniform exploration

$\frac{1}{N}\sum_{i=1}^{N} r_i(a) \quad E[r(a)] = \mu(a)$

1. Try each arm N times
2. Select highest empirical reward $\hat{\mu}(a)$ average arm $\hat{a}$
3. Play arm rest $T-NK$ times.

Rewards bounded $\mu(a) \in [0,1] \quad \forall a$

$$P(|\mu(a)-\hat{\mu}(a)| > \varepsilon) \le \underbrace{2e^{-2N\varepsilon^2}}_{\delta}$$

w.p. $\geq 1-\delta \quad |\mu(a)-\hat{\mu}(a)| \le \sqrt{\frac{1}{2N}\log\frac{2}{\delta}} = \varepsilon$ Confidence bound.

Regret $\mu(a^*)-\mu(\hat{a}) \le \hat{\mu}(a^*) + \varepsilon - \hat{\mu}(\hat{a}) + \varepsilon$

$\le 2\varepsilon \qquad \hat{\mu}(\hat{a}) \geq \hat{\mu}(a^*)$

wp $\geq 1-2\delta$

$$R(T) \le 1.NK + (T-NK)2\varepsilon$$

$$= NK + (T-NK)2\sqrt{\frac{1}{2N}\log\frac{2}{\delta}}$$

$$\le NK + T2\sqrt{\frac{1}{2N}\log\frac{2}{\delta}} \qquad \text{wp} \geq 1-K\delta$$

$N \approx \left(\frac{T}{K}\right)^{2/3}\left(2\log\frac{2}{\delta}\right)^{1/3}$

$$E[R(T)] = O\left(T^{2/3}\left(K\log\frac{1}{\delta}\right)^{1/3} + KT\delta\right)$$

$\delta = \frac{1}{T^2}$

$$= O\left(T^{2/3}(K\log T)^{1/3}\right)$$

↑ sublinear

② ε-greedy

Toss coin wp heads ε_t

If heads choose arm uniformly at random

o.w. choose empirically best arm.

$$E[R(t)] = O\left(t^{2/3} (K \log t)^{1/3}\right) \quad \text{for any } t.$$

$$\text{if } t\epsilon_t \approx NK$$

Adaptive exploration

subsequent arm and no of times it is pulled depends on previous rounds.

$n_t(a)$ - no of times a is pulled upto time t.

→ Successive Eliminate

→ Upper Confidence Bound