Contextual Active Online Model Selection with Expert Advice

Xuefeng Liu, Fangfang Xia, Rick L. Stevens, Yuxin Chen

THE UNIVERSITY OF CHICAGO

Argonne NATIONAL LABORATORY

Motivating Applications

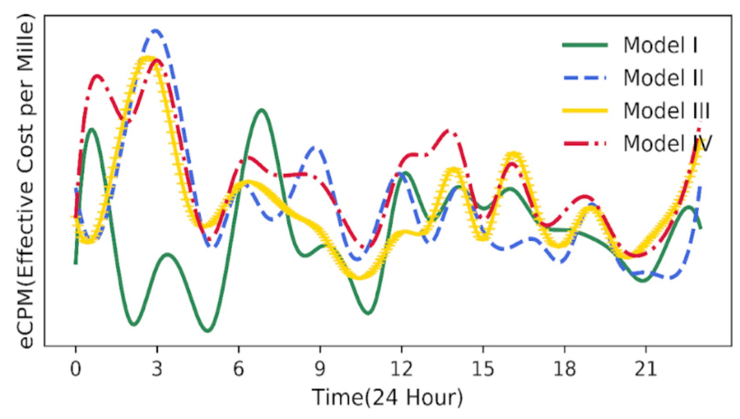

Recommend System:

Average highest performance model IV has lower performance than some other model on 80% of the time

healthcare
- drug development
- precision medicine

finance
- trading strategies
- option pricing

airline ancillary pricing
- optimal online pricing model

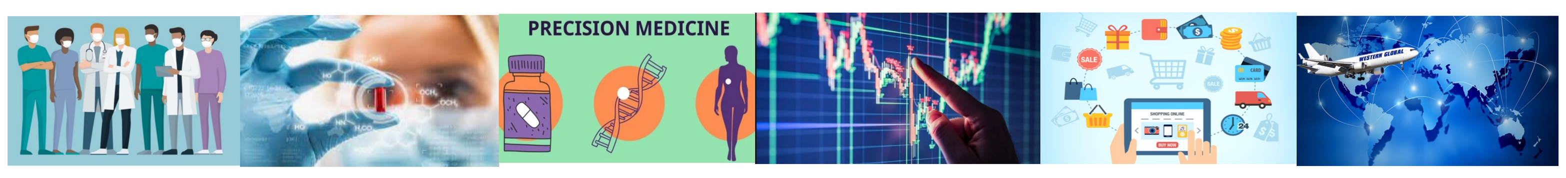

model performance depends on the context

cost-sensitive to evaluate and access the models / labels

online streaming data instead of a pool of data points

CAMS: Algorithmic Details

Input: Models $\mathcal{F}$, policies Π^*, #rounds T, budget b
Initialize loss $\tilde{\boldsymbol{L}}_0 \leftarrow 0$; query cost $C_0 \leftarrow 0$
for $t = 1, 2, ..., T$ **do**
  Receive $\boldsymbol{x}_t$
  $\eta_t \leftarrow \textsc{SetRate}(t, \boldsymbol{x}_t, |\Pi^*|)$
  Set $q_{t,i} \propto \exp\left(-\eta_t \tilde{L}_{t-1,i}\right) \forall i \in |\Pi^*|$
  $j_t \leftarrow \textsc{Recommend}(\boldsymbol{x}_t, \boldsymbol{q}_t)$
  Output $\hat{y}_{t,j_t} \sim f_{t,j_t}$ as the prediction for $\boldsymbol{x}_t$
  Set $\mathfrak{E}(\hat{\boldsymbol{y}}_t, \boldsymbol{w}_t) := \frac{1}{c} \sum_{y \in \mathcal{Y}, \bar{\ell}_t^y \in (0,1)} \bar{\ell}_t^y \log_c \frac{1}{\bar{\ell}_t^y}$ [a]
  Set query lower bound $\delta_0^t = \frac{1}{\sqrt{t}}$
  Compute $z_t = \max\{\delta_0^t, \mathfrak{E}(\hat{\boldsymbol{y}}_t, \boldsymbol{w}_t)\}$
  Sample $U_t \sim \text{Ber}(z_t)$
  if $U_t = 1$ and $C_t \leq b$ **then**
    Query the label y_t
    $C_t \leftarrow C_{t-1} + 1$
    Compute $\boldsymbol{\ell}_t$: $\ell_{t,j} = \mathbb{I}\{\hat{y}_{t,j} \neq y_t\}, \forall j \in [|\mathcal{F}|]$
    Estimate model loss: $\hat{\ell}_{t,j} = \frac{\ell_{t,j}}{z_t}, \forall j \in [|\mathcal{F}|]$
    Update $\tilde{\boldsymbol{\ell}}_t$: $\tilde{\ell}_{t,i} \leftarrow \langle \pi_i(\boldsymbol{x}_t), \hat{\ell}_{t,j} \rangle, \forall i \in [|\Pi^*|]$
    $\tilde{\boldsymbol{L}}_t = \tilde{\boldsymbol{L}}_{t-1} + \tilde{\boldsymbol{\ell}}_t$
  else
    $\tilde{\boldsymbol{L}}_t = \tilde{\boldsymbol{L}}_{t-1}$
    $C_t \leftarrow C_{t-1}$

Contextual Model Selection

Active Queries

Model Updates

procedure $\textsc{SetRate}(t, \boldsymbol{x}_t, m)$
  if STOCHASTIC **then**
    $\eta_t = \sqrt{\frac{\ln m}{t}}$
  if ADVERSARIAL **then**
    $\rho_t = 1 - \max_{\tau \in [t]} \langle \boldsymbol{w}_t, \mathbb{I}\{\hat{\boldsymbol{y}}_t = y\} \rangle$
    $\eta_t = \sqrt{\frac{1}{\sqrt{t}} + \frac{\rho_t}{c^2 \ln c}} \cdot \sqrt{\frac{\ln m}{T}}$
  return η_t

procedure $\textsc{Recommend}(\boldsymbol{x}_t, \boldsymbol{q}_t)$
  if STOCHASTIC **then**
    $\boldsymbol{w}_t = \sum_{i \in |\Pi^*|} q_{t,i} \pi_i(\boldsymbol{x}_t)$
    $j_t \leftarrow \text{maxind}(\boldsymbol{w}_t)$
  if ADVERSARIAL **then**
    $i_t \sim \boldsymbol{q}_t$
    $j_t \sim \pi_{i_t}(\boldsymbol{x}_t)$
  return j_t

[a] we denote by $\bar{\ell}_t^y := \langle \boldsymbol{w}_t, \mathbb{I}\{\hat{\boldsymbol{y}}_t \neq y\} \rangle$ as the expected loss if the true label is y, where $\boldsymbol{w}_t = \pi_{\text{maxind}(\boldsymbol{q}_t)}(\boldsymbol{x}_t)$ and $\text{maxind}(\boldsymbol{w}) := \arg\max_{j: w_j \in \boldsymbol{w}} w_j$.

Research Questions

- How to select data-adaptive models when facing heterogeneous data stream?
- How to make it labeling efficient?

We want a robust **cost-effective online-learning** algorithms that
- effectively identify best model selection policy
- works under limited labeling resources
- are adaptive to arbitrary data streams

Learning Protocol

Algorithm 1 CONTEXTUAL ACTIVE MODEL SELECTION PROTOCOL

1: Given a set of classifiers $\mathcal{F}$ and model selection policies Π
2: **for** $t = 1, 2, ..., T$ **do**
3:   The learner receives a data instance $\boldsymbol{x}_t \in \mathcal{X}$ as the context
4:   Compute the predicted label $\hat{y}_{t,j}$ for each pre-trained classifier $f_j(\boldsymbol{x}_t), j \in [k]$
5:   The learner identifies a model f_{j_t} and makes a prediction $\hat{y}_{t,j_t}$ for the instance $\boldsymbol{x}_t$
6:   **if** The learner decide to query **then**
7:     It incurs a QUERY COST and observes true label y_t
8:     It receives a (full) 0-1 loss vector $\boldsymbol{\ell}_t = \mathbb{I}_{\{\hat{\boldsymbol{y}}_t \neq y_t\}}$
9:     It can then use the queried labels to adjust its model selection criterion

Comparison against related work

Algorithm	Online bagging	Hedge	EXP3	EXP4	QBC	Model Picker	CAMS
model selection	no	yes	yes	yes	no	yes	yes
full-information	yes	yes	no	no	yes	yes	yes
active	no	no	no	no	yes	yes	yes
contextual	no	no	no	yes	no	no	yes

Theoretical Guarantees

pseudo-regret for stochastic setting

$$\overline{\mathcal{R}}_T(\mathcal{A}) = \mathbb{E}[L_T^{\mathcal{A}}] - T \min_{i \in [|\Pi^*|]} \mu_i$$

expected loss of policy *i* if recommending the most probable model $\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{x}_t, y_t}\left[\ell_{t, \text{maxind}(\pi_i(\boldsymbol{x}_t))}\right]$

expected regret for adversarial setting

$$\mathcal{R}_T(\mathcal{A}) = \mathbb{E}[L_T^{\mathcal{A}}] - \min_{i \in [|\Pi^*|]} \sum_{t=1}^{T} \tilde{\ell}_{t,i}$$

expected loss of policy *i* if randomizing the model recommendation at t $\tilde{\ell}_{t,i} := \langle \pi_i(\boldsymbol{x}_t), \boldsymbol{\ell}_t \rangle$

Query complexity and tight regret bound under

- Stochastic data streams
- Adversarial data streams
- Finite policy / model classes
- Infinite policy class with finite VC Dimension

Algorithm	Regret	Query Complexity
Exp3	$2\sqrt{Tk \log k}$	–
Exp3.p	$5.15\sqrt{nT \log \frac{n}{\delta}}$	–
Exp4	$\sqrt{2Tk \log n}$	–
Exp4.p	$6\sqrt{kT \ln \frac{n}{\delta}}$	–
Model Picker$_{\text{stochastic}}$	$62 \max_i \Delta_i k / (\lambda^2 \log k)$ $\lambda = \min_{j \in [k] \setminus \{i^*\}} \Delta_j^2 / \theta_j$	$\sqrt{2T \log k}(1 + 4\frac{c}{\Delta})$
Model Picker$_{\text{adversarial}}$	$2\sqrt{2T \log k}$	$5\sqrt{T \log k} + 2L_{T,*}$
CAMS$_{\text{stochastic}}$	$\left(\frac{\ln \frac{\lvert\Pi^*\rvert}{\gamma} + \sqrt{\ln \lvert\Pi^*\rvert \cdot 2 \ln \frac{2}{\delta}}}{\sqrt{\ln \lvert\Pi^*\rvert} \Delta}\right)^2$	$O\left(\frac{\ln T}{c \ln c}\left(\frac{\ln T}{\Delta^2} + \left(\frac{\ln \frac{\lvert\Pi^*\rvert}{\gamma}}{\sqrt{\ln \lvert\Pi^*\rvert} \Delta}\right)^2 + T\mu_{i^*}\right)\right)$
CAMS$_{\text{adversarial}}$	$2c\sqrt{\ln c / \rho_T} \cdot \sqrt{T \log \lvert\Pi^*\rvert}$	$O\left(\frac{\ln T}{c \ln c}\left(\sqrt{\frac{T \log(\lvert\Pi^*\rvert)}{\rho_T}} + \tilde{L}_{T,*}\right)\right)$
CAMS$_{\text{VCD}}$	$\left(2c\sqrt{\ln c} + 2\right)\sqrt{T \cdot \left(2d \ln \frac{eT}{d} + \ln \frac{2}{\delta}\right) / \rho_T}$	$O\left(\frac{\ln T}{c \ln c}\left(\sqrt{\frac{T \cdot \left(2d \ln \frac{eT}{d} + \ln \frac{2}{\delta}\right)}{\rho_T}} + \tilde{L}_{T,*}\right)\right)$

Experiments

Dataset characteristics

dataset	#classes	total instances	test set	stream size	classifier	policy
CIFAR10	10	60000	10000	10000	80	85
DRIFT	6	13910	3060	3000	10	11
VERTEBRAL	3	310	127	80	6	17
HIV	2	40000	4113	4000	4	20

context-free baselines: RS, QBC, IWAL, MP

contextual baselines: CQBC, CIWAL, Oracle

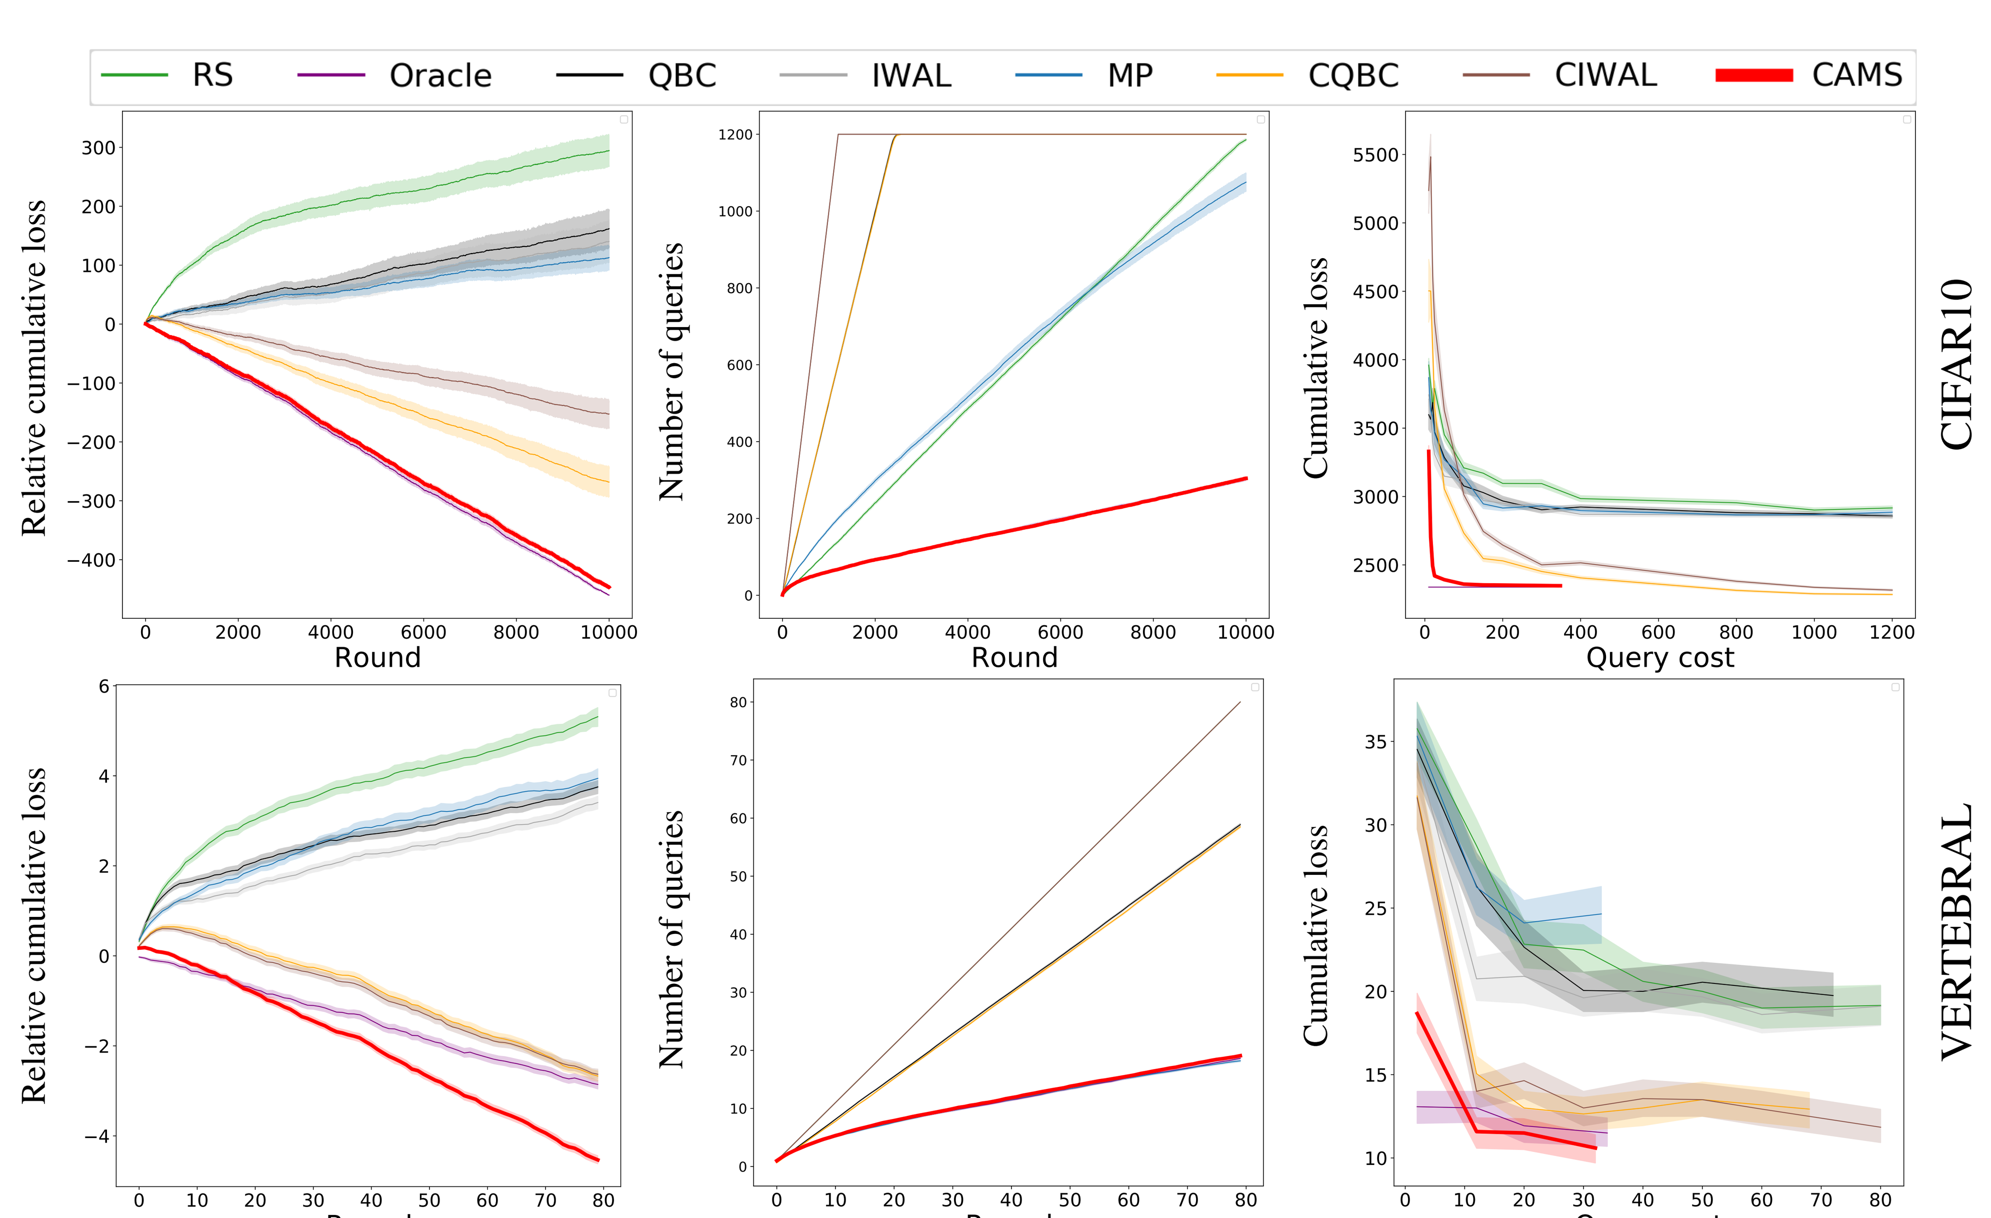

Key observations

effectively converge to the best expert

robustly recover from malicious advices

reduction to context-free setting when no experts are available

possibly outperforming all the experts including oracle

CAMS could achieve the above, while maintaining **label efficiency** and **low variance in performance**

On Vertebral, CAMS outperforms the Oracle despite 11 of the 17 experts giving malicious or random advice

check out more experimental results in paper …