Active Learning for Informative Projection Retrieval

Madalina Fiterau
mfiterau@cs.cmu.edu

Artur Dubrawski
awd@cs.cmu.edu

Motivation

Objectives:

- Noninvasive vital sign (VS) data collected in a Step-Down Unit (SDU) with alerts issued when a VS exceeds predefined thresholds
- Many alerts are artifacts, causing alarm fatigue
- Need to dismiss these artifacts
- Training classifiers for automatic artifact adjudication requires expert annotation
- We aim to reduce annotation effort

Approach:

- Regression-based Informative Projection Recovery (RIPR) facilitates expert annotation
- Requires fewer annotations to obtain an accurate classifier
- Results presented in a human-understandable form, low-dimensional projections

Outcome: Performing active learning reduces the number of alerts that need to be annotated by experts to train the artifact adjudication model. Our framework requires 48% of labels to train an accurate model, while a random forest classifier requires 89%.

Data Description

- Prospective longitudinal study recruited admissions over 8 weeks in a 24 bed trauma stepdown unit all with noninvasive VS monitoring:
  - Respiratory Rate (RR) from ECG bioimpedance
  - Systolic (SBP) and Diastolic (DBP) Blood Pressure (oscillometric)
  - Peripheral arterial oxygen saturation (SpO_2) by finger plethysmography
- VS data analyzed beyond local instability threshold values:
  - HR<40 or >140; RR<8 or >36; SBP <80 or >200; DBP>110, SpO_2<85%
  - Each alert associated with a category indicating the leading abnormal VS
  - 812 alerts of 3 types: RR, SpO_2, BP
  - 50 features computed, for each VS signal independently, during span of each alert, and a short window (4 minutes) preceding alert onset
  - Features include common statistics of each VS: mean, standard deviation, minimum, maximum, and range of values

Informative Projection Retrieval

- **IPR Problem:** Find a few simple projections of data in which alerts appear as either convincingly correct or easily dismissible
- **Difficulty:** Selecting the best set of projections and determining point assignment
- **Technique:** Machine Learning algorithm called RIPR: Regression-based Informative Projection Recovery [*]

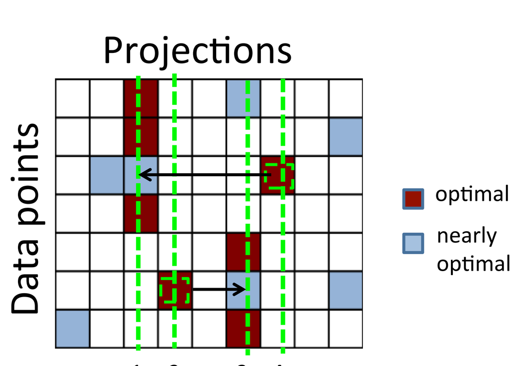

- RIPR selects a manageably small number of projections that jointly explain multiple alerts
- Each alert requires only one projection to be explained
- Low-dimensional projections allow easy interpretability
- RIPR enables automated classification of alerts

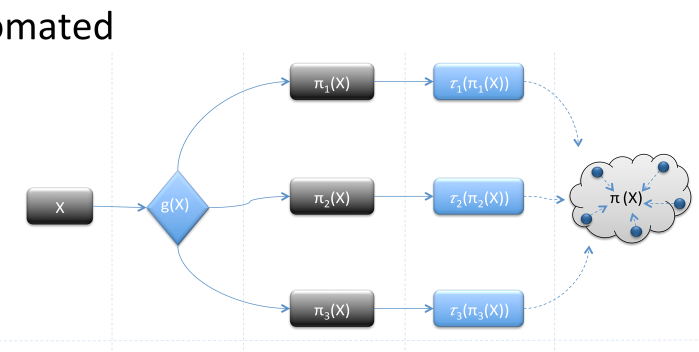

[*] M. Fiterau, A. Dubrawski, A Unified View of Informative Projection Retrieval, ICMLA 2013

Active Learning Framework

- **Challenge:** Very few existing labels, difficult to tell which projections are useful
- **Solution:** use active learning for existing models using various sample selection criteria
  -- uncertainty sampling, query by committee, information gain, conditional entropy --

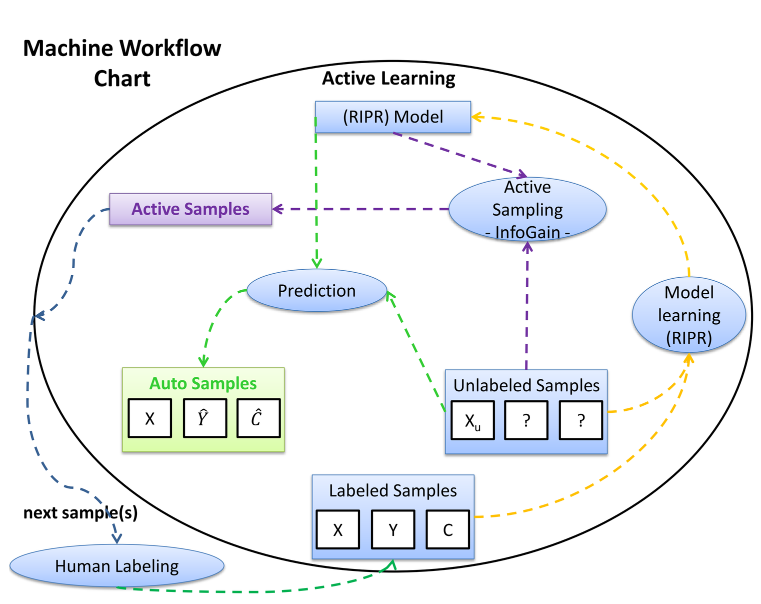

Sampling Type	Sampling Score given RIPR model
Uncertainty	$s_{uncrt}(x) = \min_{\pi \in \Pi^k_{uncrt}, \tau \in T^k_{uncrt}} \hat{h}(\tau(\pi(x))\|\pi(x))$
Query by Committee	$s_{qbc}(x) = \max_{\tau_i, \tau_j \in T^k_{qbc}} I(\tau_i(\pi_i(x)) \neq \tau_j(\pi_j(x)))$
Information Gain	$s_{ig}(x) = \hat{H}^k_{X_\ell, Y_\ell}(X^k_{u,ig}) - p(y=0)\hat{H}^k_{X_\ell \cup \{x\}, Y_\ell \cup \{0\}}(X^k_{u,ig}) - p(y=1)\hat{H}^k_{X_\ell \cup \{x\}, Y_\ell \cup \{1\}}(X^k_{u,ig}), \quad \forall x \in X^k_{u,ig}$
Low Conditional Entropy	$s_{mc}(x) = 1 - \min_{\pi \in \Pi^k_{mc}, \tau \in T^k_{mc}} \hat{h}(\tau(\pi(x))\|\pi(x))$

Experiments

Oxygen Saturation Alert Adjudication

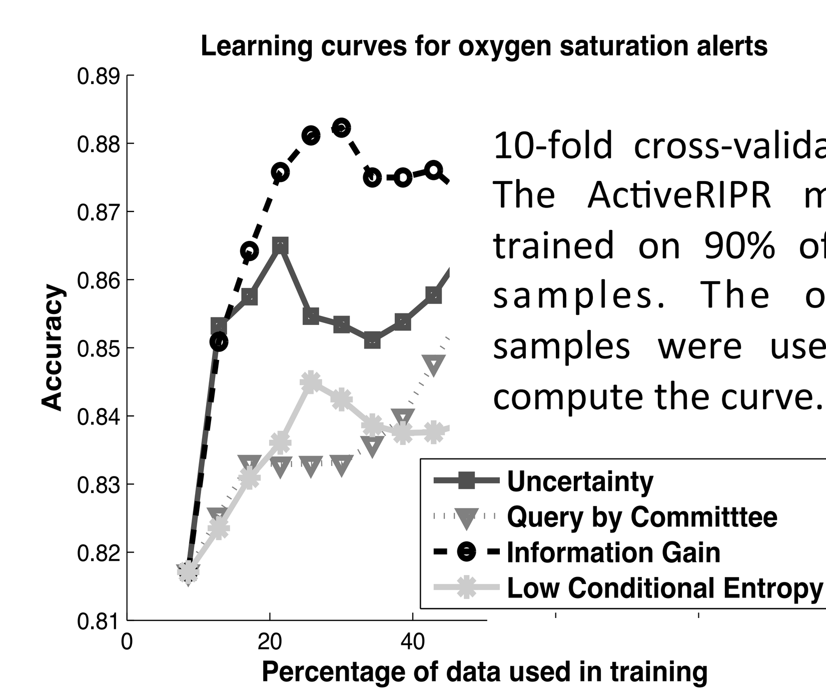

10-fold cross-validation: The ActiveRIPR model trained on 90% of the samples. The other samples were used to compute the curve.

We used ActiveRIPR to predict SpO_2 alerts, treating the expert-labeled data as the pool of samples available for active learning.

Percentage of samples needed for classification		Target accuracy 0.85	Target accuracy 0.88
Active Sampling Strategy	Uncertainty	18%	55%
	Query by Committee	46%	48%
	Information Gain	21%	25%
	Conditional Entropy	43%	46%

Very good performance in isolating SpO2 artifact, equivalent to what can be attained with 50% more annotated training data if the Active Learning protocol had not been used.

Blood Pressure Alert Adjudication

We used ActiveRIPR to predict BP alerts, using the expert-labeled pool of alerts. Leave-one-out cross validation was used due to the small amount of data available.

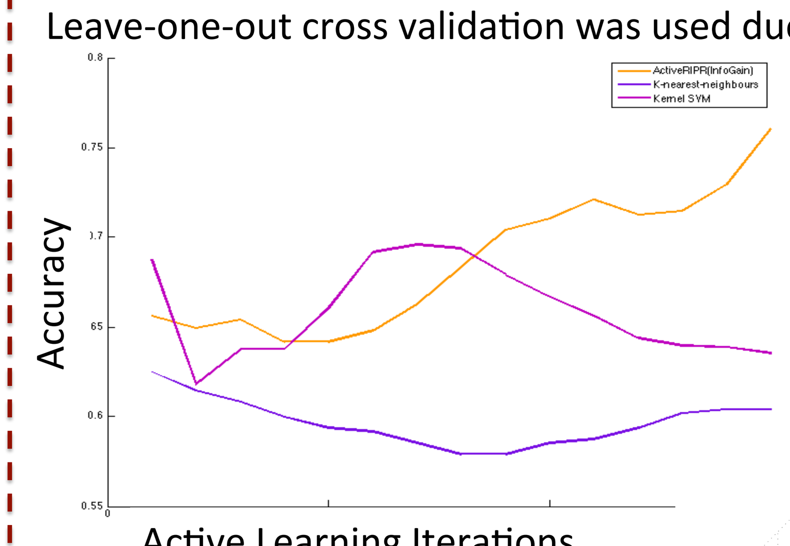

Active RIPR with the Information Gain Criterion, appears to be the most effective sampling method among tested alternatives. It outperforms other full dimensional classifiers using uncertainty sampling.

The informative projections recovered facilitate labeling

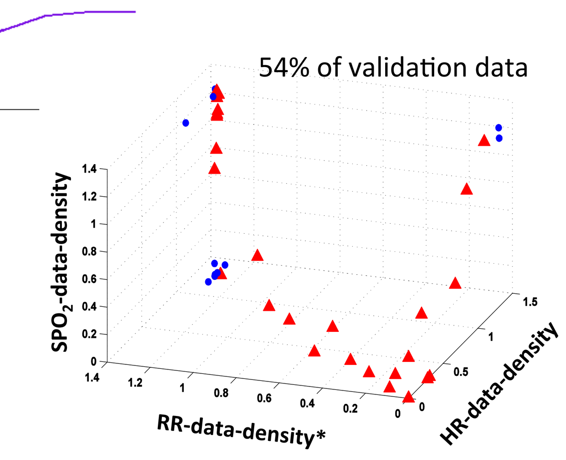

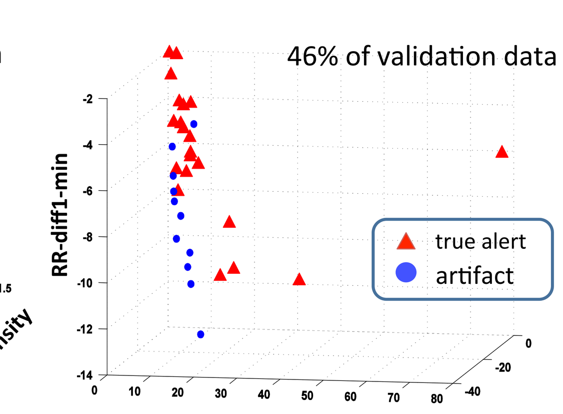

Annotation Examples

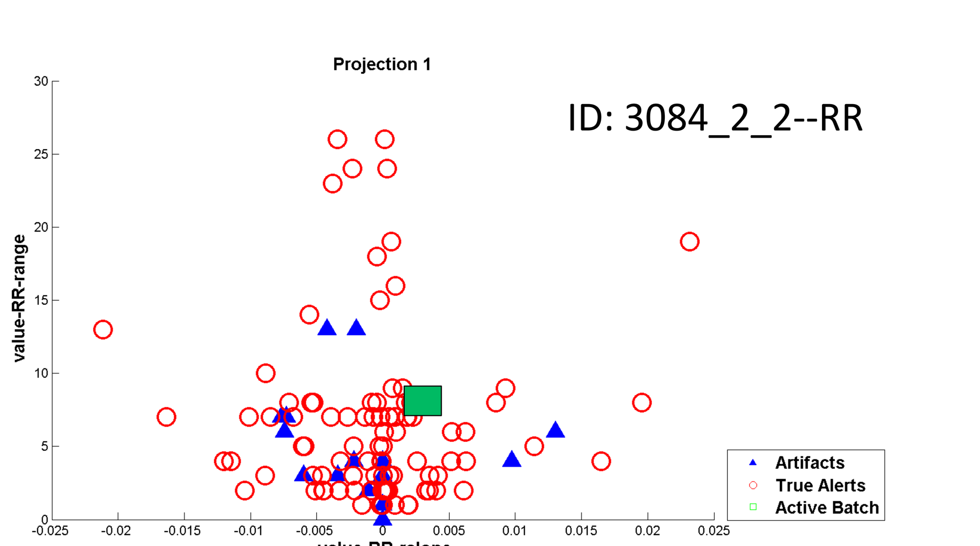

The sample *can* be confidently classified as a true alert.

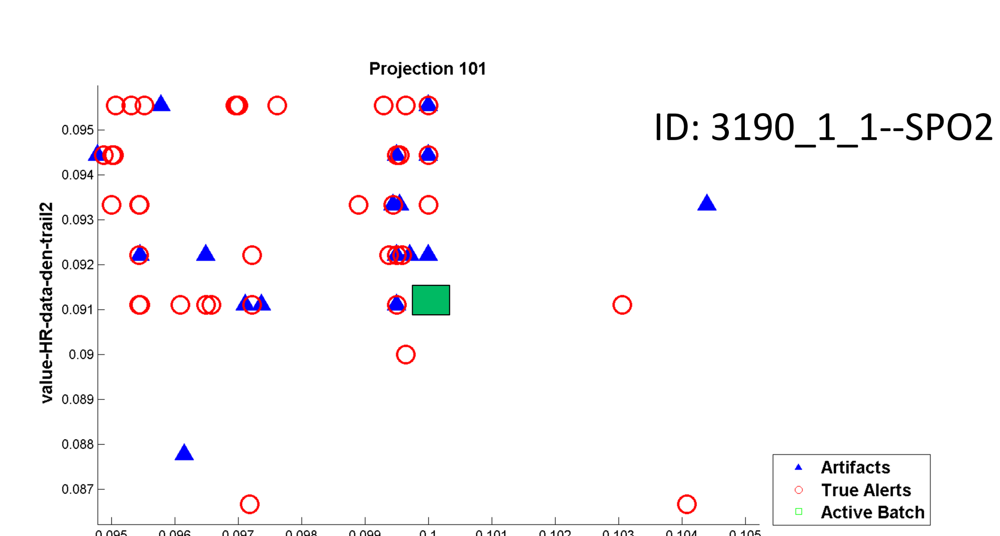

The sample *can* be somewhat confidently classified as an artifact.

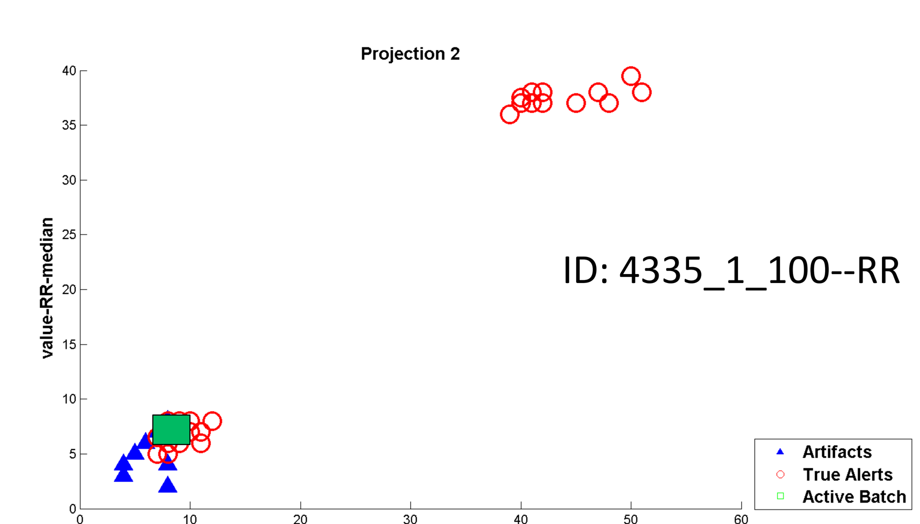

The sample *cannot* be confidently classified.

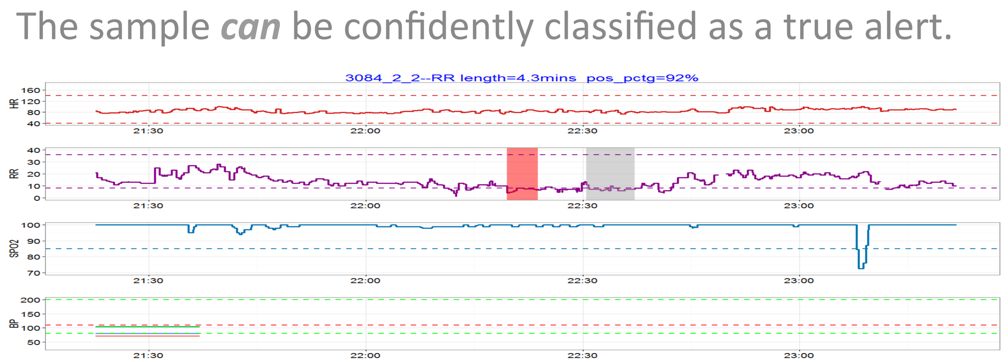

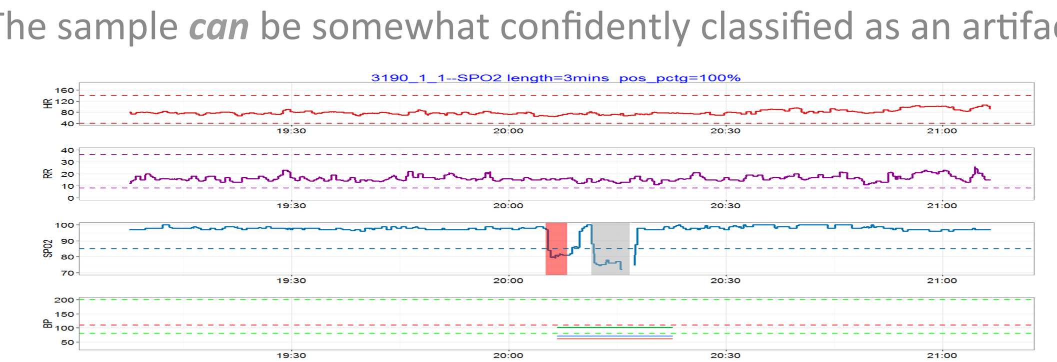

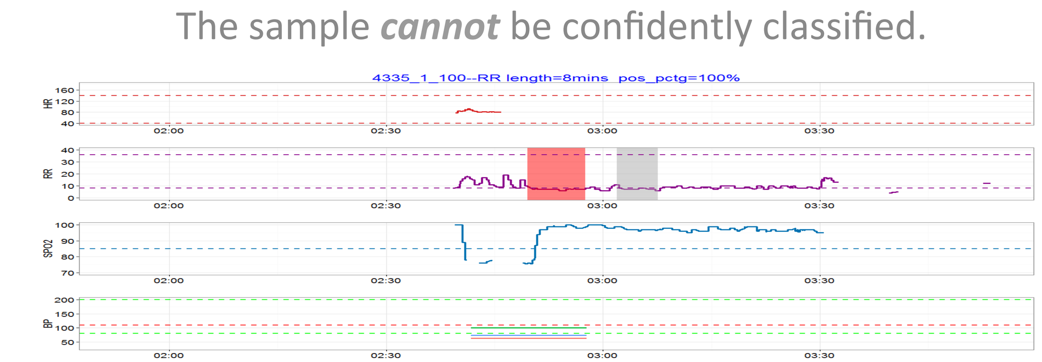

This work has been partially funded by the National Science Foundation (awards 0911032, 1320347) and the National Institutes of Health (R01NR013912)