Informative Features for Anomaly Detection

ICML @ NYC
Anomaly Detection Workshop

Allison Del Giorno, J. Andrew Bagnell, Martial Hebert
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15232
adelgior@cs.cmu.edu

THE ROBOTICS INSTITUTE
Carnegie Mellon University

Problem: Most anomaly detection algorithms are at the mercy of false positives in feature space

Causes of false alarms in anomaly detection:

- Little supervision
- Humans ignore specific changes
- Relevant features are unknown in advance

Illumination

Camera shake

Objective: Learn a feature mapping that reduces false positives by learning features that humans are uninterested in

Challenge I:
Training data may not tell the whole story*
*Or there may be no training data at all

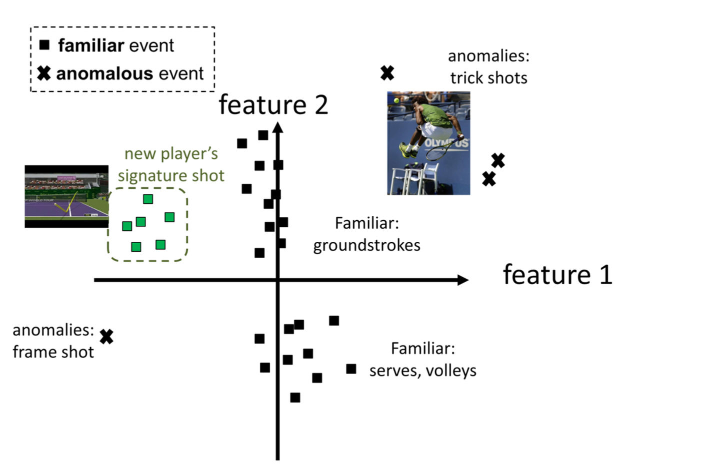

Challenge 2:
Definitions of anomalies are sensitive to features and context

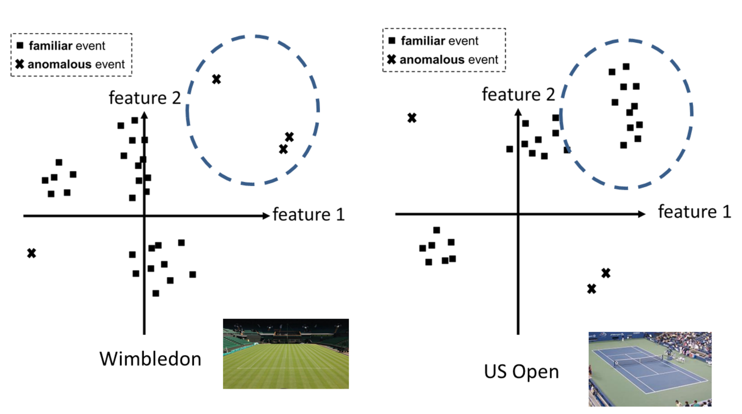

Challenge 3:
Humans and algorithms can easily disagree on the definition of anomalies

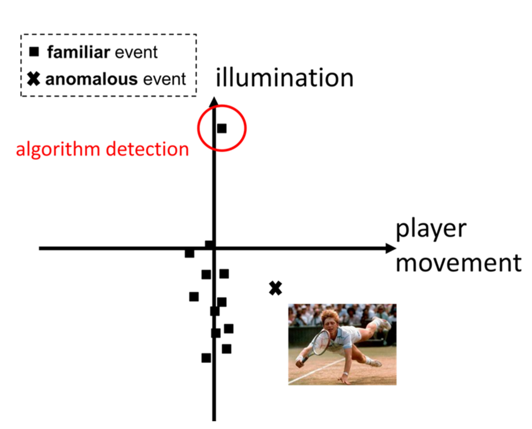

Constraints:

(1) Avoid anticipating anomalies

(2) Allow algorithms to exploit context at test time

Objective:

Learn human-defined invariances

Problem setup: Sets of normal data for training

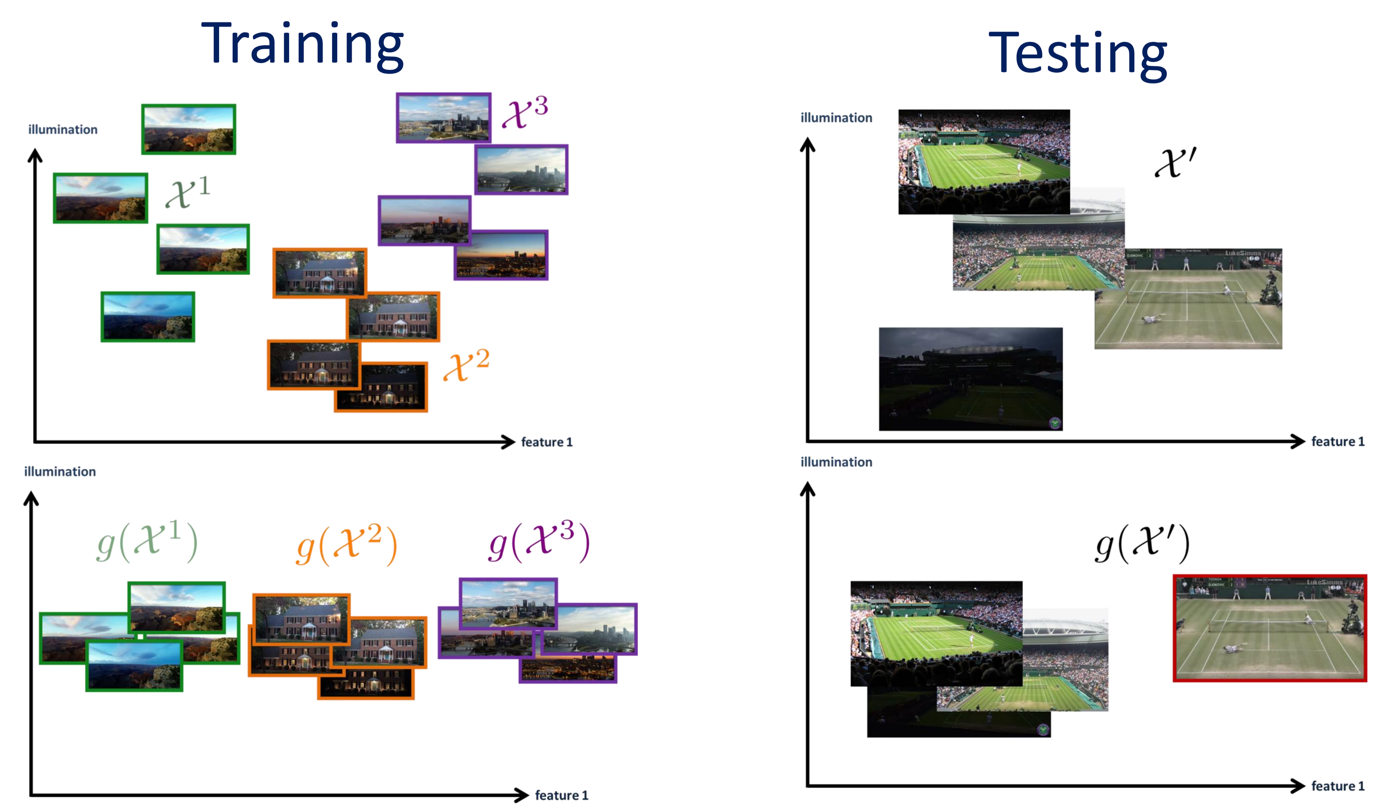

Method: A generalized eigenvalue problem learns invariance from within-set variance while preserving information across sets

Invariance: reduce *intra*set variance

$$\arg\min_g L_g^S(\mathcal{X}) = \sum_m \sum_{x,y\in\mathcal{X}^m} ||g(x)-g(y)||_2^2$$

Fidelity: preserve variance *across* sets:

$$\arg\max_g L_g^A(\mathcal{X}) = \sum_{x,y\in\mathcal{X}} ||g(x)-g(y)||_2^2$$

$$\bar{C} = C_1 + C_2 + C_3$$

$$\arg\min_{g,\lambda} L_g^S(\mathcal{X}) - \lambda L_g^A$$

$$g(x) = v^T x$$

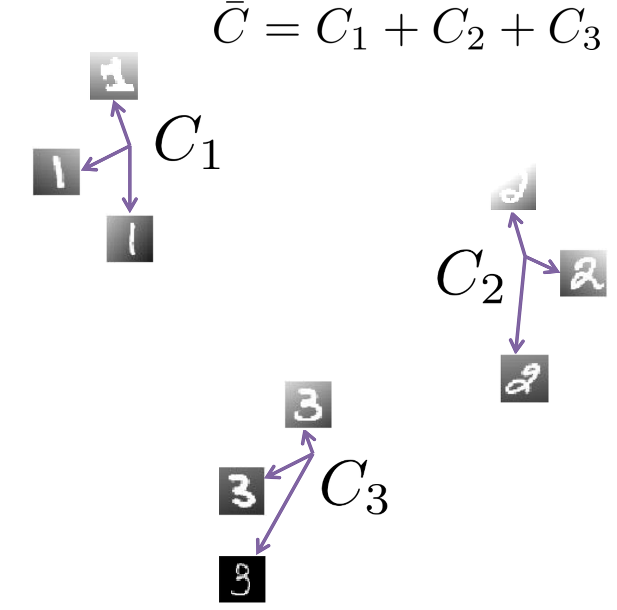

Final objective function:
Generalized eigenvalue problem

$$\arg\min_v \frac{v^T \bar{C} v}{v^T C v}$$

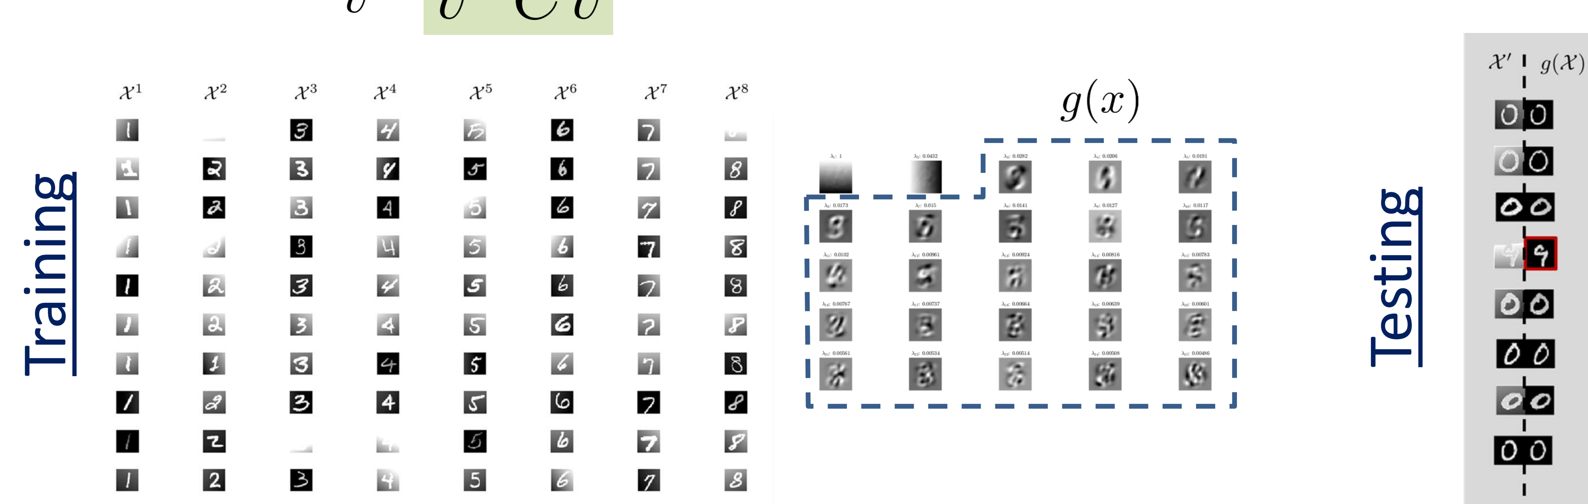

Results: Informative features improve performance using existing anomaly detection algorithms

Training

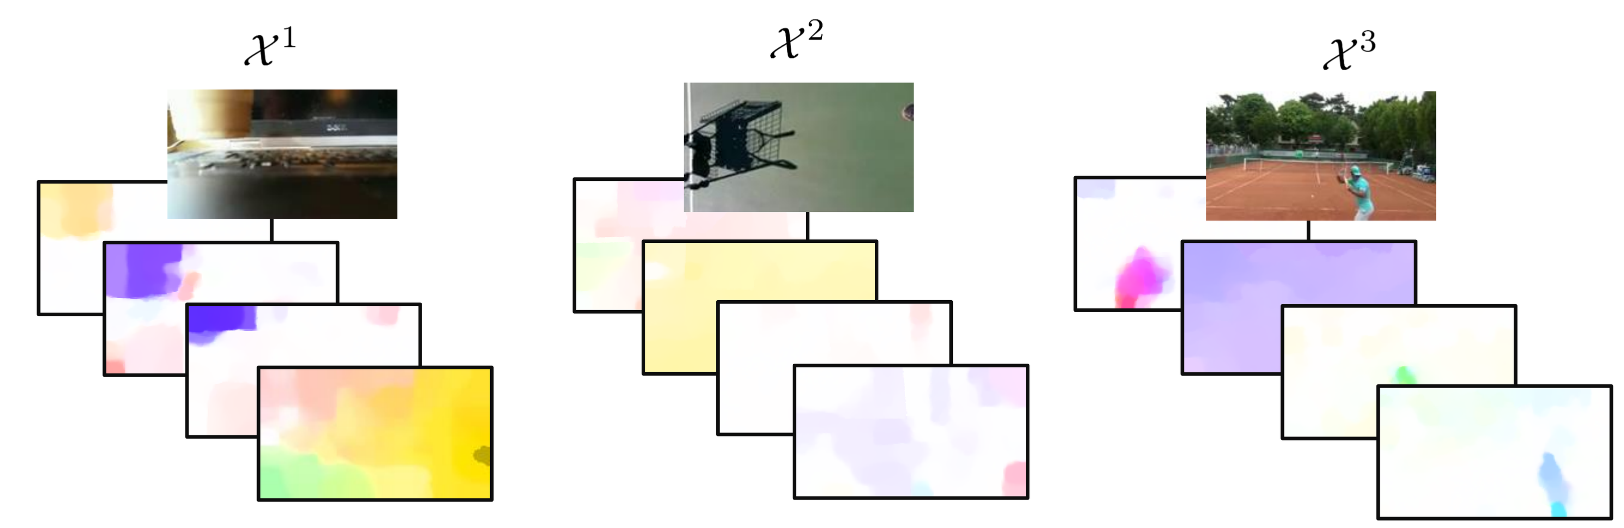

Testing

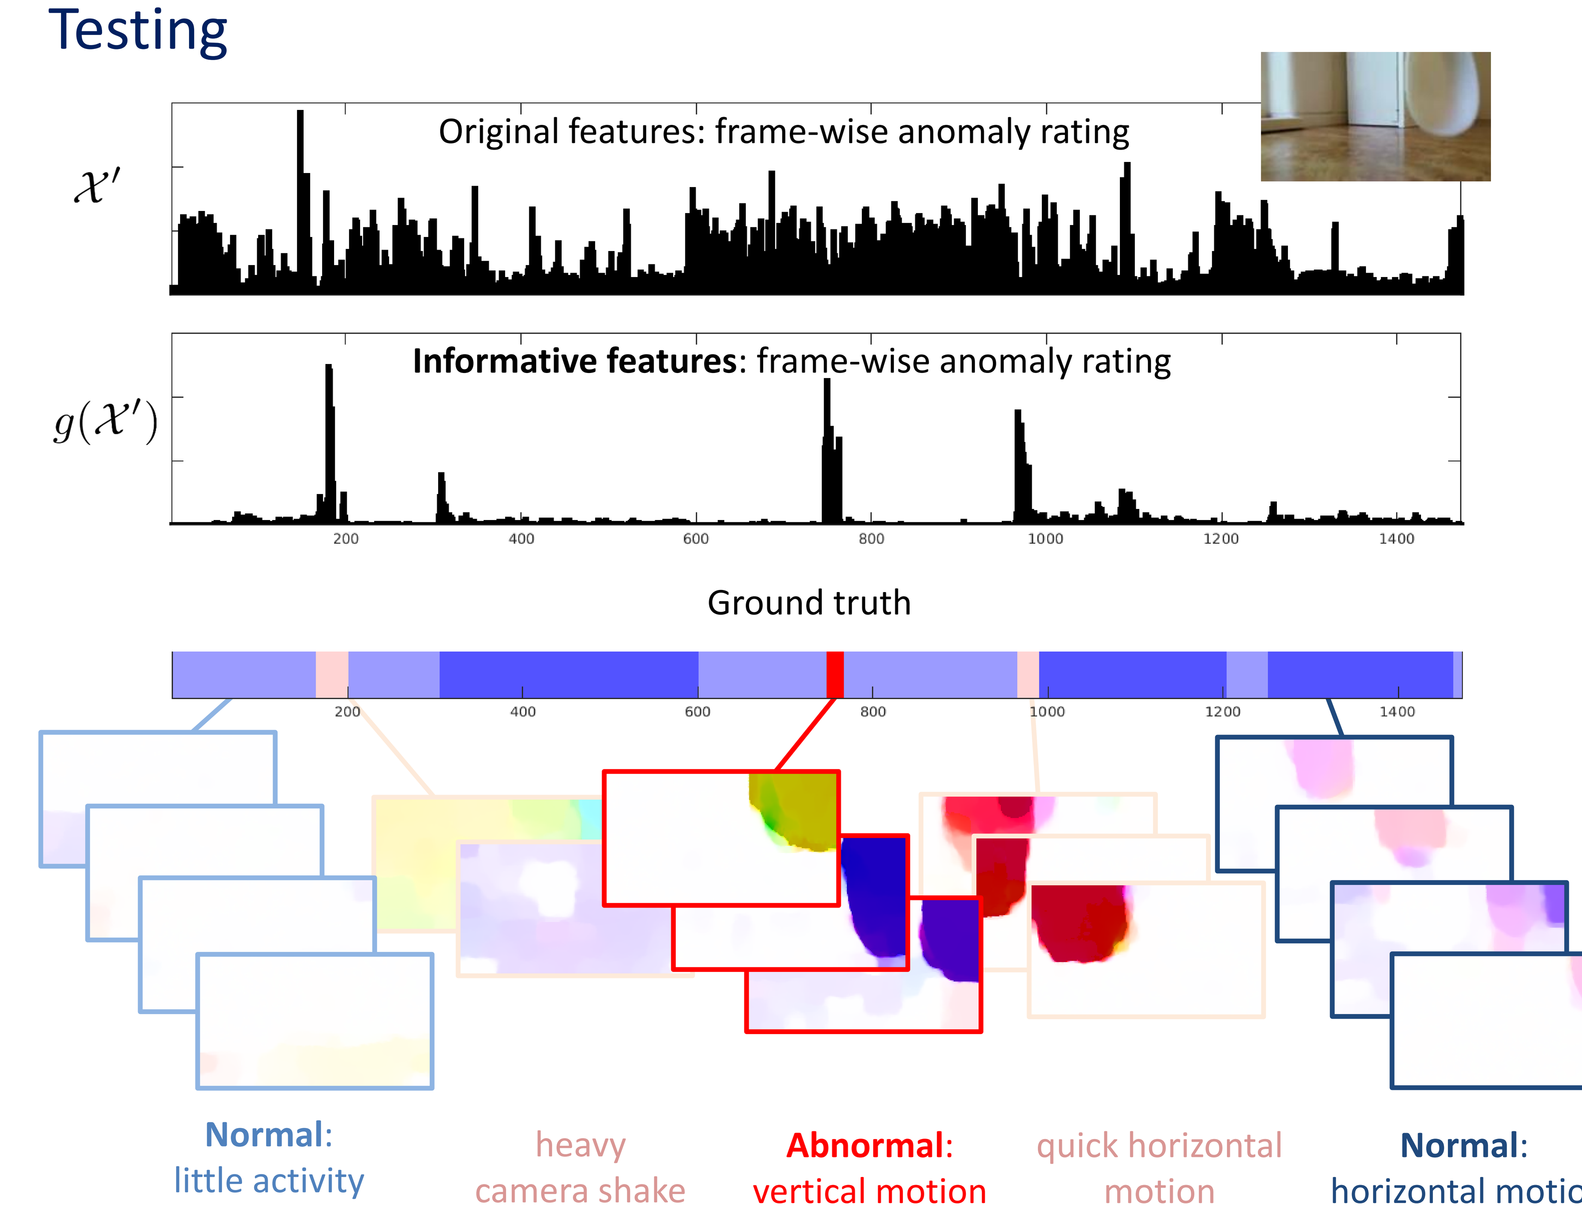

Future work includes:

- Human in-the-loop training
- Regularization for anomaly detection systems
- Nonlinear generalization
- Other ML problems (denoising, classification)
  - invariance sparse or easy to synthesize

Research supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program