

# Incorporating Privileged Information to Unsupervised Anomaly Detection

Shubhanshu Shekhar    Leman Akoglu

Carnegie Mellon University  
Heinz College of Information Systems and Public Policy  
{shubhras, lakoglu}@andrew.cmu.edu

**Abstract.** We introduce a new unsupervised anomaly detection ensemble called SPI which can harness *privileged* information—data available only for training examples but not for (future) test examples. Our ideas build on the Learning Using Privileged Information (LUPI) paradigm pioneered by Vapnik et al. [19,17], which we extend to unsupervised learning and in particular to anomaly detection. SPI (for Spotting anomalies with Privileged Information) constructs a number of frames/fragments of knowledge (i.e., density estimates) in the privileged space and *transfers* them to the anomaly scoring space through “imitation” functions that use only the partial information available for test examples. Our generalization of the LUPI paradigm to unsupervised anomaly detection shepherds the field in several key directions, including (i) *domain-knowledge-augmented detection* using expert annotations as PI, (ii) *fast detection* using computationally-demanding data as PI, and (iii) *early detection* using “historical future” data as PI. Through extensive experiments on simulated and real datasets, we show that augmenting privileged information to anomaly detection significantly improves detection performance. We also demonstrate the promise of SPI under all three settings (*i–iii*); with PI capturing expert knowledge, computationally-expensive features, and future data on three real world detection tasks.

## 1 Introduction

Outlier detection in point-cloud data has been studied extensively [1]. In this work we consider a unique setting with a much sparser literature: the problem of augmenting privileged information into unsupervised anomaly detection. Simply put, privileged information (PI) is *additional* data/knowledge/information that is available only at the learning/model building phase for (subset of) training examples, which however is unavailable for (future) test examples.

**The LUPI framework.** Learning Using Privileged Information (LUPI) has been pioneered by Vapnik et al. first in the context of SVMs [19,17] (PI-incorporated SVM is named SVM+), later generalized to neural networks [18]. The setup involves an Intelligent (or non-trivial) Teacher at learning phase, who provides the Student with *privileged* information (like explanations, metaphors, etc.), denoted  $\mathbf{x}_i^*$ , about each training example  $\mathbf{x}_i$ ,  $i = 1 \dots n$ . The key point in this paradigm is that *privileged information is not available at the test phase*

(when Student operates without guidance of Teacher). Therefore, the goal is to build models (in our case, detectors) that can leverage/incorporate such additional information but *yet, not depend on the availability of PI at test time*.

*Example:* The additional information  $\mathbf{x}_i^*$ 's belong to space  $X^*$  which is, generally speaking, different from space  $X$ . In other words, the feature spaces of vectors  $\mathbf{x}_i^*$ 's and  $\mathbf{x}_i$ 's do not overlap. As an example, consider the task of identifying cancerous biopsy images. Here the images are in pixel space  $X$ . Suppose that there is an Intelligent Teacher that can recognize patterns in such images relevant to cancer. Looking at a biopsy image, Teacher can provide a description like "Aggressive proliferation of A-cells into B-cells" or "Absence of any dynamic". Note that such descriptions are in a specialized language space  $X^*$ , different from pixel space  $X$ . Further, they would be available only for a set of examples and not when the model is to operate autonomously in the future.

*LUPI's advantages:* LUPI has been shown to (i) improve rate of convergence for learning, i.e., require asymptotically fewer examples to learn [19], as well as (ii) improve accuracy, when one can learn a model in space  $X^*$  that is not much worse than the best model in space  $X$  (i.e., PI is intelligent/non-trivial) [18]. Motivated by these advantages, LUPI has been applied to a number of problems from action recognition [13] to risk modeling [14] (expanded in §5). However, the focus of all such work has mainly been on *supervised* learning.

**LUPI for anomaly detection.** The only (perhaps straightforward) extension of LUPI to unsupervised anomaly detection has been introduced recently, generalizing SVM+ to the One-Class SVM (namely OC-SVM+) [2] for malware and bot detection. The issue is that OC-SVM is not a reliable detector since it assumes that normal points can be separated from origin in a single hyperball—experiments on numerous benchmark datasets with ground truth by Emmott et al. that compared popular anomaly detection algorithms find that OC-SVM ranks at the bottom (Table 1, pg. 4 [6]; also see our results in §4). We note that the top performer in [6] is the Isolation Forest (iForest) algorithm [11], an ensemble of randomized trees.

*Our contributions:* Motivated by LUPI's potential value to learning and the scarcity in the literature of its generalization to anomaly detection, we propose a new technique called SPI (pronounced 'spy'), for Spotting anomalies with Privileged Information. Our work *bridges the gap (for the first time) between LUPI and unsupervised ensemble based anomaly detection* that is considered state-of-the-art [6]. We summarize our main contributions as follows.

- **Study of LUPI for anomaly detection:** We analyze how LUPI can benefit anomaly detection, not only when PI is truly unavailable at test time (as in traditional setup) but also when PI is strategically and willingly avoided at test time. We argue that data/information that incurs overhead on resources (\$\$\$/storage/battery/etc.), timeliness, or vulnerability, if designated as PI, can enable resource-frugal, early, and preventive detection (expanded in §2).
- **PI-incorporated detection algorithm:** We show how to incorporate PI into ensemble based detectors and propose SPI, which constructs frames/fragments of knowledge (specifically, density estimates) in the privileged space ( $X^*$ ) and *transfers* them to the anomaly scoring space ( $X$ )

through “imitation” functions that use only the partial information available for test examples. To the best of our knowledge, ours is the first attempt to leveraging PI for improving the state-of-the-art *ensemble methods* for anomaly detection within an *unsupervised* LUPI framework. Moreover, while SPI augments PI within the tree-ensemble detector iForest [11], our solution can easily be applied to any other ensemble based detector (§3).

- **Applications:** Besides extensive simulation experiments, we employ SPI on three real-world case studies where PI respectively captures (i) expert knowledge, (ii) computationally-expensive features, and (iii) “historical future” data, which demonstrate the benefits that PI can unlock for anomaly detection in terms of accuracy, speed, and detection latency (§4).

**Reproducibility:** Implementation of SPI and real world datasets used in experiments are open-sourced at <http://www.andrew.cmu.edu/user/shubhras/SPI>.

## 2 Motivation: How can LUPI benefit anomaly detection?

The implications of the LUPI paradigm for anomaly detection is particularly exciting. Here, we discuss a number of detection scenarios and demonstrate that LUPI unlocks advantages for anomaly detection problems in multiple aspects.

In the original LUPI framework [19], privileged information (hereafter PI) is defined as data that is available *only* at training stage for training examples but *unavailable at test time* for test examples. Several anomaly detection scenarios admit this definition directly. Interestingly, PI can also be specified as *strategically “unavailable”* for anomaly detection. That is, one can willingly avoid using certain data at test time (while incorporating such data into detection models at train phase<sup>1</sup>) in order to achieve resource efficiency, speed, and robustness. We organize detection scenarios into two with PI as (truly) Unavailable vs. Strategic, and elaborate with examples below. Table 1 gives a summary.

Table 1: Types of data used in anomaly detection with various overhead on resources (\$\$\$, storage, battery, etc.), timeliness, and/or risk, *if used as privileged information can enable resource-frugal, early, as well as preventive detection.*

Properties vs. Type of Privileged Info	Unavailable vs. Strategic	need Resources	cause Delay	incur Risk
1. “historical future” data	U	n/a	n/a	n/a
2. after-the-fact data	U	n/a	n/a	n/a
3. advanced technical data	U	n/a	n/a	n/a
4. restricted-access data	U, S	✓		
5. expert knowledge	U, S	✓	✓	
6. compute-heavy data	S	✓	✓	
7. unsafe-to-collect data	S		✓	✓
8. easy-target-to-tamper data	S			✓

**Unavailable PI:** This setting includes typical scenarios, where PI is (truly) unknown for test examples.

<sup>1</sup> Note that training phase in anomaly detection does not involve the use of any labels.

1. *“historical future” data*: When training an anomaly detection model with offline/historical data that is over time (e.g., temporal features), one may use values both before *and after* time  $t$  while creating an example for each  $t$ . Such data is PI; not available when the model is deployed to operate in real-time.

2. *after-the-fact data*: In malware detection, the goal is to detect before it gets hold of and harms the system. One may have historical data for some (training) examples from past exposures, including measurements of system variables (number of disk/port read/writes, CPU usage, etc.). Such after-the-exposure measurements can be incorporated as PI.

3. *advanced technical data*: This includes scenarios where some (training) examples are well-understood but those to be detected are simply unknown. For example, the expected behavior of various types of apps on a system may be common domain knowledge that can be converted to PI, but such knowledge may not (yet) be available for new-coming apps.

**Strategic PI**: Strategic scenarios involve PI that can in principle be acquired but is willingly avoided at test time to achieve gains in resources, time, or risk.

4. *restricted-access data*: One may want to build models that do not assume access to private data or intellectual property at test time, such as source code (for apps or executables), *even if* they could be acquired through resources. Such information can also be truly unavailable, e.g. encrypted within the software.

5. *expert knowledge*: Annotations about some training examples may be available from experts, which are truly unavailable at test time. One could also strategically choose to avoid expert involvement at test time, which (a) may be costly to obtain and/or (b) cause significant delay, especially for real-time detection.

6. *compute-heavy data*: One may strategically choose not to rely on features that are computationally expensive to obtain, especially in real-time detection, but rather use such data as PI (which can be extracted offline at training phase). Such features not only cause delay but also require compute resources (which e.g., may drain batteries in detecting malware apps on cellphones).

7. *unsafe-to-collect data*: This involves cases where collecting PI at test time is unsafe/dangerous. For example, the slower a drone moves to capture high-resolution (privileged) images for surveillance, not only it causes delay but more importantly, the more susceptible it becomes to be taken down.

8. *easy-target-to-tamper data*: Finally, one may want to avoid relying on features that are easy for adversaries to tamper with. Examples to those features include self-reported data (like age, location, etc.). Such data may be available reliably for some training examples and can be used as PI.

In short, by strategically designating PI one can achieve resource, timeliness, and robustness gains for various anomaly detection tasks. Designating features that need resources as PI  $\rightarrow$  allow resource-frugal (“lazy”) detection; features that cause delay as PI  $\rightarrow$  allow early/speedy detection; and designating features that incur vulnerability as PI  $\rightarrow$  allow preventive and more robust detection.

In this subsection, we laid out a long list of scenarios that make LUPI-based learning particularly attractive for anomaly detection. In our experiments (§4) we demonstrate its premise for scenarios 1., 5. and 6. above using three real world datasets, while leaving others as what we believe interesting future investigations.

### 3 Privileged Info-Augmented Anomaly Detection

**The Learning Setting** Formally, the input for the anomaly detection model at learning phase are tuples of the form

$$\mathcal{D} = \{(\mathbf{x}_1, \mathbf{x}_1^*), (\mathbf{x}_2, \mathbf{x}_2^*), \dots, (\mathbf{x}_n, \mathbf{x}_n^*)\},$$

where  $\mathbf{x}_i = (x_i^1, \dots, x_i^d) \in X$  and  $\mathbf{x}_i^* = (x_i^{*1}, \dots, x_i^{*p}) \in X^*$ . Note that this is an unsupervised learning setting where label information, i.e.,  $y_i$ 's are not available. The privileged information is represented as a feature vector  $\mathbf{x}^* \in \mathbb{R}^p$  that is in space  $X^*$ , which is *additional to and different from* the feature space  $X$  in which the primary information is represented as a feature vector  $\mathbf{x} \in \mathbb{R}^d$ .

The important distinction from the traditional anomaly detection setting is that the input to the (trained) detector at testing phase are feature vectors

$$\{\mathbf{x}_{n+1}, \mathbf{x}_{n+2}, \dots, \mathbf{x}_{n+m}\}.$$

That is, the (future) test examples do not carry any privileged information. The anomaly detection model is to score the incoming/test examples and make decisions solely based on the primary features  $\mathbf{x} \in X$ .

In this text, we refer to space  $X^*$  as the *privileged space* and to  $X$  as the *decision space*. Here, a key assumption is that the information in the privileged space is intelligent/nontrivial, that is, it allows to create models  $f^*(\mathbf{x}^*)$  that detect anomalies with vectors  $\mathbf{x}^*$  corresponding to vectors  $\mathbf{x}$  with higher accuracy than models  $f(\mathbf{x})$ . As a result, the main question that arises which we address in this work is: "how can one use the knowledge of the information in space  $X^*$  to improve the performance of the desired model  $f(\mathbf{x})$  in space  $X$ ?"

In what follows, we present a first-cut attempt to the problem that is a natural knowledge transfer between the two feature spaces (called FT for feature transfer). We then lay out the shortcomings of such an attempt, and present our proposed solution SPI. We compare to FT (and other baselines) in experiments.

#### 3.1 First Attempt: Incorporating PI by Transfer of Features

A natural attempt to learning under privileged information that is unavailable for test examples is to treat the task as a *missing data problem*. Then, typical techniques for data imputation can be employed where missing (privileged) features are replaced with their predictions from the available (primary) features.

In this scheme, one simply maps vectors  $\mathbf{x} \in X$  into vectors  $\mathbf{x}^* \in X^*$  and then builds a detector model in the transformed space. The goal is to find the transformation of vectors  $\mathbf{x} = (x^1, \dots, x^d)$  into vectors  $\phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \dots, \phi_p(\mathbf{x}))$  that minimizes the expected risk given as

$$R(\phi) = \sum_{j=1}^p \min_{\phi_j} \int (x^{*j} - \phi_j(\mathbf{x}))^2 p(x^{*j}, \mathbf{x}) dx^{*j} d\mathbf{x}, \quad (1)$$

where  $p(x^{*j}, \mathbf{x})$  is the joint probability of coordinate  $x^{*j}$  and vector  $\mathbf{x}$ , and functions  $\phi_j(\mathbf{x})$  are defined by  $p$  regressors.

Here, one could construct approximations to functions  $\phi_j(\mathbf{x})$ ,  $j = \{1, \dots, p\}$  by solving  $p$  regression estimation problems based on the training examples

$$(\mathbf{x}_1, x_1^{*j}), \dots, (\mathbf{x}_n, x_n^{*j}), \quad j = 1, \dots, p,$$

where  $\mathbf{x}_i$ 's are input to each regression  $\phi_j$  and the  $j$ th coordinate of the corresponding vector  $\mathbf{x}_i^*$ , i.e.  $x_i^{*j}$ 's are treated as the output, by minimizing the regularized empirical loss functional

$$R(\phi_j) = \min_{\phi_j} \sum_{i=1}^n (x_i^{*j} - \phi_j(\mathbf{x}_i))^2 + \lambda_j \text{penalty}(\phi_j), \quad j = 1, \dots, p. \quad (2)$$

Having estimated the transfer functions  $\hat{\phi}_j$ 's (using linear or non-linear regression techniques), one can then learn any desired anomaly detector  $f(\hat{\phi}(\mathbf{x}))$  using the training examples, which concludes the learning phase. Note that the detector does not require access to privileged features  $\mathbf{x}^*$  and can be employed solely on primary features  $\mathbf{x}$  of the test examples  $i = n + 1, \dots, m$ .

### 3.2 Proposed SPI: Incorporating PI by Transfer of Decisions

Treating PI as missing data and predicting  $\mathbf{x}^*$  from  $\mathbf{x}$  could be a difficult task, when privileged features are complex and high dimensional (i.e.,  $p$  is large). Provided  $f^*(\mathbf{x}^*)$  is an accurate detection model, a more direct goal would be to *mimic its decisions*—the scores that  $f^*$  assigns to the training examples. Mapping *data* between two spaces, as compared to *decisions*, would be attempting to solve a more general problem, that is likely harder and unnecessarily wasteful.

The general idea behind transferring decisions/knowledge (instead of data) is to identify a small number of elements in the privileged space  $X^*$  that well-approximate the function  $f^*(\mathbf{x}^*)$ , and then try to transfer them to the decision space—through the approximation of those elements in space  $X$ . This is the knowledge transfer mechanism in LUPI by Vapnik and Izmailov [17]. They illustrated this mechanism for the (supervised) SVM classifier. We generalize this concept to unsupervised anomaly detection.

The knowledge transfer mechanism uses three building blocks of knowledge representation in AI, as listed in Table 2. We first review this concept for SVMs, followed by our proposed SPI. While SPI is clearly different in terms of the task it is addressing as well as in its approach, as we will show, it is inspired by and builds on the same fundamental mechanism.

Table 2: Three building blocks of knowledge representation in artificial intelligence, in context of SVM-LUPI for classification [17] and SPI for anomaly detection [this paper].

	SVM-LUPI	SPI (Proposed)
1. Fundamental elements of knowledge	support vectors	isolation trees
2. Frames (fragments) of the knowledge	kernel functions	tree anomaly scores
3. Structural connections of the frames	weighted sum	weighted sum (by L2R)

**Knowledge transfer for SVM:** The *fundamental elements* of knowledge in the SVM classifier are the support vectors. In this scheme, one constructs two SVMs; one in  $X$  space and another in  $X^*$  space. Without loss of generality, let  $\mathbf{x}_1, \dots, \mathbf{x}_t$  be the support vectors of SVM solution in space  $X$  and  $\mathbf{x}_1^*, \dots, \mathbf{x}_{t^*}^*$  be the support vectors of SVM solution in space  $X^*$ , where  $t$  and  $t^*$  are the respective number of support vectors.

The decision rule  $f^*$  in space  $X^*$  (which one aims to mimic) has the form

$$f^*(\mathbf{x}^*) = \sum_{k=1}^{t^*} y_k \alpha_k^* K^*(\mathbf{x}_k^*, \mathbf{x}^*) + b^*, \quad (3)$$

where  $K^*(\mathbf{x}_k^*, \mathbf{x}^*)$  is the kernel function of similarity between support vector  $\mathbf{x}_k^*$  and vector  $\mathbf{x}^* \in X^*$ , also referred as the *frames* (or *fragments*) of knowledge. Eq. (3) depicts the *structural connection* of these fragments, which is a weighted sum with learned weights  $\alpha_k^*$ 's.

The goal is to approximate each fragment of knowledge  $K^*(\mathbf{x}_k^*, \mathbf{x}^*)$ ,  $k = 1, \dots, t^*$  in  $X^*$  using the fragments of knowledge in  $X$ ; i.e., the  $t$  kernel functions  $K(\mathbf{x}_1, \mathbf{x}), \dots, K(\mathbf{x}_t, \mathbf{x})$  of the SVM trained in  $X$ . To this end, one maps  $t$ -dimensional vectors  $\mathbf{z} = (K(\mathbf{x}_1, \mathbf{x}), \dots, K(\mathbf{x}_t, \mathbf{x})) \in Z$  into  $t^*$ -dimensional vectors  $\mathbf{z}^* = (K^*(\mathbf{x}_1^*, \mathbf{x}^*), \dots, K^*(\mathbf{x}_{t^*}^*, \mathbf{x}^*)) \in Z^*$  through  $t^*$  regression estimation problems. That is, the goal is to find regressors  $\phi_1(\mathbf{z}), \dots, \phi_{t^*}(\mathbf{z})$  in  $X$  such that

$$\phi_k(\mathbf{z}_i) \approx K^*(\mathbf{x}_k^*, \mathbf{x}_i^*), \quad k = 1, \dots, t^* \quad (4)$$

for all training examples  $i = 1, \dots, n$ . For each  $k = 1, \dots, t^*$ , one can construct the approximation to function  $\phi_k$  by training a regression on the data

$$\{(\mathbf{z}_1, K^*(\mathbf{x}_k^*, \mathbf{x}_1^*)), \dots, (\mathbf{z}_n, K^*(\mathbf{x}_k^*, \mathbf{x}_n^*))\}, \quad k = 1, \dots, t^*,$$

where we regress vectors  $\mathbf{z}_i$ 's onto scalar output  $K^*(\mathbf{x}_k^*, \mathbf{x}_i^*)$ 's to obtain  $\hat{\phi}_k$ .

For the prediction of a test example  $\mathbf{x}$ , one can then replace each  $K^*(\mathbf{x}_k^*, \mathbf{x}^*)$  in Eq. (3) (which requires privileged features  $\mathbf{x}^*$ ) with  $\hat{\phi}_k(\mathbf{z})$  (which mimics it, using only the primary features  $\mathbf{x}$ —to be exact, by first transforming  $\mathbf{x}$  into  $\mathbf{z}$  through the frames  $K(\mathbf{x}_j, \mathbf{x}), j = 1, \dots, t$  in the  $X$  space).

**Knowledge transfer for SPI:** In contrast to mapping of features from space  $X$  to space  $X^*$ , knowledge transfer of decisions maps space  $Z$  to  $Z^*$  in which fragments of knowledge are represented. Next, we show how to generalize these ideas to anomaly detection with no label supervision. Figure 1 shows an overview.

To this end, we utilize a state-of-the-art ensemble technique for anomaly detection, called Isolation Forest [11] (hereafter iF, for short), which builds a set of extremely randomized trees. In essence, each tree approximates density in a random feature subspace and anomalousness of a point is quantified by the sum of such partial estimates across all trees.

In this setting, one can think of the individual trees in the ensemble to constitute the *fundamental elements* and the partial density estimates (i.e., individual anomaly scores from trees) to constitute the *fragments* of knowledge, where the structural connection of the fragments is achieved by an unweighted sum.

Similar to the scheme with SVMs, we construct two iFs; one in  $X$  space and another in  $X^*$  space. Let  $\mathcal{T} = T_1, \dots, T_t$  denote the trees in the ensemble in  $X$  and  $\mathcal{T}^* = T_1^*, \dots, T_{t^*}^*$  the trees in the ensemble in  $X^*$ , where  $t$  and  $t^*$  are the respective number of trees (prespecified by the user, typically a few 100s). Further, let  $S^*(T_k^*, \mathbf{x}^*)$  denote the anomaly score estimated by tree  $T_k^*$  for a given  $\mathbf{x}^*$  (the lower the more anomalous; refer to [11] for details of the scoring).  $S(T_k, \mathbf{x})$  is defined similarly. Then, the anomaly score  $s^*$  for a point  $\mathbf{x}^*$  in space  $X^*$  (which we aim to mimic) is written as




Fig. 1: Anomaly detection with PI illustrated. FT maps *data* between spaces (§3.1) whereas SPI (and “light” version SPI-LITE) mimic *decisions* (§3.2).

$$s^*(\mathbf{x}^*) = \sum_{k=1}^{t^*} S^*(T_k^*, \mathbf{x}^*), \quad (5)$$

which is analogous to Eq. (3). To mimic/approximate each fragment of knowledge  $S^*(T_k^*, \mathbf{x}^*)$ ,  $k = 1, \dots, t^*$  in  $X^*$  using the fragments of knowledge in  $X$ ; i.e., the  $t$  scores for  $\mathbf{x}$ :  $S(T_1, \mathbf{x}), \dots, S(T_t, \mathbf{x})$  of the IF trained in  $X$ , we estimate  $t^*$  regressors  $\phi_1(z), \dots, \phi_{t^*}(z)$  in  $X$  such that

$$\phi_k(z_i) \approx S^*(T_k^*, \mathbf{x}_i^*), \quad k = 1, \dots, t^* \quad (6)$$

for all training examples  $i = 1, \dots, n$ , where  $z_i = (S(T_1, \mathbf{x}_i), \dots, S(T_t, \mathbf{x}_i))$ . Simply put, each  $\hat{\phi}_k$  is an approximate mapping of all the  $t$  scores from the ensemble  $\mathcal{T}$  in  $X$  to an individual score (fragment of knowledge) by tree  $T_k^*$  of the ensemble  $\mathcal{T}^*$  in  $X^*$ . In practice, we learn a mapping from the leaves rather than the trees of  $\mathcal{T}$  for a more granular mapping. Specifically, we construct vectors  $z_i = (z'_{i1}, \dots, z'_{it})$  where each  $z'_{ik}$  is a size  $\ell_k$  vector in which the value at index  $\text{leaf}(T_k, \mathbf{x}_i)$  is set to  $S(T_k, \mathbf{x}_i)$  and other entries to zero. Here,  $\ell_k$  denotes the number of leaves in tree  $T_k$  and  $\text{leaf}(\cdot)$  returns the index of the leaf that  $\mathbf{x}_i$  falls into in the corresponding tree (note that  $\mathbf{x}_i$  belongs to exactly one leaf of any tree, since the trees partition the feature space).

**SPI-LITE: A “light” version.** We note that instead of mimicking each individual fragment of knowledge  $S^*(T_k^*, \mathbf{x}^*)$ ’s, one could also directly mimic the “final decision”  $s^*(\mathbf{x}^*)$ . To this end, we also introduce SPI-LITE, which estimates a *single* regressor  $\phi(z_i) \approx s^*(\mathbf{x}_i^*)$  for  $i = 1, \dots, n$  (also see Figure 1). We compare SPI and SPI-LITE empirically in §4.

**Learning to Rank (L2R) like in  $X^*$ :** An important challenge in learning to accurately mimic the scores  $s^*$ ’s in Eq. (5) is to make sure that the regressors  $\phi_k$ ’s are very accurate in their approximations in Eq. (6). Even then, it is hard to guarantee that the final ranking of points by  $\sum_{k=1}^{t^*} \hat{\phi}_k(z_i)$  would reflect their ranking by  $s^*(\mathbf{x}_i^*)$ . Our ultimate goal, after all, is to *mimic the ranking* of the ensemble in  $X^*$  space since anomaly detection is a ranking problem at its heart.

To this end, we set up an additional pairwise learning to rank objective as follows. Let us denote by  $\phi_i = (\hat{\phi}_1(z_i), \dots, \hat{\phi}_{t^*}(z_i))$  the  $t^*$ -dimensional vector of estimated knowledge fragments for each training example  $i$ . For each pair of training examples, we create a tuple of the form  $((\phi_i, \phi_j), p_{ij}^*)$  where

$$p_{ij}^* = P(s_i^* < s_j^*) = \sigma(-(s_i^* - s_j^*)), \quad (7)$$



**Algorithm 1** SPI-TRAIN: Incorporating PI to Unsupervised Anomaly Detector**Input:** training examples  $\{(\mathbf{x}_1, \mathbf{x}_1^*), \dots, (\mathbf{x}_n, \mathbf{x}_n^*)\}$ **Output:** detection model (ensemble-of-trees)  $\mathcal{T}$  in  $X$  space; regressors  $\hat{\phi}_k$ 's,  $k = 1, \dots, t^*$ ;  $\beta$  (or  $\gamma$  for kernelized L2R)

- 1: Learn  $t^*$  isolation trees  $\mathcal{T}^* = \{T_1^*, \dots, T_{t^*}^*\}$  on  $\mathbf{x}_i^*$ 's  $i = 1, \dots, n$
- 2: Learn  $t$  isolation trees  $\mathcal{T} = \{T_1, \dots, T_t\}$  on  $\mathbf{x}_i$ 's  $i = 1, \dots, n$
- 3: Construct leaf score vectors  $\mathbf{z}_i$ 's,  $i = 1, \dots, n$ , based on  $\mathcal{T}$
- 4: **for each**  $k = 1, \dots, t^*$  **do**
- 5:   Learn regressor  $\hat{\phi}_k$  of  $\mathbf{z}_i$ 's onto  $S^*(T_k^*, \mathbf{x}_i^*)$ 's
- 6:   Obtain  $\beta$  by optimizing  $C$  in (9) (or  $\gamma$  for kernelized  $C_\psi$ )
- 7: **end for**

**Algorithm 2** SPI-TEST: PI-Augmented Unsupervised Anomaly Detection**Input:** test examples  $\{\mathbf{x}_{n+1}, \dots, \mathbf{x}_{n+m}\}$ ;  $\mathcal{T}$ ,  $\hat{\phi}_k$ 's  $k = 1, \dots, t^*$ ,  $\beta$  (or  $\gamma$  if kernelized)**Output:** estimated anomaly scores  $\{s_{n+1}, \dots, s_{n+m}\}$  for all test examples

- 1: **for each** test example  $\mathbf{x}_e$ ,  $e = n+1, \dots, n+m$  **do**
- 2:   Construct leaf score vector  $\mathbf{z}_e = (\mathbf{z}'_{e1}, \dots, \mathbf{z}'_{et})$  where entry in each  $\mathbf{z}'_{ek}$  for index  $\text{leaf}(T_k, \mathbf{x}_e)$  is set to  $S(T_k, \mathbf{x}_e)$  and to 0 o.w., for  $k = 1, \dots, t$
- 3:   Construct  $\phi_e = (\hat{\phi}_1(\mathbf{z}_e), \dots, \hat{\phi}_{t^*}(\mathbf{z}_e))$
- 4:   Estimate anomaly score as  $s_e = \beta \phi_e^T$  (or  $s_e = \sum_{l=1}^n \gamma_l K(\phi_l, \phi_e)$  if kernelized)
- 5: **end for**

which is the probability that  $i$  is ranked ahead of  $j$  by anomalousness in  $X^*$  space (recall that lower  $s^*$  is more anomalous), where  $\sigma(v) = 1/(1 + e^{-v})$  is the sigmoid function. Notice that the larger the gap between the anomaly scores of  $i$  and  $j$ , the larger this probability gets (i.e., more surely  $i$  ranks above  $j$ ).

Given the training pair tuples above, our goal of learning-to-rank is to estimate  $\beta \in \mathbb{R}^{t^*}$ , such that

$$p_{ij} = \sigma(\Delta_{ij}) = \sigma(\beta \phi_i^T - \beta \phi_j^T) = \sigma(-\hat{s}_i^* + \hat{s}_j^*) \approx p_{ij}^*, \quad \forall i, j \in \{1, \dots, n\}. \quad (8)$$

We then utilize the cross entropy as our cost function over all  $(i, j)$  pairs, as

$$\min_{\beta} C = \sum_{(i,j)} -p_{ij}^* \log(p_{ij}) - (1 - p_{ij}^*) \log(1 - p_{ij}) = \sum_{(i,j)} -p_{ij}^* \Delta_{ij} + \log(1 + e^{\Delta_{ij}}) \quad (9)$$

where  $p_{ij}^*$ 's are given as input to the learning as specified in Eq. (7) and  $p_{ij}$  is denoted in Eq. (8) and is parameterized by  $\beta$  that is to be estimated.

The objective function in (9) is convex and can be solved via a gradient-based optimization, where  $\frac{dC}{d\beta} = \sum_{(i,j)} (p_{ij} - p_{ij}^*) (\phi_i - \phi_j)$  (details omitted for brevity). More importantly, in case the linear mapping  $s_i^* \approx \beta \phi_i^T$  is not sufficiently accurate to capture the desired pairwise rankings, the objective can be *kernelized* to learn a *non-linear* mapping that is likely more accurate. The idea is to write  $\beta_\psi = \sum_{l=1}^n \gamma_l \psi(\phi_l)$  (in the transformed space) as a weighted linear combination of (transformed) training examples, for feature transformation function  $\psi(\cdot)$  and parameter vector  $\gamma \in \mathbb{R}^n$  to be estimated. Then,  $\Delta_{ij}$  in objective (9) in the transformed space can be written as

$$\Delta_{ij} = \sum_{l=1}^n \gamma_l [\psi(\phi_l) \psi(\phi_i)^T - \psi(\phi_l) \psi(\phi_j)^T] = \sum_{l=1}^n \gamma_l [K(\phi_l, \phi_i) - K(\phi_l, \phi_j)]. \quad (10)$$

The kernelized objective, denoted  $C_\psi$ , can also be solved through gradient-based optimization where we can show partial derivatives (w.r.t. each  $\gamma_l$ ) to be equal to  $\frac{\partial C_\psi}{\partial \gamma_l} = \sum_{(i,j)} (p_{ij} - p_{ij}^*) [K(\phi_l, \phi_i) - K(\phi_l, \phi_j)]$ . Given the estimated  $\gamma_l$ 's, prediction of score is done by  $\sum_{l=1}^n \gamma_l K(\phi_l, \phi_e)$  for any (test) example  $e$ .

**The SPI Algorithm:** We outline the steps of SPI for both training and testing (i.e., detection) in Algo. 1 and Algo. 2, respectively. Note that the test-time detection no longer relies on the availability of privileged features for the test examples, but yet be able to leverage/incorporate them through its training.

## 4 Experiments

We design experiments to evaluate our methods in two different settings:

1. **Benchmark Evaluation:** We show the effectiveness of augmenting PI (see Table 3) on 17 publicly available benchmark datasets.<sup>2</sup>
2. **Real-world Use Cases:** We conduct experiments on LingSpam<sup>3</sup> and BotOrNot<sup>4</sup> datasets to show that (i) domain-expert knowledge as PI improves spam detection, (ii) compute-expensive PI enables fast detection at test time, and (iii) “historical future” PI allows early detection of bots.

**Baselines** We compare both SPI and SPI-LITE to the following baselines:

1. IF( $X$ -only): Isolation Forest [11] serves as a simple baseline that operates solely in decision space  $X$ . PI is not used neither for modeling nor detection.
2. OC-SVM+ (PI-incorporated): OC+ for short, is an extension of (unsupervised) One-Class SVM that incorporates PI as introduced in [2].
3. FT(PI-incorporated): This is the direct feature transfer method that incorporates PI by learning a mapping  $X \rightarrow X^*$  as we introduced in §3.1.
- \* IF\* ( $X^*$ -only): IF that operates in  $X^*$  space. We report performance by IF\* only for reference, since PI is unavailable at test time.

### 4.1 Benchmark Evaluation

The benchmark datasets do not have an explicit PI representation. Therefore, in our experiments we introduce PI as explained below.

**Generating privileged representation.** For each dataset, we introduce PI by perturbing normal observations. We designate a small random fraction (= 0.1) of  $n$  normal data points as anomalies. Then, we randomly select a subset of  $p$  attributes and add zero-mean Gaussian noise to the designated anomalies along the selected subset of attributes with matching variances of the selected features. The  $p$  selected features represent PI since anomalies stand-out in this subspace due to added noise, while the rest of the  $d$  attributes represent  $X$  space. Using normal observations allows us to control for features that could be used as PI. Thus we discard the actual anomalies from these datasets where PI is unknown.

We construct 4 versions per dataset with varying fraction  $\gamma$  of perturbed features (PI) retained in  $X^*$  space. In particular, each set has  $\gamma p$  features in  $X^*$ , and  $(1 - \gamma)p + d$  features in  $X$  for  $\gamma \in \{0.9, 0.7, 0.5, 0.3\}$ .

<sup>2</sup> <http://agents.fel.cvut.cz/stegodata/Loda.zip>

<sup>3</sup> <http://csmining.org/index.php/ling-spam-datasets.html>

<sup>4</sup> <https://botometer.iuni.iu.edu/bot-repository/datasets/caverlee-2011/caverlee-2011.zip>

Table 3: Mean Average Precision (MAP) on benchmark datasets (avg’ed over 5 runs) for  $\gamma = 0.7$ . Numbers in parentheses indicate rank of each algorithm on each dataset.  $iF^*$  (for reference only) reports MAP in the  $X^*$  space.

Datasets	$p+d$	$n$	$iF$	OC+	FT	SPI-LITE	SPI	$iF^*$
breast-cancer	30	357	0.1279 (4)	0.0935 (6)	0.0974 (5)	0.4574 (3)	<b>0.5746</b> (2)	0.6773 (1)
ionosphere	33	225	0.0519 (4)	<b>0.2914</b> (1)	0.0590 (3)	0.0512 (5)	0.0470 (6)	0.0905 (2)
letter-recognition	617	4197	0.0889 (6)	0.1473 (4)	0.0908 (5)	0.3799 (3)	<b>0.6413</b> (2)	0.9662 (1)
multiple-features	649	1200	0.1609 (5)	0.1271 (6)	0.2044 (4)	0.6589 (3)	<b>0.8548</b> (2)	1.0000 (1)
wall-following-robot	24	2923	0.1946 (5)	0.2172 (4)	0.1848 (6)	0.4331 (3)	<b>0.5987</b> (2)	0.7538 (1)
cardiotocography	27	1831	0.2669 (5)	0.6107 (4)	0.2552 (6)	0.6609 (3)	<b>0.6946</b> (2)	0.8081 (1)
isolet	617	4497	0.1533 (5)	0.1561 (4)	0.1303 (6)	0.5084 (3)	<b>0.7124</b> (2)	0.9691 (1)
libras	90	216	0.1368 (5)	0.4479 (4)	0.0585 (6)	0.5175 (3)	<b>0.6806</b> (2)	1.0000 (1)
parkinsons	22	147	0.0701 (6)	0.0964 (4)	0.0714 (5)	0.1556 (3)	<b>0.1976</b> (1)	0.1778 (2)
statlog-satimage	36	3594	0.2108 (6)	0.5347 (5)	0.5804 (4)	0.9167 (3)	<b>0.9480</b> (2)	0.9942 (1)
gisette	4971	3500	0.1231 (4)	0.0814 (6)	0.0977 (5)	0.5593 (3)	<b>0.8769</b> (2)	0.9997 (1)
waveform-1	21	3304	0.1322 (4)	0.1481 (3)	0.0841 (6)	0.1234 (5)	<b>0.1556</b> (2)	0.4877 (1)
madelon	500	1300	0.7562 (5)	0.1167 (6)	<b>0.9973</b> (2)	0.9233 (4)	0.9925 (3)	1.0000 (1)
synthetic-control	60	400	0.3207 (6)	0.7889 (4)	0.6870 (5)	0.8103 (3)	<b>0.8539</b> (2)	0.9889 (1)
waveform-2	21	3304	0.1271 (5)	<b>0.2828</b> (2)	0.1014 (6)	0.1778 (3)	0.1772 (4)	0.2944 (1)
statlog-vehicle	18	629	0.1137 (6)	0.3146 (5)	0.6326 (4)	0.6561 (3)	<b>0.7336</b> (2)	1.0000 (1)
statlog-segment	18	1320	0.1250 (6)	0.2323 (4)	0.1868 (5)	0.3304 (3)	<b>0.3875</b> (2)	0.7399 (1)
(Average Rank)			(5.11)	(4.23)	(4.88)	(3.29)	(2.35)	(1.11)

**Results** We report the results on perturbed datasets with  $\gamma = 0.7^5$  as fraction of features retained in space  $X^*$ . Table 3 reports mean Average Precision (area under the precision-recall curve) against 17 datasets for different methods. The results are averaged across 5 independent runs on stratified train-test splits.

Our SPI outperforms competition in detection performance in most of the datasets. To compare the methods statistically, we use the non-parametric Friedman test [5] based on the average ranks. Table 3 reports the ranks (in parentheses) on each dataset as well as the average ranks. With  $p$ -value =  $2.16 \times 10^{-11}$ , we reject the null hypothesis that all the methods are equivalent using Friedman test. We proceed with Nemenyi post-hoc test to compare the algorithms pairwise and to find out the ones that differ significantly. The test identifies performance of two algorithms to be significantly different if their average ranks differ by at least the ‘‘critical difference’’ (CD). In our case, comparing 6 methods on 17 datasets at significance level  $\alpha = 0.05$ ,  $CD = 1.82$ .




Fig. 2: Average rank of algorithms (w.r.t. MAP) and comparison by the Nemenyi test. Groups of methods not significantly different (at  $p$ -val = 0.05) are connected with horizontal lines. CD depicts critical distance required to reject equivalence. Note that SPI is significantly better than the baselines.

<sup>5</sup> The results with  $\gamma \in \{0.9, 0.5, 0.3\}$  are similar and reported in the supplementary material available at <http://www.andrew.cmu.edu/user/shubhras/SPI>

Results of the post-hoc test are summarized through a graphical representation in Figure 2. We find that SPI is significantly better than all the baselines. We also notice that SPI has no significant difference from  $iF^*$  which uses PI at test time, demonstrating its effectiveness in augmenting PI. While all the baselines are comparable to SPI-LITE, its average rank is better (also see last row in Table 3), followed by other PI-incorporated detectors, and lastly  $iF$  with no PI.

Average Precision (AP) is a widely-accepted metric to quantify overall performance of ranking methods like anomaly detectors. We also report average rank of the algorithms against other popular metrics including AUC of ROC curve, NDCG@10 and PRECISION@10 in Figure 3. Notice that the results are consistent across measures, SPI and SPI-LITE performing among the best.




Fig. 3: SPI and SPI-LITE outperform competition w.r.t. different evaluation metrics. Average rank (bars) across benchmark datasets.  $iF^*$  shown for reference.

## 4.2 Real-world Use Cases

**Data description.** LingSpam dataset<sup>3</sup> consists of 2412 non-spam and 481 spam email messages from a linguistics mailing-list. We evaluate two use cases (1) domain-expert knowledge as PI and (2) compute-expensive PI on LingSpam.

BotOrNot dataset<sup>4</sup> is collected from Twitter during December 30, 2009 to August 2, 2010. It contains 22,223 content polluters (bots) and 19,276 legitimate users, along with their number of followings over time and tweets. For our experiments, we select accounts with age less than 10 days (for early detection task) at the beginning of dataset collection. The subset contains 901 legitimate (human) accounts and 4535 bots. We create 10 sets containing all the legitimate and a random 10% sample of the bots. We evaluate use case (3) “historical future” as PI and report the results averaged over these sets.

### Case 1: Domain-expert Knowledge as PI for Email Spam Detection.

**$X^*$  space:** The Linguistic Inquiry and Word Count (LIWC) software<sup>6</sup> is a widely used text analysis tool in social sciences. It uses a manually-curated keyword dictionary to categorize text into 90 psycholinguistic classes. Construction of LIWC dictionary relies exclusively on human experts which is a slow and evolving process. For the LingSpam dataset, we use the percentage of word counts in each class (assigned by LIWC software) as the privileged features.

**$X$  space:** The bag-of-words model is widely used as feature representation in text analysis. As such, we use the term frequencies for our email corpus as the primary features.

Figure 4 shows the detection performance<sup>7</sup> of algorithms in ROC curves

<sup>6</sup> <https://liwc.wpengine.com/>

<sup>7</sup> See [supplementary material](#) quantifying the performance of methods against other ranking metrics

(averaged over 15 independent runs on stratified train-test splits). We find that  $iF$ , which does not leverage PI but operates solely in  $X$  space, is significantly worse than most PI-incorporated methods. OC-SVM+ is nearly as poor as  $iF$  despite using PI—this is potentially due to OC-SVM being a poor anomaly detector in the first place, as shown in [6] and as we argued in §1. All knowledge transfer methods, SPI, SPI-LITE, and FT, perform similarly on this case study, and are as good as  $iF^*$ , directly using  $X^*$ .




Fig. 4: Detection performance on Case 1: using expert knowledge as PI. Legend depicts the AUC values. PI-incorporated detectors (except OC-SVM+) outperform non-PI  $iF$  and achieve similar performance to  $iF^*$ .

### Case 2: Compute-Expensive Features as PI for Email Spam Detection.

**$X^*$  space:** Beyond bag-of-words, one can use syntactic features to capture *stylistic* differences between spam and non-spam emails. To this end, we extract features from the parse trees of emails using the StanfordParser<sup>8</sup>. The parser provides the taxonomy (tree) of Part-of-Speech (PoS) tags for each sentence, based on which we construct (i) PoS bi-gram frequencies, and (ii) quantitative features (width, height, and horizontal/vertical imbalance) of the parse tree.

On average, StanfordParser requires 66 seconds<sup>9</sup> to parse and extract features from a single raw email in LingSpam. Since the features are computationally demanding, we incorporate those as PI to facilitate faster detection at test time.

**$X$  space:** We use the term frequencies as the primary features as in Case 1.

Figure 5 (a) shows the detection performance<sup>7</sup> of methods in terms of AUC under ROC. We find that  $iF^*$  using (privileged) syntactic features achieves lower AUC of  $\sim 0.65$  as compared to  $\sim 0.83$  using (privileged) LIWC features in Case 1. Accordingly, all methods perform relatively lower, suggesting that the syntactic features are less informative of spam than psycholinguistic ones. Nonetheless, we observe that the performance ordering remains consistent, where  $iF$  ranks at the bottom and SPI and SPI-LITE get closest to  $iF^*$ .

Figure 5 (b) shows the comparison of wall-clock time required by each detector to compute the anomaly scores at test time for varying fraction of test data. On average, SPI achieves  $5500\times$  speed-up over  $iF^*$  that employs the parser at test time. This is a considerable improvement of response time for comparable accuracy. Also notice the inset plot showing the AUC vs. total test time, where our proposed SPI and SPI-LITE are closest to the ideal point at the top left.

### Case 3: “Historical Future” as PI for Twitter Bot Detection.

We use temporal data from the activity and network evolution of an account to capture behavioral differences between a human and a bot. We construct temporal features including volume, rate-of-change, and lag-autocorrelations of

<sup>8</sup> <https://nlp.stanford.edu/software/lex-parser.shtml>

<sup>9</sup> Using a single thread on 2.2 GHz Intel Core i7 CPU with 8 cores and 16GB RAM

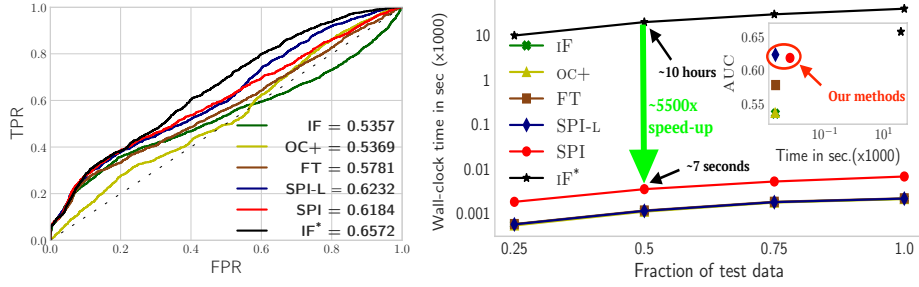


Fig. 5: Comparison of detectors on Case 2: using computationally-expensive features as PI. (a) detection performance, legend depicts AUC values; and (b) wall-clock time required (in seconds, note the logarithmic scale) vs. test data size [inset plot on top right: AUC vs. time (methods depicted with symbols)].

the number of followings. We also extract temporal features from text such as count of tweets, links, hash-tags and mentions.

**X\* space:** All the temporal features within  $f_t$  days in the future (relative to detection at time  $t$ ) constitute privileged features. Such future values would not be available at any test time point but can be found in historical data.

**X space:** Temporal features within  $h_t$  days in the past as well as static user features (from screen name and profile description) constitute primary features.

Figure 6 (a) reports the detection performance of algorithms in terms of ROC curves (averaged over 10 sets) at time  $t = 2$  days after the data collection started; for  $h_t = 2$ ,  $f_t = 7$ .<sup>10</sup> The findings are similar to other cases: SPI and SPI-LITE outperform the competing methods in terms of AUC and OC-SVM+ performs similar to non-PI iF; demonstrating that knowledge transfer based methods are more suitable for real-world use cases.

Figure 6 (b) compares the detection performance of SPI and iF over time; for detection at  $t = \{0, 1, 2, 3, 4\}$ . As time passes, historical data grows as  $h_t = \{0, 1, 2, 3, 4\}$  where “historical future” data is fixed at  $f_t = 7$  for PI-incorporated methods. Notice that at time  $t = 1$ , SPI achieves similar detection performance to iF’s performance at  $t = 2$  that uses more historical data of 2 days. As such, SPI enables 24 hours early detection as compared to non-PI iF for the same accuracy. Notice that with the increase in historical data, the performances of both methods improve, as expected. At the same time, that of SPI improves faster, ultimately reaching a higher saturation level, specifically  $\sim 7\%$  higher relative to iF. Moreover, SPI gets close to iF\*’s level in just around 3 days.

## 5 Related Work

We review the history of LUPI, follow up and related work on learning with side/hidden information, as well as LUPI-based anomaly detection.

**Learning Under Privileged Information:** The LUPI paradigm is introduced by Vapnik and Vashist [19] as the SVM+ method, where, Teacher provides Student not only with (training) examples but also explanations, comparisons,

<sup>10</sup> Same conclusions can be drawn for  $f_t \in \{1, 3, 5, 7\}$  (see [supplementary material](#)).

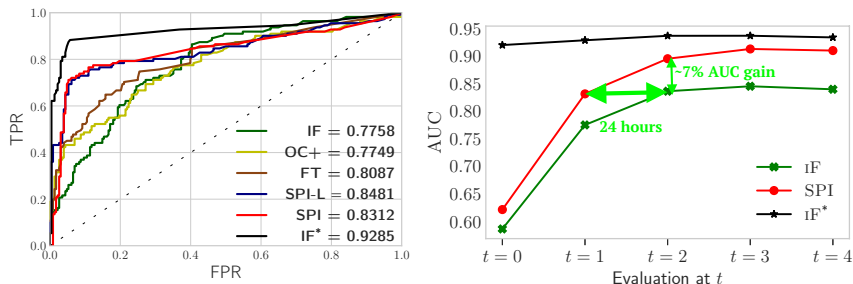


Fig. 6: Comparison of detectors on Case 3: using “historical future” data as PI. (a) SPI outperforms competition in performance and is closest to  $iF^*$ ’s; (b) SPI achieves same detection performance as  $iF$  24 hours earlier, and gets close to  $iF^*$  in 3 days of history.

metaphors, etc. which accelerate the learning process. Roughly speaking, PI adjusts Student’s concept of similarity between training examples and reduces the amount of data required for learning. Lapin et al. [10] showed that learning with PI is a particular instance of importance weighting in SVMs. Another such mechanism was introduced more recently by Vapnik and Izmailov [17], where knowledge is transferred from the space of PI to the space where the decision function is built. The general idea is to specify a small number of fundamental concepts of knowledge in the privileged space and then try to transfer them; i.e., construct additional features in decision space via e.g., regression techniques in decision space. Importantly, the knowledge transfer mechanism is not restricted to SVMs, but generalizes, e.g. to neural networks [18].

LUPI has been applied to a number of different settings including clustering [7,12], metric learning [8], learning to rank [15], malware and bot detection [2,3], risk modeling [14], as well as recognizing objects [16], actions and events [13].

*Learning with Side/Hidden Information:* Several other work, particularly in computer vision [4,20], propose methods to learn with data that is unavailable at test time referred as side and hidden information (e.g., text descriptions or tags for general images, facial expression annotations for face images, etc.). In addition, Jonschkowski et al. [9] describe various patterns of learning with side information. All of these work focus on supervised learning problems.

**LUPI-based Anomaly Detection:** With the exception of One-Class SVM (OC-SVM+) [2], which is a direct extension of Vapnik’s (supervised) SVM+, the LUPI framework has been utilized only for supervised learning problems. While anomaly detection has been studied extensively [1], we are unaware of any work other than [2] leveraging privileged information for unsupervised anomaly detection. Motivated by this along with the premises of the LUPI paradigm, we are the first to design a new technique that ties LUPI with unsupervised tree-based ensemble methods, which are considered state-of-the-art for anomaly detection.

## 6 Conclusion

We introduced SPI, a new ensemble approach that leverages privileged information (data available only for training examples) for unsupervised anomaly detection. Our work builds on the LUPI paradigm, and to the best of our



knowledge, is the first attempt to incorporating PI to improve the state-of-the-art ensemble detectors. We validated the effectiveness of our method on both benchmark datasets as well as three real-world case studies. We showed that SPI and SPI-LITE consistently outperform the baselines. Our case studies leveraged a variety of privileged information—“historical future”, complex features, expert knowledge—and verified that SPI can unlock multiple benefits for anomaly detection in terms of detection latency, speed, as well as accuracy.

**Acknowledgments** This research is sponsored by NSF CAREER 1452425 and IIS 1408287. Any conclusions expressed in this material are of the authors and do not necessarily reflect the views, expressed or implied, of the funding parties.

## References

1. C. C. Aggarwal. *Outlier Analysis*. Springer, 2013.
2. E. Burnaev and D. Smolyakov. One-class SVM with privileged information and its application to malware detection. In *ICDM Workshops*, pages 273–280, 2016.
3. Z. B. Celik, P. McDaniel, R. Izmailov, N. Papernot, and A. Swami. Extending detection with forensic information. *arXiv:1603.09638*, 2016.
4. J. Chen, X. Liu, and S. Lyu. Boosting with side information. In *ACCV*, 2012.
5. J. Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7(Jan):1–30, 2006.
6. A. Emmott, S. Das, T. Dietterich, A. Fern, and W.-K. Wong. Systematic construction of anomaly detection benchmarks from real data. In *KDD ODD*, 2013.
7. J. Feyereisl and U. Aickelin. Privileged information for data clustering. *Inf. Sci.*, 194:4–23, 2012.
8. S. Fouad, P. Tino, S. Raychaudhury, and P. Schneider. Incorporating privileged information through metric learning. *IEEE Neural Net. Learning Sys.*, 24(7), 2013.
9. R. Jonschkowski, S. Höfer, and O. Brock. Patterns for learning with side information. *arXiv:1511.06429*, 2015.
10. M. Lapin, M. Hein, and B. Schiele. Learning using privileged information: Svm+ and weighted svm. *Neural Networks*, 53:95–108, 2014.
11. F. T. Liu, K. M. Ting, and Z.-H. Zhou. Isolation forest. In *ICDM*, 2008.
12. R. M. Marcacini, M. A. Domingues, E. R. Hruschka, and S. O. Rezende. Privileged information for hierarchical document clustering: A metric learning approach. In *ICPR*, pages 3636–3641, 2014.
13. L. Niu, W. Li, and D. Xu. Exploiting privileged information from web data for action and event recognition. *Intern. J. of Comp. Vision*, 118(2):130–150, 2016.
14. B. Ribeiro, C. Silva, N. Chen, A. Vieira, and J. C. das Neves. Enhanced default risk models with SVM+. *Expert Syst. Appl.*, 39(11):10140–10152, 2012.
15. V. Sharmanska, N. Quadrianto, and C. H. Lampert. Learning to rank using privileged information. In *ICCV*, pages 825–832, 2013.
16. V. Sharmanska, N. Quadrianto, and C. H. Lampert. Learning to transfer privileged information. *arXiv:1410.0389*, 2014.
17. V. Vapnik and R. Izmailov. Learning with intelligent teacher: Similarity control and knowledge transfer. In *Stat. Learning and Data Sci.*, pages 3–32, 2015.
18. V. Vapnik and R. Izmailov. Knowledge transfer in svm and neural networks. *Ann. Math. Artif. Intell.*, 81(1-2):3–19, 2017.
19. V. Vapnik and A. Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 22(5-6):544–557, 2009.
20. Z. Wang and Q. Ji. Classifier learning with hidden information. In *CVPR*, 2015.