

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ZERO-SHOT OUTLIER DETECTION VIA SYNTHETICALLY PRE- TRAINED TRANSFORMERS: MODEL SELECTION BYGONE!

Anonymous authors

Paper under double-blind review

ABSTRACT

Outlier detection (OD) has a vast literature as it finds numerous applications in environmental monitoring, security, manufacturing, and finance to name a few. Being an inherently *unsupervised* task, model selection is a key bottleneck for OD (both algorithm and hyperparameter selection) without label supervision. There is a long list of techniques to choose from – both classical algorithms and deep neural architectures – and while several studies report their hyperparameter sensitivity, the literature remains quite slim on unsupervised model selection—limiting the effective use of OD in practice. In this paper we present FoMo-OD, for zero/0-shot OD exploring a transformative new direction that *bypasses* the hurdle of model selection altogether (!), thus breaking new ground. The fundamental idea behind FoMo-OD is the Prior-data Fitted Networks, recently introduced by Müller et al. (2022), which trains a Transformer model on a large body of *synthetically* generated data from a prior data distribution. In essence, FoMo-OD is a **pretrained Foundation Model for zero/0-shot OD on tabular data**, which can directly predict the (outlier/inlier) label of any test data at inference time, by merely a *single forward pass*—making obsolete the need for choosing an algorithm/architecture and tuning its associated hyperparameters, besides requiring no training of model parameters when given a new OD dataset. Extensive experiments on **57** public benchmark datasets against **26** baseline methods show that FoMo-OD performs statistically no different from the **2nd** top baseline, while significantly outperforming the majority of the baselines, with an average inference time of **7.7 ms** per test sample.

1 INTRODUCTION

Outlier detection (OD) finds applications in many domains such as security, environmental monitoring, finance, and so on. This popularity brings along a large literature that offers a plethora of detection algorithms to choose from given a new OD task. These techniques, however, exhibit several hyperparameters (HPs) that need careful tuning to which they are often quite sensitive (Ma et al., 2023). What makes it notoriously difficult to achieve effective OD performance in practice is *model selection* (both algorithm and HPs) in the *absence of any labels*, as most tasks are unsupervised.¹

In fact, while deep learning and modern architectures have revolutionized many areas of machine learning (ML), it has not quite been the case for OD—mainly because deep OD models (Pang et al., 2021) exhibit many more HPs (for architecture, regularization, and optimization) that detection performance is sensitive to (Ding et al., 2022), as compared to classical methods with only a few HPs.

Large foundation models have stirred up most recent advances in ML, which are (pre-)trained on massive amounts of data. The most notable progress has been in natural languages and vision, thanks to the admirable quantity and quality of public text and image datasets. In contrast, public (benchmark) datasets for OD is minuscule in comparison (Han et al., 2022; Zhao et al., 2021; Steinbuss and Böhm, 2021). Another obstacle for foundation models for tabular OD has been the non-shared feature spaces of different datasets, unlike the shared pixel or word spaces for images and text.

Recently, the introduction of Prior-data Fitted Networks (PFNs) has marked a milestone as a new approach to ML on tabular data (Müller et al., 2022). PFNs are based on Bayesian non-parametrics

¹While semi-/supervised settings of OD exist, unsupervised OD is preferable in most domains for the capacity to detect novel/emergent types of anomalies, beyond just the known types.

Table 1: Comparison of methods across datasets. (top row) Rank w.r.t. AUROC performance avg.’ed over 57 datasets is presented for FoMo-0D (with $D = 100$), **top-10 baselines** with default HPs, and **top-4⁵** baselines with performance **avg.**’ed over varying HPs (denoted w/ ^{avg}); followed by p -values of the pairwise Wilcoxon signed rank test, comparing FoMo-0D to each baseline (from top to bottom) over All (57) datasets, those (42) w/ $d \leq 100$ and (46) w/ $d \leq 500$ dimensions. FoMo-0D performs as well as (**i.e., statistically no different from**) the **2nd best model** (kNN , w/ $p = 0.106$) across All datasets, while it is **comparable to** ($p > 0.05$) **or better than** ($p > 0.95$) **all baselines** over datasets w/ $d \leq 100$ (aligned w/ pretraining where $D = 100$) *and* $d \leq 500$ (generalizing beyond pretraining).

	FoMo-0D	DTE-NP	kNN	ICL	DTE-C	LOF	CBLOF	Feat.Bag.	SLAD	DDPM	OCSVM	DTE-NP ^{avg}	kNN^{avg}	ICL ^{avg}	DTE-C ^{avg}		
Rank(avg)	11.886		7.553	9.018	10.851	11.36	12.316	13.342	13.386	12.982	14.061	13.851		9.079	11.105	12.991	22.263
All	-		0.016	0.106	0.462	0.454	0.585	0.750	0.823	0.759	0.901	0.895		0.112	0.315	0.670	1.000
$d \leq 100$	-		0.415	0.700	0.949	0.953	0.970	0.971	0.996	0.876	0.980	0.978		0.752	0.860	0.958	1.000
$d \leq 500$	-		0.220	0.569	0.827	0.894	0.960	0.968	0.994	0.910	0.960	0.979		0.607	0.756	0.846	1.000

and meta-learning on large quantities of *synthetically* simulated data from a data prior. The key idea is to compute a posterior predictive distribution (PPD) for a test point given the training data as input context. To approximate the PPD, a Transformer (Vaswani et al., 2017) is pre-trained to mimic the PPD via simulating numerous training datasets from a (general, complex) data prior. For inference, the fresh training set along with the test samples are passed to the (frozen) pre-trained PFN, which outputs the predictions in a *single forward pass*, requiring no model training or model selection. Variants of PFN are shown to match the performance of tree-based models on small classification datasets (Hollmann et al., 2023) and in time series forecasting with limited data (Dooley et al., 2023).

In this paper, we capitalize on these ideas and introduce FoMo-0D; a prior-data fitted Foundation Model for zero- or Q-shot Outlier Detection (for the “Fear of Missing out”-liers). The implication and “gift” of PFNs for unsupervised OD goes beyond those for supervised learning: it helps *bypass* not only model (parameter) training, but most importantly, the notoriously-hard task of model (hyperparameter) selection altogether. As such, FoMo-0D unlocks zero-shot OD on a new dataset without the need for any algorithm or HP selection. During inference, data is used only as input *context* to FoMo-0D, and *not* for parameter training or HP tuning. Arguably, this is a potential game changer for unsupervised OD, especially for practitioners. Figure 1 illustrates the new FoMo-0D paradigm versus the typical OD setting.

In designing FoMo-0D, we simply use Gaussian mixture models as a simple yet effective tabular data prior, to capture general and diverse inlier data distributions, following current literature (Hollmann et al., 2023; Zhao et al., 2021). We combine these with simulated outlier types common in the real-world; namely local and global subspace outliers (Steinbuss and Böhm, 2021). While the data prior can be extended to comprise more complex data distributions (e.g. through the use of Bayesian Neural Networks (BNNs; (Neal, 2012)) and Structural Causal Models (SCMs; (Pearl, 2009)) as in (Hollmann et al., 2023)), and additional outlier types can be included (e.g. dependency, contextual, etc. outliers), as we show in the experiments, even with the relatively straightforward prior that we employed, FoMo-0D achieves remarkable performance. As shown in Table 1, FoMo-0D pretrained on synthetic datasets with up to 100 dimensions performs *statistically no different* from all 26 state-of-the-art baselines (all p -values > 0.2) on 46 benchmark datasets with dimensionality $d \leq 500$, while *significantly outperforming the majority* of the baselines (with $p > 0.95$) (see Appendix Tables 12.1&12.2). Further, FoMo-0D takes a mere average of 7.7 ms to infer a test sample since a new dataset requires a single forward pass for inference and no training overhead.

Our contributions: We summarize the main contributions of our work as follows.

- **A Foundation Model for Tabular OD:** We present FoMo-0D, *the first foundation model for zero-shot OD* on tabular datasets. FoMo-0D is a Prior-data Fitted Network (PFN) (Müller et al., 2022) that is pretrained on many synthetically generated datasets drawn from a novel data prior that we introduce to capture various inlier and outlier distributions. The pretrained FoMo-0D can then directly compute the posterior predictive distribution (PPD) of test points in a new dataset.
- **Unsupervised Outlier Model Selection Made Obsolete:** The most outstanding property of FoMo-0D is its *zero-shot inference* on a new dataset via a single forward pass, fully abolishing the need not only for model training on a new dataset, but importantly also the notorious task of algorithm selection and hyperparameter tuning in the absence of labeled data.
- **Scalable Pre-training Design:** To unlock the premise of large-scale pretraining on numerous large datasets, (1) we implement a new mechanism to speed up sample-to-sample attention from

108 quadratic to *linear time* complexity—enabling *larger datasets*; and (2) we scale up on-the-fly data
 109 synthesis through data transformation—enabling *more datasets* in less time.

- 110 • **Fast Inference at Detection Time:** Thanks to a pretrained prior-data fitted Transformer, FoMo-
 111 OD bypasses both model (parameter) training and selection, both of which can be slow for modern
 112 deep OD models with many hyper/parameters. Rather, it takes *fraction of a second* to label a
 113 test point through a single forward pass that can be parallelized across test samples. Such speedy
 114 inference also unlocks the potential for deploying FoMo-OD in *real time* on data streams.
- 115 • **Effectiveness:** We evaluate FoMo-OD on **57** public benchmark datasets (Han et al., 2022) from
 116 diverse domains and compare against **26** baselines from classical to modern (Livernoche et al.,
 117 2024), where FoMo-OD significantly outperforms the majority of the baselines while performing
 118 statistically no different from the top *2nd* baseline, at the fraction of the compute cost.

119 As FoMo-OD proposes a paradigm shift for OD, abolishing model training and selection altogether,
 120 while delivering unreasonable effectiveness on benchmark datasets even with a basic data prior, we
 121 expect FoMo-OD will trigger further work in both research and practice. To this end, we make all of
 122 our codebase for synthetic data generation, model training, and our pretrained FoMo-OD checkpoints,
 123 openly available at <https://anonymous.4open.science/r/PFN40D>.

125 2 PROBLEM AND PRELIMINARIES

127 2.1 SEMI-SUPERVISED OUTLIER DETECTION

129 Outlier detection (OD) methods can be categorized based on the availability of labeled data. In
 130 supervised OD, the task is similar to binary classification with imbalanced classes (as outliers typically
 131 make up only a small portion of the overall data). The more difficult unsupervised setting assumes
 132 the “contaminated” training data contains both inliers and outliers, but without any labels. A semi-
 133 supervised or one-class classification approach lies between these two extremes, where only inlier
 134 data is available for training, but unknown outliers may appear during inference. Semi-supervised
 135 OD is used in practice where it is easy to gather inlier data, but learning from known, labeled outliers
 136 is undesirable because outliers are hard to collect and/or new, unknown outlier types are likely to
 137 arise in future test data that renders learning only from the known outliers suboptimal/risky.

138 Note that semi-supervised OD may be a *misnomer* from the supervised ML perspective, where
 139 semi-supervised classification assumes the presence of some labeled instances from **all** classes in
 140 the training data. As such, model selection continues to be as difficult for semi-supervised OD as
 141 unsupervised OD, where no labeled outliers exist in the input/training data in both settings.

142 We focus on semi-supervised OD. Formally, let $\mathcal{D}_{\text{in}} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ denote the input data
 143 containing only inliers $\mathbf{x}_i \in \mathbb{R}^d$, where $y_i = 0 \forall i \in [n]$, and $\mathcal{D}_{\text{test}}$ depicts the test data comprising
 144 both inliers and outliers. The task is to assign labels to $\mathbf{x}_i \in \mathcal{D}_{\text{test}}$ given the inlier-only input \mathcal{D}_{in} .

146 2.2 BACKGROUND ON PRIOR-DATA FITTED NETWORKS

148 **Posterior Predictive Distribution (PPD):** In the Bayesian framework for supervised learning, the
 149 prior defines a hypotheses space Φ which expresses our beliefs about the data distribution before
 150 seeing any data. Each hypothesis $\phi \in \Phi$ describes a mechanism by which the data is generated. The
 151 posterior predictive distribution $p(\cdot | \mathbf{x}_{\text{test}}, \mathcal{D}_{\text{train}})$ provides a framework for making prediction on
 152 new, unseen test data \mathbf{x}_{test} , conditioned on observed training data $\mathcal{D}_{\text{train}} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$.
 153 Based on Bayes’ Theorem, the PPD can be derived by the integration over the space of hypotheses Φ :

$$154 p(y_{\text{test}} | \mathbf{x}_{\text{test}}, \mathcal{D}_{\text{train}}) = \int_{\Phi} p(y_{\text{test}} | \mathbf{x}_{\text{test}}, \phi) p(\mathcal{D}_{\text{train}} | \phi) p(\phi) d\phi, \quad (1)$$

156 where $p(\phi)$ denotes the prior probability and $p(\mathcal{D} | \phi)$ is the likelihood of the data \mathcal{D} given ϕ .

158 **PFNs and PPD Approximation:** As obtaining the above PPD is generally intractable, Prior-data
 159 Fitted Networks (PFNs) are proposed to approximate the PPD (Müller et al., 2022). Unlike traditional
 160 machine learning models that are trained directly on observed datasets, PFNs are pre-trained offline
 161 on simulated datasets that are generated according to a prior distribution. Specifically, it contains the
 pre-training and inference stages described as the following.

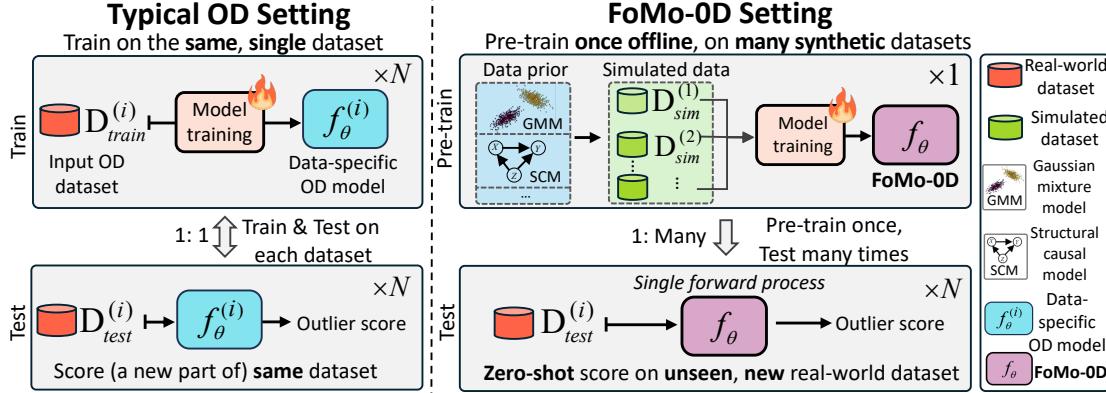


Figure 1: (best in color) Comparison of typical OD vs. the FoMo-OD settings. Given a new un/semi-supervised OD dataset, FoMo-OD not only eliminates the need for model training, but most importantly, also abolishes the onerous task of model selection (algorithm and hyperparameters) w/out labels.

Pre-training on synthetic data. At the beginning of the pre-training stage, massive synthetic training datasets are generated, by first sampling a hypothesis (i.e., the generating mechanism) $\phi \sim p(\phi)$, and then sampling a dataset $\mathcal{D} \sim p(\mathcal{D}|\phi)$. For training purposes, each dataset \mathcal{D} can be split as $\mathcal{D}_{test} \subset \mathcal{D}$ and $\mathcal{D}_{train} = \mathcal{D} \setminus \mathcal{D}_{test}$. Thus the PFN with parameters θ can be optimized by making predictions on data points in \mathcal{D}_{test} . For a test point $(\mathbf{x}_{test}, y_{test}) \in \mathcal{D}_{test}$, the training loss is formulated as

$$\mathcal{L} = \mathbb{E}_{(\{\mathbf{x}_{test}, y_{test}\}) \cup \mathcal{D}_{train} \sim p(\mathcal{D})} [-\log q_\theta(y_{test} | \mathbf{x}_{test}, \mathcal{D}_{train})]. \quad (2)$$

The above loss can also be interpreted as minimizing the expected KL divergence between $p(\cdot | \mathbf{x}, \mathcal{D})$ and $q_\theta(\cdot | \mathbf{x}, \mathcal{D})$ (Müller et al., 2022). In practice, a PFN model q_θ is typically implemented by a Transformer-based architecture (Vaswani et al., 2017), which takes $(\mathbf{x}_{test}, \mathcal{D}_{train})$ as input, where $\mathbf{x}_{test} \in \mathcal{D}_{test}$ and \mathcal{D}_{train} contains an arbitrary number of instances. The output is the conditional class probabilities for \mathbf{x}_{test} . As the whole training set \mathcal{D}_{train} is passed as input/context to the Transformer, it learns to predict class labels through sample-to-sample attention.

Inference on real-world data. In the inference stage, a fresh real-word dataset \mathcal{D}_{train} and some test instance \mathbf{x}_{test} are fed into the (frozen) pre-trained model, which computes the PPD $q_\theta(\cdot | \mathbf{x}_{test}, \mathcal{D}_{train})$ in a single forward process. Importantly, PFNs do not require gradient-based parameter tuning on data observed at inference time, where the training and prediction are delivered through a one-step forward process *in less than a second* (Hollmann et al., 2023).

In summary, PFNs are trained once offline, and can be used many times for zero-shot inference when new datasets with different characteristics are input. The main benefit is that **no training or tuning** is required at the inference stage. This type of learning ability is also termed as in-context learning (ICL) (Xie et al., 2021), which was shown to be an effective paradigm for various tasks in NLP with the stream of large language models (Brown et al., 2020). In fact, ICL with PFNs is recently shown to be a promising paradigm for supervised classification on tabular datasets (Hollmann et al., 2023).

3 FoMo-OD: A NEW PFN FOR 0-SHOT OD – MODEL SELECTION BYGONE!

Inspired by the recent PFNs (Müller et al., 2022) and their successful applications in supervised classification (Hollmann et al., 2023) and time series forecasting (Dooley et al., 2023), we propose FoMo-OD, a prior-data fitted Foundation Model for 0-shot Outlier Detection. FoMo-OD is (pre)trained on a large body of synthetically generated OD datasets toward zero-shot inference on a new dataset. Most notable of our zero-shot FoMo-OD is its elimination of the need not only for model training on a new dataset, but especially also for model selection (both algorithm and HPs), which is notoriously-hard without any labeled data. By breaking such new ground, and its effectiveness on many benchmark datasets compared to classical and modern baselines, we expect FoMo-OD will become a milestone in future research and practice of OD. The new FoMo-OD paradigm (right) versus the typical OD setting (left) is illustrated in Figure 1.

In the following we describe our OD data prior, training of FoMo-OD on prior-simulated datasets, inference on new datasets, and our specific model architecture and improvements for scalable training.

216 3.1 DESIGNING A DATA PRIOR FOR OUTLIER DETECTION
217

218 Arguably, what has triggered the recent breakthroughs in NLP and CV is the massive amounts of
219 datasets available for (pre)training, along with high-capacity model architectures. In comparison
220 to the natural language and image domains, the quantity (and quality) of publicly available tabular
221 OD datasets is minuscule. Even in the presence of large quantities of data, in training their Chronos
222 foundation models for time series forecasting, Ansari et al. (2024) show that using synthetic data in
223 combination with real-world data improves the overall zero-shot performance. For these reasons, we
224 design a new data prior from which we simulate numerous OD datasets for pretraining FoMo-OD.

225 Ideally the data prior should reflect distributions as general and diverse as seen in real-world datasets,
226 however, “*finding a prior supporting a large enough subset of possible [data generating] functions*
227 *isn’t trivial*” (Nagler, 2023). Surprisingly, in contrast, our initial attempt has been sufficient to achieve
228 remarkable performance even with a relatively straightforward and simple-to-implement data prior,
229 which we describe next.

230 **Inlier synthesis:** We simulate inliers by simply drawing from a Gaussian Mixture Model (GMM)
231 with m -clusters in d -dimensions, with centers $\mu_{jk} \in [-5, 5]$, $j \in [m]$, $k \in [d]$ and *diagonal*² Σ_j with
232 entries in $[-5, 5]$. In each step of every epoch during pretraining, we create batch size B different
233 GMMs with varying $m \leq M$ and $d \leq D$ chosen uniformly at random from $[M]$ and $[D]$, respectively.
234 From each GMM, we draw a set of S inliers, defined as instances within the 90%-ile of the GMM.

235 **Outlier synthesis:** Following the previous literature on outlier synthesis (Han et al., 2022), we
236 generate *subspace* outliers by first drawing a subset of dimensions \mathcal{K} at random, where $|\mathcal{K}| \leq d$, and
237 then generate S points from the corresponding “inflated” GMMs, which share the same centers μ_j ’s
238 with the original GMM but with the inflated (diagonal) covariances $5 \times \Sigma_{j,kk}$ ’s for $k \in \mathcal{K}$. Outliers
239 are defined as points outside the 90%-ile of the original GMM. We label each sample based on its
240 Mahalanobis distance computed analytically (see Property B.2 in the Appendix).

241 Specifically, we simulate datasets containing $2S = 10,000$ samples (half inlier, half outlier) from
242 the two corresponding GMMs (original and inflated) with up to $M = 5$ clusters and up to $D = 100$
243 dimensions. Example 2-d synthetic datasets are illustrated in Appendix A.

244 **Remarks:** We emphasize once again that our model is not trained on **any** real-world data and rather,
245 on purely synthetic data (although future work can combine existing benchmark OD datasets with
246 synthesized data, as was done for Chronos (Ansari et al., 2024)). Notably, our GMM-based data
247 prior can be seen as extremely basic. While it has been our intent to extend our preliminary attempt
248 toward designing a sophisticated data prior for OD, we found to our surprise that even with such an
249 elementary prior, FoMo-OD performs remarkably well against numerous SOTA baselines. Therefore,
250 we present FoMo-OD using this effortless approach for its simplicity to showcase the prowess of PFNs
251 for OD. Future work can employ BNNs and SCMs (Hollmann et al., 2023), and other outlier types
252 (contextual, dependency, etc. (Steinbuss and Böhm, 2021)) toward a more comprehensive data prior.

253
254 3.2 (PRE)TRAINING AND INFERENCE
255

256 **Model (Pre)Training (Once, Offline):** FoMo-OD is a Prior-data Fitted Network (PFN, see Section
257 2.2) based on the Transformer architecture. In the synthetic prior-data fitting phase, it is trained
258 on datasets drawn from our OD data prior for tabular data that we introduced in Section 3.1. Each
259 dataset is simulated from a different GMM configuration based on randomly drawn parameters, and
260 consists of varying number of training samples and dimensions to capture the diversity in real-world
261 tabular datasets. Detailed steps are outlined in Algo. 1 in Appendix C.2, and described as follows.

262 Each time, we first draw a hypothesis (i.e. GMM configuration) uniformly at random, that is
263 $\phi = \{d \in [D], m \in [M], \{\mu_j\}_{j=1}^m \in [-5, 5]^d, \{\Sigma_j\}_{j=1}^m; \text{diag}(\Sigma_j) \in [-5, 5]^d\}$, and then generate
264 a dataset $\mathcal{D} = \{\mathcal{D}_{\text{in}}, \mathcal{D}_{\text{out}}\}$ containing synthetic inlier and outlier samples from the drawn hypothesis
265 and its variance-inflated variant, respectively.

266 We optimize FoMo-OD’s parameters θ to make predictions on $\mathcal{D}_{\text{test}} = \{\mathcal{D}_{\text{test}}^{\text{in}}, \mathcal{D}_{\text{test}}^{\text{out}}\}$, conditioned
267 on the inlier-only training data $\mathcal{D}_{\text{train}} \subset \mathcal{D}_{\text{in}}$ based on the cross-entropy loss (see Eq. (2)). During

268 2In our early experiments, we found no difference in terms of test performance on synthetic datasets between
269 using diagonal and non-diagonal Σ , however, it is easier to compute the inverse of diagonal Σ for generation.

270 training, $\mathcal{D}_{\text{test}}$ contains a *balanced* number of inlier and outlier samples, where $\mathcal{D}_{\text{test}}^{\text{in}} = \mathcal{D}_{\text{in}} \setminus \mathcal{D}_{\text{train}}$,
 271 and $\mathcal{D}_{\text{test}}^{\text{out}} \subset \mathcal{D}_{\text{out}}$ contains an equal number of samples as $\mathcal{D}_{\text{test}}^{\text{in}}$. To vary the training data size, we
 272 subsample $\mathcal{D}_{\text{train}}$ of randomly drawn size $n \in [n_L, n_U]$, where n_L and n_U denote the lower and
 273 upper bounds. In our current implementation, we set $n_L = 500$, and $n_U = 5,000$.

274 FoMo-0D is trained on 200,000 batches (200 epochs \times 1,000 steps/epoch) of $B = 8$ generated
 275 datasets in each batch. While this pretraining phase can be expensive, it is done *only once, offline*.
 276 Moreover, we introduce several scalability improvements to speed up pretraining, as discussed later
 277 in Section 3.3. Full details on the training and implementation of FoMo-0D are given in Appendix C.
 278

279 **Zero-shot Inference (on Unseen Dataset):** During the inference phase, our pretrained-in-advance
 280 FoMo-0D can be employed on any unseen real-world dataset. In fact, we apply the same single
 281 pretrained network on all benchmark datasets in our experiments in this paper.

282 Specifically, for a new semi-supervised OD task with inlier-only training data $\mathcal{D}_{\text{train}}$ and mixed
 283 test data $\mathcal{D}_{\text{test}}$, feeding $\langle \mathcal{D}_{\text{train}}, \mathbf{x}_{\text{test}} \rangle$ as input to FoMo-0D (for each $\mathbf{x}_{\text{test}} \in \mathcal{D}_{\text{test}}$ separately) yields
 284 the PPD $q_{\theta}(y|\mathbf{x}_{\text{test}}, \mathcal{D}_{\text{train}})$ in a *single forward pass*. As such, FoMo-0D performs model “training”
 285 and prediction *simultaneously* at test time. In fact, as the entire training data is passed as context,
 286 FoMo-0D leverages in-context learning (ICL) (Xie et al., 2021; Garg et al., 2022) for inference.
 287 The **key** contribution of FoMo-0D goes beyond eliminating gradient-based model training for a new
 288 dataset: because no model training is required, one thus neither needs to choose any specific OD
 289 model to train, nor grapple with tuning any hyperparameters of the said model—rendering model
 290 selection an obsolete concern for the future of OD. Additionally, the speedy, easily parallelizable
 291 inference (for *less-than-a-second* per test sample) is then the “icing on the cake”.

292 For a visual summary, Figure 1 (right) illustrates (top) pretrain & (bottom) test phases of FoMo-0D.
 293

294 3.3 ARCHITECTURE AND SCALABILITY

295 **Architecture and sample-to-sample attention:** Like existing PFNs in the literature, FoMo-0D is
 296 based on the Transformer architecture (Vaswani et al., 2017), encoding each sample’s feature vector as
 297 a token, and allowing token representations to attend to each other, hence enabling *sample-to-sample*
 298 *attention*. We also adopt the three adaptations of TabPFN (Hollmann et al., 2023), which (1) computes
 299 self-attention among all the training samples but only *cross*-attention from test samples to the training
 300 samples, (2) enables variable feature dimensionality by zero-padding, and (3) randomly rotates input
 301 samples while omitting positional encodings to achieve model invariance to sample permutations in
 302 the dataset. We defer the architecture details to the original papers.

303 Given $\mathcal{D}_{\text{train}} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, each self-attention layer outputs n embeddings $\{\mathbf{z}_i\}_{i=1}^n$; where the
 304 i -th token is mapped via linear transformations to a key \mathbf{k}_i , query \mathbf{q}_i and value \mathbf{v}_i based on which the
 305 i -th output is computed by weighing all \mathbf{v}_j ’s by the normalized dot product between \mathbf{q}_i and all the
 306 \mathbf{k}_j ’s (i.e. sample-to-sample dot product similarity) as

$$308 \quad \mathbf{z}_i = \sum_{j=1}^n \text{softmax}(\{\langle \mathbf{q}_i, \mathbf{k}_j \rangle\}_{j=1}^n)_j \cdot \mathbf{v}_j . \quad (3)$$

310 The sample-to-sample attention is intriguing from the perspective of OD: Many classical OD algo-
 311 rithms (Aggarwal, 2013) are based on nonparametrics; in particular, they make use of the distances
 312 to the k *nearest* neighbors (k NNs) of a point to compute its outliersness (where k is a critical hy-
 313 perparameter (HP)). One can think of FoMo-0D as mimicking non-parametric models but by using
 314 parametric attention mechanisms. Interestingly, PFNs are much more robust and flexible than k NN
 315 based OD approaches, for (1) sample-to-sample relations are not pre-specified but rather learned
 316 through attention weights, and thus (2) they are not limited to just the nearest neighbors but rather can
 317 *learn which* training points are worth attending to, and last but not least (3) as attention is dataset-wide
 318 across all points, there is no need for specifying a cut-off HP value like k , to which most k NN based
 319 OD techniques are sensitive to (Aggarwal and Sathe, 2015; Campos et al., 2016; Goldstein and
 320 Uchida, 2016; Ding et al., 2022)—to reiterate, algorithm & HP selection is bygone with FoMo-0D.

321 While intuitively beneficial for OD, “vanilla” attention among the training samples incurs quadratic
 322 complexity. To be able to seize the benefits with scale, we incorporate a scalable architecture to our
 323 design, as we describe next. The scale up also unlocks a larger context (i.e. dataset) size for FoMo-0D,
 enabling its pretraining on larger datasets for potentially better generalization.

Scaling up attention with “routers”: The $\mathcal{O}(n^2)$ quadratic sample complexity at pretraining presents an obstacle for achieving high performance at inference. From dataset size perspective, it limits pretraining to relatively small training datasets. From context size perspective, it limits in-context learning that typically benefits from longer context lengths (Xie et al., 2021).

Toward a high-performance pretrained model, we scale up FoMo-0D’s attention via the “router mechanism” of Zhang and Yan (2023). As shown in Figure 2, the main idea is to learn a small number ($R \ll n$) of “routers” or representatives, which gather information from all n samples and then distribute the information back to the n output embeddings, creating what-looks-like a “bottleneck” attention mechanism—reducing complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}(2Rn) = \mathcal{O}(n)$. This design allows FoMo-0D training to **scale linearly** with respect to both dimensionality d and also dataset size n .

Concretely, the representatives first aggregate information from all samples by serving as query in multi-head self-attention (MSA) and the embedding array of all samples becomes both key and value:

$$\mathbf{M} = \text{MSA}_1(\mathbf{R}, \mathbf{Z}, \mathbf{Z}), \quad (4)$$

where $\mathbf{R} \in \mathbb{R}^{R \times d}$ depicts the *learnable* vector array of representatives and \mathbf{M} denotes the aggregated messages. Then, the routers distribute the received information among samples by using the sample embeddings as query and the aggregated messages as both key and value:

$$\hat{\mathbf{Z}} = \text{MSA}_2(\mathbf{Z}, \mathbf{M}, \mathbf{M}). \quad (5)$$

Finally, we obtain $\bar{\mathbf{Z}} = \text{LayerNorm}(\hat{\mathbf{Z}} + \mathbf{Z})$ after layer normalization. Note that the test samples only attend to the training samples’ embeddings, computed in the described manner across layers, which finally feed into the prediction head for estimating each test sample’s PPD at the output layer.

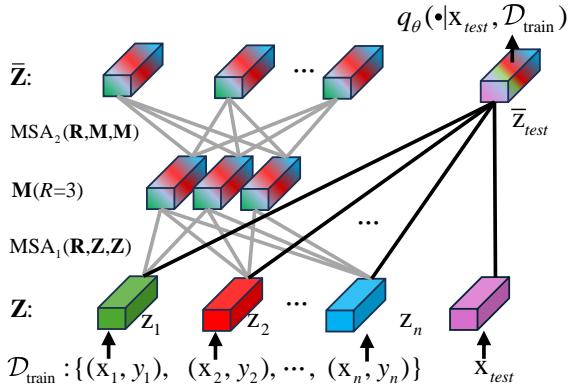


Figure 2: FoMo-0D architecture employs the “router mechanism” for scalable attention.

Scaling up (pre)training data synthesis with linear transforms: Besides the scalability challenge associated with architecture/attention, another computational challenge in pretraining FoMo-0D arises from drawing samples from the data prior. That is, generating samples from a pre-specified data distribution requires considerable time, especially in high dimensions³, provided the large number of datasets we sample (concretely, a batch size of 8 datasets over 1,000 steps each for 200 epochs).

To give an idea, sampling a dataset with $n = 10,000$ points in $d = 100$ dimensions using 10 CPUs in parallel takes ≈ 0.4 seconds (see Appendix Figure 7). Across 200 training epochs with 1,000 steps each, it adds up to more than 177 hours just to generate 1,6 million datasets on-the-fly. Of course, one can trade storage with compute-time by generating all these datasets apriori via massive parallelism. Nevertheless, synthetic data generation demands considerable time (and/or storage).

To scale up data synthesis, FoMo-0D employs two distinct strategies. **First**, we propose *reuse at epoch level*: that is, one can reuse the same 8K unique datasets at every epoch, or in general, the same $8K \times P$ datasets periodically at every P epochs. A larger P would lead to more diversity in terms of the overall pretraining data used.

Second, and more innovatively, we propose *reuse at dataset level via transformation*: that is, having generated one unique dataset $\mathbf{X} \in \mathbb{R}^{n \times d}$ from a GMM, we propose a linear transform $T(\mathbf{x})$ of the form $\mathbf{W}\mathbf{x} + \mathbf{b}$ for randomly drawn parameters $\mathbf{W} \in \mathbb{R}^{d \times d}$ and $\mathbf{b} \in \mathbb{R}^d$ (see Appendix B.1).⁴ This simple yet efficient transformation creates a new dataset, akin to one being drawn from another GMM with centers $T(\mu_j) = \mathbf{W}\mu_j + \mathbf{b}$ and covariance $T(\Sigma_j) = \mathbf{W}\Sigma_j\mathbf{W}^T, \forall j \in [m]$. Note that we do not actually materialize these parameters but only transform the dataset. As we show in the following, such transformations preserve the Mahalanobis distances as well as the percentile thresholds for labeling points as inlier/outlier. Details and proofs are given in Appendix B.

³This is because the inverse of the $(d \times d)$ covariance matrix plays a crucial role in the process of generating samples from a GMM, which has $\mathcal{O}(d^3)$ time complexity. (It is also partly the reason why we use diagonal Σ_j ’s in our data prior.) In addition, Mahalanobis distance for labeling inliers/outliers also requires the inverse.

⁴In practice, we apply the linear transform on the subspace of inflated features only, wherein inliers and outliers are defined, which remains to be a multi-variate GMM.

378 **Lemma 1** Linear transform T with invertible \mathbf{W} on \mathcal{G}_m^d preserves Mahalanobis distances.

379 **Lemma 2** Linear transform T with invertible \mathbf{W} on \mathcal{G}_m^d preserves the percentiles of the GMM.

380
381 The implication of these lemmas is that a linear transformation of a dataset from a GMM retains the
382 identity of the inliers and outliers, i.e. no relabeling is required. Moreover, notice that as a byproduct
383 we obtain a transformed dataset as though it is drawn from a GMM with a *non-diagonal* covariance
384 matrix which, besides the time savings, offers a slightly more complex data prior.

385 To reach 8K unique datasets for each epoch, we first generate 500 datasets from different GMMs (with
386 varying configurations), and then employ 15 different linear transformations to each unique dataset
387 by varying \mathbf{W} and \mathbf{b} . Drawing each (\mathbf{W}, \mathbf{b}) takes ≈ 0.02 seconds, while the matrix-matrix product
388 of \mathbf{X} ($n \times d$) and \mathbf{W} ($d \times d$) takes negligible time (for $d \leq 100$). Thus, obtaining a transformed
389 dataset offers $20\times$ speed-up compared to generating one (0.02 vs. 0.4 seconds).

390 4 EXPERIMENTS

391 4.1 SETUP

392 We present the experiment setup briefly, including important notes on data synthesis, real-world
393 datasets, baselines, metrics and HPs. For additional details, we refer to Appendix D.

394 **Pre-training Dataset Synthesis:** During pretraining, we generate unique GMM datasets by first
395 drawing a configuration, including dimensionality $d \in [D]$, number of components $m \in [M]$, centers
396 $\{\mu_j\}_{j=1}^m$ (each $\mu_j \in [-5, 5]^d$) and covariances $\{\Sigma_j\}_{j=1}^m$ ($\text{diag}(\Sigma_j) \in [-5, 5]^d$). We set $M = 5$
397 and vary $D \in \{20, 100\}$ to study pretraining with relatively small and high dimensional datasets,
398 respectively. We synthesize inliers and outliers as described in Section 3.1.

400 **Real-world Benchmark Datasets:** While pretraining is purely on synthetic datasets, we evaluate
401 FoMo-OD on 57 real-world datasets from the ADBench benchmark (Han et al., 2022) (see Table 15).

402 We use 5 train/test splits of each dataset via different seeds and report mean performance and standard
403 deviation. Note that the baselines require model re-training and inference for each $\mathcal{D}_{\text{train}}/\mathcal{D}_{\text{test}}$ split,
404 while FoMo-OD uses the splits only for inference as $\mathcal{D}_{\text{train}}$ is merely passed as context.

405 **Baselines:** We compare FoMo-OD against 26 baselines, from classical/shallow methods to modern/deep
406 models. The baselines are imported from one of the latest papers that proposed the SOTA
407 diffusion-based model DTE (Livernoche et al., 2024), and its three variants; DTE-C, DTE-IG, and
408 DTE-NP. We defer to the original paper for additional details.

409 **Model Implementation:** We trained our final model for 200,000 steps with a batch size of 8 datasets.
410 That is, our FoMo-OD is trained on 1,600,000 synthetically generated datasets. This training takes
411 about 25 hours on 1 GPU (Nvidia RTX A6000). Each dataset had a fixed size of 10,000 samples,
412 with $|\mathcal{D}_{\text{train}}| \in [n_L = 500, n_U = 5000]$, and the rest used as $\mathcal{D}_{\text{test}}$ with *balanced* number of inliers
413 and outliers. Other implementation details of FoMo-OD, including the training algorithm, model
414 architecture, data synthesis and reuse, and hardware are provided in Appendix C.

415 **Metrics and Hypothesis Testing:** Detection performance is w.r.t. 3 widely-used metrics for OD:
416 AUROC; area under ROC curve, AUPR; area under Precision-Recall curve, and F1 score; using
417 threshold at the true number of outliers in the test data (varies by dataset).

418 To compare methods, we compute their rank on each dataset (lower is better), and present average
419 rank across datasets. This is an alternative to the average metric (e.g. AUROC), which is not
420 meaningful when task difficulties and hence metric values vary widely. In addition, we perform
421 significance tests to compare two methods statistically, using the one-sided paired Wilcoxon signed
422 rank test (Demšar, 2006) between FoMo-OD and a baseline based on the performances across all
423 datasets and report the p -values. We consider results to be significant at 0.05 following convention.

424 **Hyperparameters (HPs):** Importantly, Livernoche et al. (2024) picked for each baseline the best-
425 performing set of HPs as recommended by the authors in their original paper. As for their own DTE,
426 which behaves similar to k NN, they use $k = 5$ and set the *same* k for the k NN baseline (Ramaswamy
427 et al., 2000) to be consistent. However, it is well known that k NN is sensitive to the value of k
428 (Aggarwal and Sathe, 2015), and so are many other OD models to their respective HPs (Campos
429 et al., 2016; Goldstein and Uchida, 2016; Zhao et al., 2021; Ding et al., 2022).

Therefore, we compare to the performance results of these baselines as imported from DTE’s Tables 13, 14 and 15, respectively for AUROC, F1, and AUPR (Livernoche et al., 2024). In addition, we also compare to the **top-4**⁵ best performing baselines (in order: DTE-NP, kNN, ICL, and DTE-C) on their *average* performance across a list of different HP settings (which reflects their *expected* performance under HP values selected at random, in the absence of any other prior knowledge), which is the recommended approach by Goldstein and Uchida (2016) “*to get a fair evaluation when comparing [OD] algorithms*”. We annotate the method name with ^{avg} for the version with performance averaged over varying HPs. The detailed list of HP values for each top baseline is given in Appendix D.4.

Overall, we compare FoMo-0D to 30 baselines; 26 from Livernoche et al. (2024) and ^{avg} of the top-4.

4.2 RESULTS

Detection performance: Table 1 presented the comparison of FoMo-0D w/ $D = 100$ to all baselines w.r.t. average rank across datasets as well as pairwise Wilcoxon signed rank tests based on AUROC (for full results on all datasets and all metrics, see Appendix G). We find that among 30 baselines and 2 variants of FoMo-0D (w/ $D = 100$ and $D = 20$), FoMo-0D w/ $D = 100$ *performs as well as the 2nd best model* (k NN with default HP; $k = 5$) on all datasets. While DTE-NP outperforms FoMo-0D with author-recommended $k = 5$, we find that DTE-NP^{avg} is on par with FoMo-0D.

Against all other baselines, we obtain notably large p -values. Typically, $p > 0.05$ implies no statistical difference between two methods. On the other hand, the large p -values we obtain that are often larger than 0.50 suggest that the odds are tilted towards FoMo-0D to outperform.

FoMo-0D w/ $D = 100$ performs statistically no different from **all** baselines on datasets with $d \leq 100$ (i.e., “at its own game” when pretraining data dimensions align with real-world datasets), while it *outperforms the majority of baselines* ($p > 0.95$). These results also hold on datasets with $d \leq 500$.

Table 2 shows similar results for FoMo-0D w/ $D = 20$, which is pretrained on datasets with considerably fewer dimensions. Even in this limited setting, its performance is remarkable: against 30 baselines, it performs on par with the *3rd* best baseline (ICL, with default HP). The p -value is even larger (0.437) when compared to ICL^{avg}. Moreover, on datasets with $d \leq 20$ which align with its pretraining data, all p -values are larger than 0.5, where it outperforms the top *5th* baseline and the majority of others. These are outstanding results for a model pretrained purely on synthetic datasets from a simple data prior in small dimensions, showcasing the prowess of PFNs for OD.

Table 2: Comparison of methods across datasets. (top row) Rank w.r.t. AUROC performance ^{avg}’ed over 57 datasets is presented for FoMo-0D (with $D = 20$), **top-10** baselines with default HPs, and **top-4**⁵ baselines with performance ^{avg}’ed over varying HPs (denoted w/ ^{avg}); followed by p -values of the pairwise Wilcoxon signed rank test, comparing FoMo-0D to each baseline (from top to bottom) over All (57) datasets, those (24) w/ $d \leq 20$ and (38) datasets w/ $d \leq 50$ dimensions, respectively. Even with small $D = 20$, FoMo-0D performs as well as (i.e., statistically no different at 0.05 from) the *top 3rd baseline* (ICL, w/ $p = 0.089$) across All datasets, while it *outperforms the top 5th (LOF) and onward baselines* over datasets w/ $d \leq 20$ (aligned w/ pretraining where $D = 20$) and $d \leq 50$ (generalizing beyond pretraining). (setting: $D = 20$, $P = 50$, $R = 500$, train/inference context size=5K, no data transformation)

	FoMo-0D	DTE-NP	kNN	ICL	DTE-C	LOF	CBLOF	Feat.Bag.	SLAD	DDPM	OCSVM	DTE-NP ^{avg}	kNN ^{avg}	ICL ^{avg}	DTE-C ^{avg}
Rank(avg)	12.59	7.19	8.57	10.34	10.79	11.82	12.81	12.8	12.52	13.50	13.34	8.60	10.63	12.44	21.43
All	-	<u>0.001</u>	<u>0.019</u>	0.089	0.159	0.394	0.434	0.703	0.516	0.752	0.679	<u>0.007</u>	0.062	0.437	1.0
$d \leq 20$	-	0.572	0.789	0.968	0.616	0.993	0.989	1.0	0.978	0.906	0.992	0.813	0.924	0.999	1.0
$d \leq 50$	-	0.347	0.794	0.893	0.946	0.997	0.988	1.0	0.963	0.994	0.986	0.574	0.847	0.995	1.0

Figure 3 shows the distribution of ranks across datasets for each of the 32 methods. While paired significant tests are the most conclusive, FoMo-0D achieves relatively small average rank as well as notably low ranks across datasets that is also visually better than the majority of the baselines.

Running time: Table 3 presents the total training time and the average inference time per test sample, as measured on our largest benchmark dataset, for FoMo-0D and the top-3 baselines. Given

⁵To rank the baselines, we compute the 26×26 pairwise p -values based on the Wilcoxon signed rank test, as shown in Appendix Figure 16, and rank the baselines w.r.t. their mean p -value.

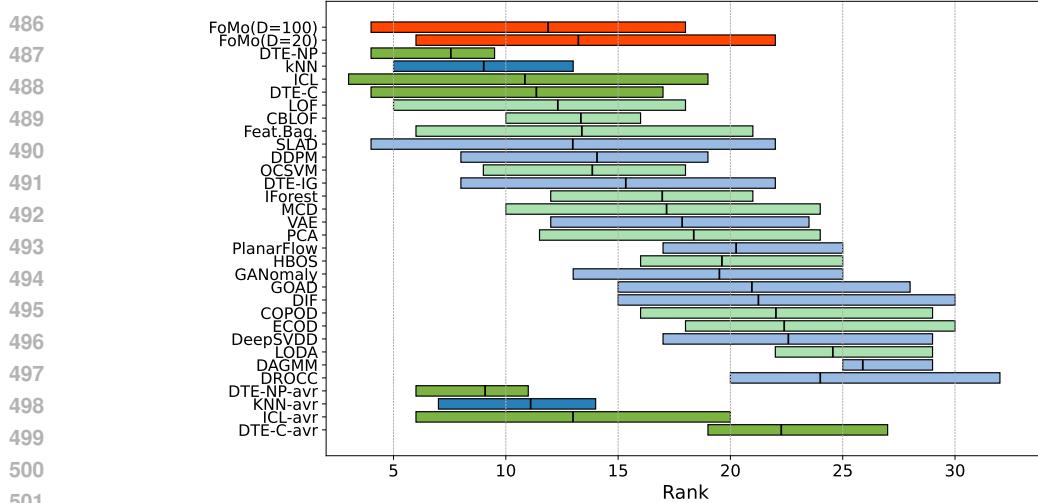


Figure 3: (best in color) Rank (w.r.t. AUROC performance, lower is better) distribution across datasets shown via boxplots for (from top to bottom) FoMo-OD in red, all 26 baselines as ordered by mean p -value⁵ (shallow and deep baselines resp. in green and blue), and top-4 baselines’ avg variants.

a new dataset, FoMo-OD bypasses model training (and HP tuning) and directly performs inference, with an average of 7.7 ms per sample (see Appendix Figure 6). In comparison, all baseline methods need to train on each individual dataset preceding inference. This training time can be high for deep learning based models like ICL, and further compounded with training *multiple* models for hyperparameter tuning purposes. Even for non-parametric and/or shallow models like kNN and DTE-NP (which queries k nearest neighbors), the training involves various data pre-processing steps such as constructing a tree-like data structure for fast (often approximate) kNN distance querying.

Table 3: Training and inference time (in milliseconds) comparison between FoMo-OD and the top-3⁵ baselines (w/ default HPs, *excluding* the time for model selection/hyperparameter optimization) on our largest dataset (namely, `donors`, see Appendix Table 15).

Method	FoMo-OD	DTE-NP	k NN	ICL
Training time (total)	none	56.83	1433.74	186461.48
Inference time (per sample)	7.7	0.76	0.17	0.01

4.3 ABLATION ANALYSES

Due to space limits, we present the detailed ablation analyses in Appendix E. We discuss the effect of D in E.1, the cost and performance of varying R in E.2 and E.3, the context size in E.4, the reuse periodicity P in E.5, the effect of data transformation T on performance and speed up in E.6 and E.7, data diversity and prolonged training in E.8, and quantile transformation on ADBench in E.9.

5 RELATED WORK

Due to space limits, we present the detailed related work in Appendix J.

6 CONCLUSION

We introduced FoMo-OD, **the first foundation model for outlier detection** (OD) on tabular data. It capitalizes on the in-context learning ability of a Transformer model pretrained on a large number of synthetic datasets that can then perform zero-shot inference on a new dataset by directly passing it as input context. FoMo-OD breaks new ground by fully abolishing notoriously-hard model selection. Further, FoMo-OD offers extremely fast inference thanks to a mere single forward pass. Against **26** baselines on **57** public datasets from diverse domains, FoMo-OD performs on par with the *2nd* best baseline, while significantly outperforming the majority of baselines. We leave improving the data prior and extending beyond tabular OD, among others, as future directions. For a detailed discussion on limitations and future directions, we refer to Appendix K.

540 REFERENCES
541

- 542 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
543 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical
544 report. *arXiv preprint arXiv:2303.08774* (2023).
- 545 Steven Adriaensen, Herilalaina Rakotoarison, Samuel Müller, and Frank Hutter. 2024. Efficient
546 bayesian learning curve extrapolation using prior-data fitted networks. *Advances in Neural
547 Information Processing Systems* 36 (2024).
- 548 Charu C. Aggarwal. 2013. *Outlier Analysis*. Springer.
- 549 Charu C. Aggarwal and Saket Sathe. 2015. Theoretical foundations and algorithms for outlier
550 ensembles. *Acm sigkdd explorations newsletter* 17, 1 (2015), 24–47.
- 551 Charu C. Aggarwal and Saket Sathe. 2017. *Outlier Ensembles - An Introduction*. Springer.
- 552 Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. 2019. Ganomaly: Semi-supervised
553 anomaly detection via adversarial training. In ACCV. Springer, 622–637.
- 554 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel
555 Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual
556 language model for few-shot learning. *Advances in neural information processing systems* 35
557 (2022), 23716–23736.
- 558 Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen,
559 Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al.
560 2024. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815* (2024).
- 561 Lion Bergman and Yedid Hoshen. 2020. Classification-Based Anomaly Detection for General Data.
562 In *International Conference on Learning Representations*.
- 563 Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. 2000. LOF: Identifying
564 Density-Based Local Outliers. In *International Conference on Management of Data*.
- 565 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
566 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel
567 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler,
568 Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray,
569 Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever,
570 and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural
571 Information Processing Systems*, Vol. 33. 1877–1901.
- 572 Guilherme O Campos, Arthur Zimek, Jörg Sander, Ricardo JGB Campello, Barbora Micenková,
573 Erich Schubert, Ira Assent, and Michael E Houle. 2016. On the evaluation of unsupervised outlier
574 detection: measures, datasets, and an empirical study. *Data mining and knowledge discovery* 30
575 (2016), 891–927.
- 576 George Casella and Roger Berger. 2024. *Statistical inference*. CRC Press.
- 577 Raghavendra Chalapathy and Sanjay Chawla. 2019. Deep learning for anomaly detection: A survey.
578 *arXiv preprint arXiv:1901.03407* (2019).
- 579 Janez Demšar. 2006. Statistical comparisons of classifiers over multiple data sets. *The Journal of
580 Machine learning research* 7 (2006), 1–30.
- 581 Xueying Ding, Lingxiao Zhao, and Leman Akoglu. 2022. Hyperparameter sensitivity in deep outlier
582 detection: Analysis and a scalable hyper-ensemble solution. *Advances in Neural Information
583 Processing Systems* 35 (2022), 9603–9616.
- 584 Xueying Ding, Yue Zhao, and Leman Akoglu. 2024. Fast Unsupervised Deep Outlier Model Selection
585 with Hypernetworks. *ACM SIGKDD* (2024).

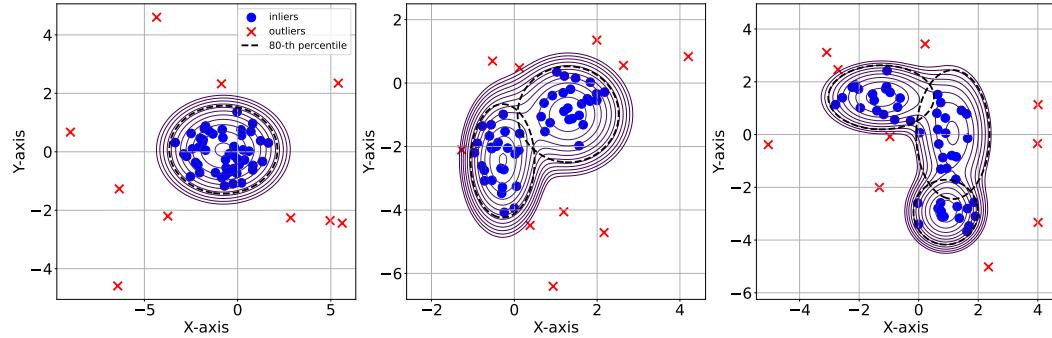
- 594 Samuel Dooley, Gurnoor Singh Khurana, Chirag Mohapatra, Siddartha Venkat Naidu, and Colin
 595 White. 2023. ForecastPFN: Synthetically-Trained Zero-Shot Forecasting. In *Thirty-seventh Con-*
 596 *ference on Neural Information Processing Systems*.
- 597
- 598 Xuefeng Du, Yiyou Sun, Jerry Zhu, and Yixuan Li. 2024. Dream the impossible: Outlier imagination
 599 with diffusion models. *Advances in Neural Information Processing Systems* 36 (2024).
- 600 Sepideh Esmaeilpour, Bing Liu, Eric Robertson, and Lei Shu. 2022. Zero-shot out-of-distribution
 601 detection based on the pre-trained model CLIP. In *Proceedings of the AAAI conference on artificial*
 602 *intelligence*, Vol. 36. 6568–6576.
- 603
- 604 Benjamin Feuer, Niv Cohen, and Chinmay Hegde. 2023. Scaling TabPFN: Sketching and Feature
 605 Selection for Tabular Prior-Data Fitted Networks. In *NeurIPS 2023 Second Table Representation*
 606 *Learning Workshop*.
- 607 Benjamin Feuer, Robin Tibor Schirrmeyer, Valeria Cherepanova, Chinmay Hegde, Frank Hutter,
 608 Micah Goldblum, Niv Cohen, and Colin White. 2024. TuneTables: Context Optimization for
 609 Scalable Prior-Data Fitted Networks. arXiv:2402.11137 [cs.LG]
- 610
- 611 Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. 2022. What can transformers
 612 learn in-context? a case study of simple function classes. *Advances in Neural Information*
 613 *Processing Systems* (2022).
- 614 Markus Goldstein and Andreas Dengel. 2012. Histogram-based outlier score (HBOS): A fast
 615 unsupervised anomaly detection algorithm. *KI-2012: poster and demo track 1* (2012), 59–63.
- 616
- 617 Markus Goldstein and Seiichi Uchida. 2016. A comparative evaluation of unsupervised anomaly
 618 detection algorithms for multivariate data. *PloS one* 11, 4 (2016).
- 619
- 620 Sachin Goyal, Aditi Raghunathan, Moksh Jain, Harsha Simhadri, and Prateek Jain. 2020. DROCC:
 621 Deep Robust One-Class Classification. In *Proceedings of the 37th International Conference on*
 622 *Machine Learning (ICML'20)*. JMLR.org, Article 348, 11 pages.
- 623
- 624 Zhaopeng Gu, Bingke Zhu, Guibo Zhu, Yingying Chen, Ming Tang, and Jinqiao Wang. 2024.
 625 AnomalyGPT: Detecting industrial anomalies using large vision-language models. In *Proceedings*
 626 *of the AAAI Conference on Artificial Intelligence*, Vol. 38. 1932–1940.
- 627
- 628 A. Gut. 2009. *An Intermediate Course in Probability*. Springer New York.
- 629
- 630 Songqiao Han, Xiyang Hu, Hailiang Huang, Minqi Jiang, and Yue Zhao. 2022. Adbench: Anomaly
 631 detection benchmark. *Advances in Neural Information Processing Systems* 35 (2022).
- 632
- 633 Haoyang He, Jiangning Zhang, Hongxu Chen, Xuhai Chen, Zhishan Li, Xu Chen, Yabiao Wang,
 634 Chengjie Wang, and Lei Xie. 2024. A diffusion-based framework for multi-class anomaly detection.
 635 In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 8472–8480.
- 636
- 637 Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances*
 638 *in Neural Information Processing Systems* 33 (2020), 6840–6851.
- 639
- 640 Hadi Hojjati, Thi Kieu Khanh Ho, and Narges Armanfar. 2022. Self-supervised anomaly detection:
 641 A survey and outlook. *arXiv preprint arXiv:2205.05173* (2022).
- 642
- 643 Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. 2023. TabPFN: A
 644 Transformer That Solves Small Tabular Classification Problems in a Second. In *The Eleventh*
 645 *International Conference on Learning Representations*.
- 646
- 647 Jongheon Jeong, Yang Zou, Taewan Kim, Dongqing Zhang, Avinash Ravichandran, and Onkar
 648 Dabeer. 2023. Winclip: Zero-/few-shot anomaly classification and segmentation. In *Proceedings*
 649 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 19606–19616.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child,
 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language
 models. *arXiv preprint arXiv:2001.08361* (2020).

- 648 Diederik P Kingma. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
 649
- 650 Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization.
 651 *arXiv:1412.6980 [cs.LG]*
- 652 Aodong Li, Chen Qiu, Marius Kloft, Padhraic Smyth, Maja Rudolph, and Stephan Mandt. 2023.
 653 Zero-Shot Anomaly Detection via Batch Normalization. In *Thirty-seventh Conference on Neural*
 654 *Information Processing Systems*.
- 655 Aodong Li, Yunhan Zhao, Chen Qiu, Marius Kloft, Padhraic Smyth, Maja Rudolph, and Stephan
 656 Mandt. 2024. Anomaly Detection of Tabular Data Using LLMs. *arXiv preprint arXiv:2406.16308*
 657 (2024).
- 658 Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2008. Isolation Forest. In *2008 Eighth IEEE*
 659 *International Conference on Data Mining*. 413–422.
- 660 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning.
 661 *Advances in neural information processing systems* 36 (2024).
- 662 Victor Livernoche, Vineet Jain, Yashar Hezaveh, and Siamak Ravanbakhsh. 2024. On Diffusion
 663 Modeling for Anomaly Detection.. In *ICLR*.
- 664 Philipp Liznerski, Lukas Ruff, Robert A Vandermeulen, Billy Joe Franks, Klaus-Robert Müller, and
 665 Marius Kloft. 2022. Exposing outlier exposure: What can be learned from few, one, and zero
 666 outlier images. *arXiv preprint arXiv:2205.11474* (2022).
- 667 Junwei Ma, Valentin Thomas, Guangwei Yu, and Anthony L. Caterini. 2024. In-Context Data
 668 Distillation with TabPFN. *CoRR* abs/2402.06971 (2024).
- 669 Martin Q Ma, Yue Zhao, Xiaorong Zhang, and Leman Akoglu. 2023. The need for unsupervised
 670 outlier model selection: A review and evaluation of internal evaluation strategies. *ACM SIGKDD*
 671 *Explorations Newsletter* 25, 1 (2023), 19–35.
- 672 Samuel Müller, Matthias Feurer, Noah Hollmann, and Frank Hutter. 2023. PFNs4BO: in-context
 673 learning for Bayesian optimization. In *International Conference on Machine Learning*.
- 674 Samuel Müller, Noah Hollmann, Sebastian Pineda-Arango, Josif Grabocka, and Frank Hutter. 2022.
 675 Transformers Can Do Bayesian Inference. *ICLR* (2022).
- 676 Thomas Nagler. 2023. Statistical Foundations of Prior-Data Fitted Networks.. In *ICML (Proceedings*
 677 *of Machine Learning Research, Vol. 202)*. PMLR.
- 678 Radford M Neal. 2012. *Bayesian learning for neural networks*. Vol. 118. Springer Science &
 679 Business Media.
- 680 Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. 2021. Deep learning for
 681 anomaly detection: A review. *ACM computing surveys (CSUR)* 54, 2 (2021), 1–38.
- 682 Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. 2021. Deep learning on a data diet:
 683 Finding important examples early in training. *Advances in neural information processing systems*
 684 34 (2021), 20596–20607.
- 685 Judea Pearl. 2009. *Causality*. Cambridge University Press.
- 686 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 687 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable
 688 visual models from natural language supervision. In *International conference on machine learning*.
- 689 Sridhar Ramaswamy, Rajeev Rastogi, and Kyuseok Shim. 2000. Efficient algorithms for mining
 690 outliers from large data sets. In *ACM SIGMOD international conference on management of data*.
- 691 Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen,
 692 and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International conference on*
 693 *machine learning*. Pmlr, 8821–8831.

- 702 Allan Raventós, Mansheej Paul, Feng Chen, and Surya Ganguli. 2024. Pretraining task diversity and
 703 the emergence of non-bayesian in-context learning for regression. *Advances in Neural Information
 704 Processing Systems* (2024).
- 705 Danilo Rezende and Shakir Mohamed. 2015. Variational Inference with Normalizing Flows. In
 706 *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine
 707 Learning Research, Vol. 37)*. PMLR, Lille, France, 1530–1538.
- 708 Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek,
 709 Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. 2021. A unifying review of deep
 710 and shallow anomaly detection. *Proc. IEEE* 109, 5 (2021), 756–795.
- 711 Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander
 712 Binder, Emmanuel Müller, and Marius Kloft. 2018. Deep One-Class Classification. In *International
 713 Conference on Machine Learning*. 4393–4402.
- 714 Tom Shenkar and Lior Wolf. 2022. Anomaly Detection for Tabular Data with Internal Contrastive
 715 Learning. In *International Conference on Learning Representations*.
- 716 Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari Morcos. 2022. Beyond
 717 neural scaling laws: beating power law scaling via data pruning. *Advances in Neural Information
 718 Processing Systems* 35 (2022), 19523–19536.
- 719 Georg Steinbuss and Klemens Böhm. 2021. Benchmarking unsupervised outlier detection with
 720 realistic synthetic data. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 15, 4
 721 (2021), 1–20.
- 722 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz
 723 Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*. 5998–6008.
- 724 Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. An explanation of
 725 in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080* (2021).
- 726 Hongzuo Xu, Yijie Wang, Juhui Wei, Songlei Jian, Yizhou Li, and Ning Liu. 2023. Fascinating
 727 supervisory signals and where to find them: Deep anomaly detection with scale learning. In
 728 *International Conference on Machine Learning*. PMLR, 38655–38673.
- 729 Jaemin Yoo, Tiancheng Zhao, and Leman Akoglu. 2023. Data Augmentation is a Hyperparameter:
 730 Cherry-picked Self-Supervision for Unsupervised Anomaly Detection is Creating the Illusion of
 731 Success. *Trans. Mach. Learn. Res.* 2023 (2023).
- 732 Sangwoong Yoon, Young-Uk Jin, Yung-Kyun Noh, and Frank Park. 2023. Energy-based models
 733 for anomaly detection: A manifold diffusion recovery approach. *Advances in Neural Information
 734 Processing Systems* 36 (2023), 49445–49466.
- 735 Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. 2022. Scaling vision trans-
 736 formers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- 737 Yunhao Zhang and Junchi Yan. 2023. Crossformer: Transformer Utilizing Cross-Dimension Depen-
 738 dency for Multivariate Time Series Forecasting.. In *ICLR*.
- 739 Yue Zhao and Leman Akoglu. 2024. Toward unsupervised outlier model selection. In *International
 740 Conference on Automated Machine Learning (AutoML)*.
- 741 Yue Zhao, Ryan Rossi, and Leman Akoglu. 2021. Automatic unsupervised outlier model selection.
 742 *Advances in Neural Information Processing Systems* 34 (2021), 4489–4502.
- 743 Yue Zhao, Sean Zhang, and Leman Akoglu. 2022. Toward unsupervised outlier model selection. In
 744 *2022 IEEE International Conference on Data Mining (ICDM)*. IEEE, 773–782.
- 745 Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Dae ki Cho, and Haifeng
 746 Chen. 2018. Deep Autoencoding Gaussian Mixture Model for Unsupervised Anomaly Detection.
 747 In *International Conference on Learning Representations*.

756 **A ILLUSTRATION OF SYNTHETIC DATA IN 2-d**
 757

758 We visualize our synthetic data in Figure 4, with 3 randomly created 2-d GMMs with the number of
 759 clusters ($N = 1, 2, 3$). We choose the 80th percentile as the criterion, such that inliers are samples
 760 drawn from the GMM and within the 80th percentile, and outliers are samples drawn from the
 761 inflated GMMs and outside of the 80th percentile.



763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 801 802 803 804 805 806 807 808 809

Figure 4: Illustration of synthetic data in 2D with 80th percentile as the criterion.

778 **B LINEAR TRANSFORM FOR SCALABLE GMM DATA SYNTHESIS**
 779

780 **B.1 DEFINITIONS**

782 **Definition 1 (Gaussian Mixture Model)** We denote an m -cluster d -dimension Gaussian Mixture
 783 Model as $\mathcal{G}_m^d = \{(w_j, \mu_j, \Sigma_j)\}_{j=1}^m$, which is the weighted sum of m Gaussian distributions:

$$785 p(\mathbf{x}) = \sum_{j=1}^m w_j \cdot g(\mathbf{x}|\mu_j, \Sigma_j), \quad (6)$$

788 where $w_j \in \mathbb{R}^+$ is the weight for the j -th Gaussian $\mathcal{N}(\mu_j, \Sigma_j)$ with $\sum_{j=1}^m w_j = 1$, and $g(\cdot|\mu_j, \Sigma_j)$
 789 is the density of the j -th component/cluster, with mean/center $\mu_j \in \mathbb{R}^d$ and covariance $\Sigma_j \in \mathbb{R}^{d \times d}$
 790 being positive semi-definite, such that $\mathbf{x}^T \Sigma_j \mathbf{x} \geq 0$, for all $\mathbf{x} \in \mathbb{R}^d$.

792 **Definition 2 (Linear Transform)** We denote a linear transformation T in \mathbb{R}^d as:

$$793 T(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b}, \quad (7)$$

795 where $\mathbf{x} \in \mathbb{R}^d$, and $\mathbf{W} \in \mathbb{R}^{d \times d}$, $\mathbf{b} \in \mathbb{R}^d$ are the parameters of T .

797 **Definition 3 (Mahalanobis Distance)** The Mahalanobis distance dist_M between a point $\mathbf{x} \in \mathbb{R}^d$
 798 and a Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ is defined as:

$$800 \text{dist}_M(\mathbf{x}) = \sqrt{(\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu)}. \quad (8)$$

802 **Definition 4 (χ_d^2 -distribution)** The Chi-squared distribution χ_d^2 with d degrees of freedom is the
 803 distribution of the sum of squares of d independent standard Normal random variables.

804 **B.2 PROPERTIES**

806 **Property B.1 (Lemma 5.3.2 (Casella and Berger, 2024))** If $Z \sim \mathcal{N}(0, 1)$, then $Z^2 \sim \chi_1^2$; If
 807 X_1, \dots, X_d are independent and $X_i \sim \chi_1^2$, then $\sum_{i=1}^d X_i \sim \chi_d^2$.

809 **Property B.2** The squared Mahalanobis distance $\text{dist}_M^2(\mathbf{x}) \sim \chi_d^2$, with $\mathbf{x} \sim \mathcal{N}(\mu, \Sigma)$.

810 *Proof:* If $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, then we have $\mathbf{z} = \Sigma^{-\frac{1}{2}}(\mathbf{x} - \boldsymbol{\mu}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ (Gut, 2009), such that:
 811

$$\text{dist}_M^2(\mathbf{x}) = \mathbf{z}^T \mathbf{z} = \sum_{i=1}^d z_i^2 \quad (9)$$

812 where z_i are independent standard Normal random variables. We have $\sum_{i=1}^d z_i^2 \sim \chi_d^2$ from Prop-
 813 erty B.1, which completes the proof.
 814

815 B.3 LEMMAS

816 **Lemma 1** *Linear transform T with invertible \mathbf{W} on \mathcal{G}_m^d preserves Mahalanobis distances.*

817 *Proof:* We denote the transformed GMM as $T(\mathcal{G}_m^d) = \{(w_j, \mathbf{W}\boldsymbol{\mu}_j + \mathbf{b}, \mathbf{W}\Sigma_j\mathbf{W}^T)\}_{j=1}^m$, then with
 818 $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}_j, \Sigma_j)$, for the transformed point $T(\mathbf{x})$ we have:
 819

$$\text{dist}_M(T(\mathbf{x})) = \sqrt{(T(\mathbf{x}) - (\mathbf{W}\boldsymbol{\mu}_j + \mathbf{b}))^T(\mathbf{W}\Sigma\mathbf{W}^T)^{-1}(T(\mathbf{x}) - (\mathbf{W}\boldsymbol{\mu}_j + \mathbf{b}))} \quad (10)$$

$$= \sqrt{(\mathbf{W}(\mathbf{x} - \boldsymbol{\mu}_j))^T(\mathbf{W}\Sigma\mathbf{W}^T)^{-1}(\mathbf{W}(\mathbf{x} - \boldsymbol{\mu}_j))} \quad (11)$$

$$= \sqrt{(\mathbf{x} - \boldsymbol{\mu}_j)^T \mathbf{W}^T (\mathbf{W}^T)^{-1} \Sigma^{-1} \mathbf{W}^{-1} \mathbf{W} (\mathbf{x} - \boldsymbol{\mu}_j)} \quad (12)$$

$$= \sqrt{(\mathbf{x} - \boldsymbol{\mu}_j)^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}_j)} = \text{dist}_M(\mathbf{x}). \quad (13)$$

■

820 **Lemma 2** *Linear transform T with invertible \mathbf{W} on \mathcal{G}_m^d preserves the percentiles of the GMM.*

821 *Proof:* Let $\chi_d^2(\alpha)$ denote the α -th percentile of χ_d^2 , such that for $X \sim \chi_d^2$:

$$\text{Prob}(X \leq \chi_d^2(n)) = \frac{\alpha}{100}. \quad (14)$$

822 Based on Property B.2, we have $\text{Prob}(\text{dist}_M^2(\mathbf{x}) \leq \chi_d^2(\alpha)) = \frac{\alpha}{100}$.
 823

824 Let $\mathbf{x} \sim \mathcal{G}_m^d$, such that $\text{dist}_M^2(\mathbf{x}) > \chi_d^2(\alpha)$ for all $\mathcal{N}_j(\boldsymbol{\mu}_j, \Sigma_j)$, which indicates that \mathbf{x} is outside the
 825 α -th percentile of \mathcal{G}_m^d . Since $\text{dist}_M(\mathbf{x})$ is preserved under T (see Lemma 1), then we conclude that
 826 the linear transform T with invertible \mathbf{W} preserves the percentiles of the GMM. ■
 827

828 C IMPLEMENTATION DETAILS

829 C.1 HARDWARE

830 We base our experiments on a NVIDIA RTX A6000 GPU with AMD EPYC 7742 64-Core Processors.
 831

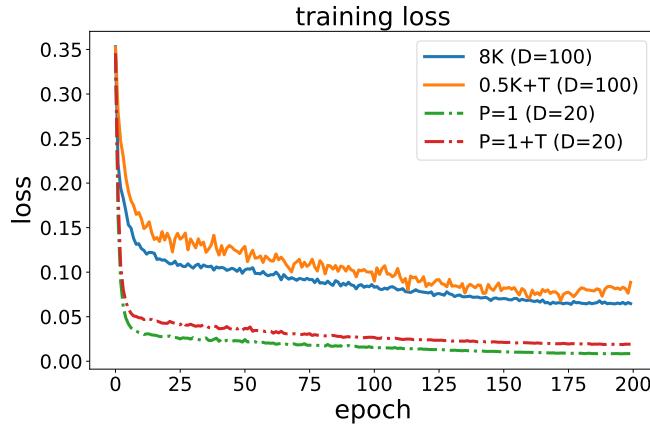
832 C.2 TRAINING AND INFERENCE

833 We train our models for 200 epochs with the Adam optimizer (Kingma and Ba, 2017) and a
 834 learning_rate = 0.001, and test with the model corresponding to the lowest training loss.
 835 The size of our $D = \{20, 100\}$ model is 4.87M and 4.89M parameters, respectively. We show the
 836 training process of PFNs and our model in Algorithm 1.
 837

838 **Dealing with varying dimensions and dataset size** For an input with d features, we follow Müller
 839 et al. (2022) and deal with $d < D$ by rescaling the input with $\frac{D}{d}$ and padding the features to size D
 840 with 0, and randomly sample D features out of d if $d > D$. In addition, FoMo-0D uses context size of
 841 5K at inference, where we randomly sample (5K–1) points as $\mathcal{D}_{\text{train}}$ from datasets with $n > 5K$ for
 842 each test sample $\mathbf{x} \in \mathcal{D}_{\text{test}}$.
 843

844 **Model architecture** We use a 4-layer Transformer with hidden dimension $\text{h_dim} = 256$, a linear
 845 layer ($\mathbb{R}^D \rightarrow \mathbb{R}^{\text{h_dim}}$) as the embedding layer and a 2-layer MLP ($\mathbb{R}^{\text{h_dim}} \rightarrow \mathbb{R}^2$) as the classification
 846 layer for inlier vs. outlier. For each Transformer layer, we use num_head = 4 for each attention
 847 module and $R = 500$ for the router-based attention (Figure 2).
 848

864
865 **Algorithm 1:** Prior-fitting of a PFN (Müller et al., 2022) and ours
866 **Input** : A prior distribution over datasets $p(\mathcal{D})$, from which samples can be drawn and the number of
867 datasets Q to draw for one epoch, the number of training epochs E , the periodicity P , the number
868 of unique datasets q , linear transformation T .
869 **Output** : A model q_θ that will approximate the PPD
870 1 Initialize the neural network q_θ ;
871 2 Initialize the epoch-level collection $\mathcal{C}_E = []$;
872 3 for $i \leftarrow 1$ to E do
873 if $i \leq P$ then
874 Initialize an empty buffer $\mathcal{B}_i = []$;
875 Initialize the dataset-level collection $\mathcal{C}_q = []$;
876 for $j \leftarrow 1$ to Q do
877 if $j \leq q$ then
878 Step 1: sample $D_j := \mathcal{D}_{\text{train}} \cup \{(\mathbf{x}_k, y_k)\}_{i=k}^{|\mathcal{D}_{\text{test}}|} \sim p(\mathcal{D})$;
879 $\mathcal{C}_q \leftarrow \mathcal{C}_q + [D_j]$
880 end
881 else
882 $j \leftarrow j \bmod q$
883 $D_j \leftarrow T(\mathcal{C}_q[j])$
884 end
885 Step 2: compute stochastic loss approximation $\bar{\ell}_\theta = \sum_{k=1}^{|\mathcal{D}_{\text{test}}|} (-\log q_\theta(y_k | \mathbf{x}_k, \mathcal{D}_{\text{train}}))$;
886 Step 3: update parameters θ with stochastic gradient descent on $\nabla_\theta \bar{\ell}_\theta$;
887 $\mathcal{B}_i \leftarrow \mathcal{B}_i + [D_j]$
888 end
889 $\mathcal{C}_E \leftarrow \mathcal{C}_E + [\mathcal{B}_i]$
890 end
891 else
892 $i \leftarrow i \bmod P$
893 $\mathcal{B}_i \leftarrow \mathcal{C}_E[i]$
894 for $j \leftarrow 1$ to Q do
895 $D_j \leftarrow T(\mathcal{B}_i[j])$
896 Perform Step 2 and Step 3
897 end
898 end
899 end
900
901



910 Figure 5: (best in color) Training loss of FoMo-0D ($D = 100$) with 8K unique datasets/epoch (in
911 blue) and using 0.5K unique + 7.5K transformed datasets/epoch (in orange), and FoMo-0D ($D = 20$)
912 with $P = 1$ (in green) and $P = 1$ with transformation (in red) over 200 epochs.

913
914 **Training loss** In Figure 5, we plot the training loss of our $D = 100$ model trained with 8K unique
915 datasets/epoch (denoted as “8K”) versus 0.5K unique + 7.5K transformed datasets/epoch (denoted
916 as “0.5K+T”), together with the $D = 20$ model trained with reuse periodicity $P = 1$ (denoted as
917 “P=1”, reusing the same 8K datasets across epochs) and $P = 1$ with transformation (denoted as

“P=1+T”, transforming the 8K datasets across epochs). Notice that the loss with transformation is slightly higher than no transformation (i.e., $D = 100$, “0.5K+T” vs. “8K”, and $D = 20$, “P=1+T” vs. “P=1”) across all 200 epochs, which is reasonable since the transformed datasets have non-diagonal covariances that make the learning task harder and thus result in a higher training loss. The training losses of FoMo-OD with $D = 100$ are also higher than with $D = 20$ since the subspace OD tasks are harder in higher dimensions.

Inference time Figure 8 (left) showed the inference time of FoMo-OD on CPU, comparing typical attention versus the router-based attention (with $R = 500$ routers) under varying context sizes from 1K to 10K. The time is measured on CPU to clearly showcase the scalability trends; *quadratic* without routers and *linear* with routers.

Figure 6 shows the inference time on GPU. Notice that the time is much lower (in milliseconds), thanks to the Transformer architecture taking advantage of GPU parallelism, while the compute time for attention without routers continues to grow faster than that with routers.

In implementation, FoMo-OD (with $R = 500$ routers) uses inference context size of 5K by default, which takes about 7.7 ms per test sample on average.

D DETAILED EXPERIMENT SETUP

D.1 PRE-TRAINING DATASET SYNTHESIS

During pretraining, we generate unique GMM datasets by first drawing a configuration, including dimensionality $d \in [D]$, number of components $m \in [M]$, centers $\{\mu_j\}_{j=1}^m$ (each $\mu_j \in [-5, 5]^d$) and covariances $\{\Sigma_j\}_{j=1}^m$ ($\text{diag}(\Sigma_j) \in [-5, 5]^d$). We set $M = 5$ and vary $D \in \{20, 100\}$ to study pretraining with relatively small and high dimensional datasets, respectively. We synthesize inliers and outliers as described in Section 3.1.

We then sample $S = 5,000$ points that are within the 90th percentile of the GMM. To synthesize outliers, we “inflate” a *subset* of dimensions by randomly choosing $|\mathcal{K}| \in [D]$ dimensions and multiplying the corresponding variances by $\times 5$ (following (Han et al., 2022)), i.e. $5 \times \Sigma_{j,kk}$ ’s for $k \in \mathcal{K}$, and then draw $\tilde{S} = 5,000$ samples from the inflated GMM that are outside the 90th percentile of the original GMM.

To speed up data synthesis via linear transformations, we first draw 500 unique datasets using $m \in [5]$ and $d \in \{1, 2, \dots, 100\}$ (i.e. 5×100) and transform each one $15 \times$ using varying parameters (\mathbf{W}, \mathbf{b}) as described in Section 3.3.⁶ This yields 8K unique datasets (500 original and 7,500 transformed) to use at one training epoch (over 1,000 steps with batch size $B = 8$). We repeat this process at each epoch, drawing 500 new datasets and transforming them to reach 8K datasets per epoch.

D.2 REAL-WORLD BENCHMARK DATASETS

While pretraining is purely on synthetic datasets, we evaluate FoMo-OD on 57 real-world datasets from the ADBench benchmark (Han et al., 2022) (see Table 15). They consist of 47 popular tabular outlier detection datasets, as well as 10 newly-constructed tabular datasets created from images and natural language tasks by using pretrained models to extract embeddings. We defer to the original paper for the details on these benchmark datasets.

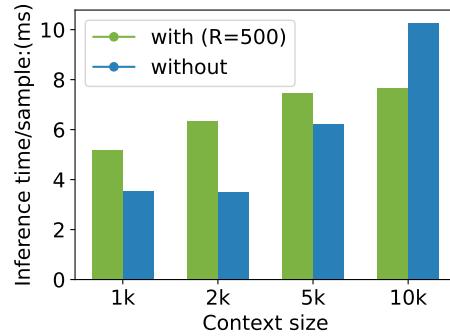


Figure 6: Inference time of FoMo-OD on *GPU* with vs. w/out router-based attention under varying context size.

⁶It is important to ensure that the eigenvalues of \mathbf{W} (i.e. variances) are not too small such that the dataset does not flatten in any direction. To this end, we draw a random orthonormal basis $\mathbf{U} \in [-1, 1]^{d \times d}$ and a diagonal Λ with eigenvalues $\lambda_{kk} \in ([-1, -0.1] \cup [0.1, 1])^d$, and obtain $\mathbf{W} = \mathbf{U}\Lambda\mathbf{U}^T$. We also use $\mathbf{b} \in [-1, 1]^d$.

We compare to DTE (Livernoche et al., 2024) and baselines therein as described next, thus, following their semi-supervised OD setup we split each dataset five times into train/test using five different seeds and report the mean performance and its standard deviation. In particular, each random split designates 50% of the inliers as $\mathcal{D}_{\text{train}}$, while $\mathcal{D}_{\text{test}}$ contains the rest of the inliers and all the outlier samples. Note that while the baseline methods require model re-training and inference for each $\mathcal{D}_{\text{train}}/\mathcal{D}_{\text{test}}$ split, FoMo-0D uses the splits only for inference as $\mathcal{D}_{\text{train}}$ is merely passed as context.

Before passing the datasets as input to FoMo-0D, we perform a quantile transform such that the features follow a Normal distribution, to better align with the pretraining data from GMMs.

D.3 BASELINES

We compare FoMo-0D against **26** baselines, from classical/shallow methods to modern/deep models. Our baselines include all the baselines imported from one of the latest papers that proposed the SOTA diffusion-based model DTE (Livernoche et al., 2024), and its three variants; DTE-C, DTE-IG, and DTE-NP. Their baselines comprise all those in ADBench (Han et al., 2022); both classical ones (k NN (Ramaswamy et al., 2000), LOF (Breunig et al., 2000), iForest (Liu et al., 2008), HBOS (Goldstein and Dengel, 2012), etc.) and deep models (DeepSVDD (Ruff et al., 2018), DAGMM (Zong et al., 2018), DROCC (Goyal et al., 2020), etc.). They also include more recent approaches based on self-supervised learning (GOAD (Bergman and Hoshen, 2020), ICL (Shenkar and Wolf, 2022), SLAD (Xu et al., 2023), etc.), besides the four additional generative baselines: normalizing planar flows (Rezende and Mohamed, 2015), DDPM (Ho et al., 2020), VAE (Kingma, 2013) and GANomaly (Akcay et al., 2019). We defer to the original paper for additional details. Overall, our 26 baselines consist of the most recent, SOTA approaches for OD that span a diverse family (nonparametric, self-supervised, generative, etc.).

D.4 HYPERPARAMETERS FOR BASELINES

Table 4 gives the list of HP values we used to study the HP sensitivity/performance variability of the (from top to bottom) top-4 baselines.

Table 4: Top-4 baselines (from top to bottom) and hyperparameter (HP) configurations.

Baseline	Hyperparameters
DTE-NP	$k \in \{5, 10, 20, 40, 50\}$
k NN	$k \in \{5, 10, 20, 40, 50\}$
ICL	$\text{learning_rate} \in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$
DTE-C	$k \in \{5, 10, 20, 40, 50\}$

D.5 RANKING THE 26 BASELINES

Figure 16 presents the visualization of the p -values of the pairwise Wilcoxon signed rank test w.r.t. AUROC among the baseline methods used by Livernoche et al. (2024). We rank these 26 baselines based on their mean p -value (i.e., row-wise average) against the other baselines.

D.6 COMPARISON OF TOP-4 BASELINE VARIANTS WITH VARYING HP CONFIGURATIONS

Figure 17, 18, 19, 20 give the p -values, respectively comparing the variants of the top-4 baselines (DTE-NP, k NN, ICL, DTE-C) among themselves using different HP configurations, as well as the $^{\text{avg}}$ model with the average performance across HPs. (Specifically for ICL, learning_rate (lr) $\in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$; and for others, #nearest-neighbors $k \in \{5, 10, 20, 40, 50\}$). We find that for ICL, $\text{lr} = 10^{-3}$ or 10^{-4} are preferable while those that are too small or too large perform poorly. For others, small $k \in \{5, 10\}$ tend to outperform larger $k \in \{40, 50\}$. Note that Livernoche et al. (2024) used $k = 5$ in their paper that proposed DTE (and variants) as well as the k NN baseline for fair comparison, while the DTE^{avg} and $k\text{NN}^{\text{avg}}$ models across HP configurations perform subpar.

1026 D.7 SAMPLING TIME OF d -DIMENSIONAL GMM
1027

1028 Figure 7 shows the sampling time of drawing 10,000
 1029 points from different GMMS with increasing dimen-
 1030 sionality $d = \{10, 20, \dots, 200\}$. We parallelize the
 1031 sampling process over 10 CPUs, where each CPU
 1032 draws 1000 samples.

1033 We observe that the sampling time grows nonlinearly
 1034 as the number of dimensions increases, which sug-
 1035 gests that it may incur considerable computational
 1036 overhead to directly draw from the data prior over
 1037 hundreds of thousands of training steps, motivating
 1038 the use of our proposed on-the-fly linear transforma-
 1039 tion T for scalability.

1040

1041
1042

E ABLATION ANALYSES

1043 In this section, we perform various ablations to study the effect of different design choices in FoMo-
 1044 0D; namely, **E.1** maximum pretraining data dimensionality D , the number of routers R on **E.2** cost
 1045 and **E.3** performance, **E.4** context size (both for training and inference), **E.5** number of unique
 1046 datasets used for pretraining (i.e., reuse periodicity P), data transformation T during synthesis on **E.6**
 1047 performance and **E.7** speed up, **E.8** data diversity and prolonged training, and finally, **E.9** quantile
 1048 transforming the benchmark datasets preceding inference.

1049 Unless stated otherwise, most ablation results are performed using FoMo-0D with $D = 20$, as it is
 1050 faster to pretrain under these many varying settings.

1052

1053
1054E.1 EFFECT OF PRETRAINING DIMENSIONALITY D

1055 *How does FoMo-0D’s generalization performance change by increasing dimensionality of the
 1056 pretraining data?*

1057 We start by comparing FoMo-0D pretrained on datasets with up to $D = 20$ versus $D = 100$
 1058 dimensions. Note that learning on higher dimensional datasets is harder, as evident from the relatively
 1059 larger pretraining loss as shown in Appendix Figure 5. While the statement is accurate in general, it
 1060 is also partly because subspace outliers “hide” better in higher dimensions.

1061 Comparing Table 1 ($D = 100$) with Table 2 ($D = 20$) w.r.t. p -values over All datasets, we find that
 1062 FoMo-0D at larger scale does better, where **all** p -values are larger for $D = 100$ than $D = 20$. We
 1063 find that FoMo-0D with $D = 20$ performs well on datasets with $d \leq 20$ (i.e., “on its own game”),
 1064 however beyond its pretraining setting, e.g. on datasets with $d \leq 50$, $D = 100$ is superior to $D = 20$
 1065 as shown in Appendix Table 8.

1066

1067
1068

E.2 EFFECT OF ROUTERS ON COST

1069 *What is the running time and memory cost of FoMo-0D with & w/out router-based attention?*

1070 Figure 8(left) shows the average inference time per test sample, comparing FoMo-0D using a router-
 1071 based attention mechanism with $R = 500$ routers (in green) versus FoMo-0D using typical attention
 1072 without any routers (in blue). As inference context size increases, running time for traditional
 1073 attention grows quadratically while router mechanism scales linearly.⁷

1074 Similarly, memory cost with routers is considerably lower when using routers, especially for larger
 1075 context sizes, as shown in Figure 8(middle).

1077

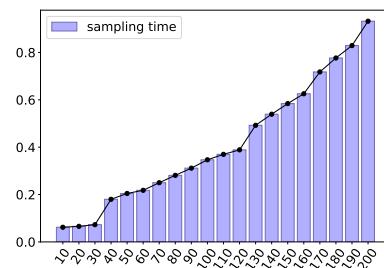
1078
1079

Figure 7: Sampling time (in seconds) of 10,000 points from GMMS with varying number of dimensions.

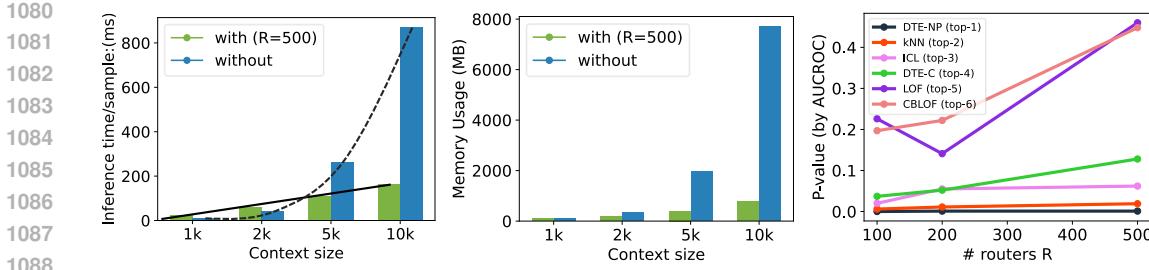


Figure 8: FoMo-0D w/ router mechanism saves time and memory while more #routers perform better, offering a cost-performance trade-off: (left) inference-time (ms) per sample and (middle) memory cost (MB) with & w/out routers by varying context size; (right) performance (based on p -value against top baselines, higher is better) vs. number of routers. (setting: $D = 20$, $P = 1$)

E.3 EFFECT OF ROUTERS ON PERFORMANCE

What is the impact of the number R of routers (or representatives) on performance?

Router-based mechanism allows to trade-off running time with expressiveness of the attention and hence performance. Figure 8(right) shows the p -values of the Wilcoxon signed rank test as the number of routers R is increased from 100 to 200 and 500, comparing FoMo-0D to each of the top-6 baselines. We notice that FoMo-0D performance tends to increase monotonically with more routers.

E.4 EFFECT OF CONTEXT SIZE

What is the impact of context size, both during model pretraining as well as during inference?

To study how performance changes by context size, we train FoMo-0D with varying context size in $\{1K, 2K, 5K\}$ and employ each pretrained model for inference with varying context size in $\{1K, 2K, 5K, 10K\}$. Table 5 shows the results, where performance is depicted by the average rank of FoMo-0D (the lower, the better).

Table 5: Average rank (based on comparison to 30 baselines w.r.t. AUROC) of FoMo-0D across datasets under different context sizes for training and inference. Smaller ranks imply better performance. (setting: $D = 20$, $R = 500$, $P = 1$)

	Infer:1K	Infer:2K	Infer:5K	Infer:10K
Train:1K	13.816	14.623	15.193	15.439
Train:2K	13.079	13.219	13.439	13.561
Train:5K	13.088	13.211	13.307	13.430

We find that training with a larger context improves performance at any inference context size. On the other hand, perhaps counter-intuitively, FoMo-0D with smaller inference context size does better. We conjecture that is because the #routers-to-context size ratio increases with a larger context size at inference, limiting the expressive power of the “bottleneck” attention mechanism. The pairwise statistical tests among the $3 \times 4 = 12$ models support these observations, as shown in Figure 9. Interestingly, when training context size is large enough at 5K, inference with 10K samples generalizes beyond training with no significant difference (at 0.05) from other inference context sizes.

E.5 EFFECT OF NUMBER OF UNIQUE DATASETS

How do FoMo-0D performances compare when pretrained on unique vs. reused datasets, via varying periodicity P ?

Next we study the effect of dataset reuse at epoch level (w/out transformation) on performance as presented in Section 3.3. We vary reuse periodicity P in $\{1, 50, 100\}$, and accordingly, increase the number of unique datasets used for pretraining across epochs. As shown in Table 6, FoMo-0D ($w/D = 20$) performs similarly with varying dataset reuse. In fact, it is competitive even with $P = 1$, remaining no different from the 3rd best baseline (ICL) across All (57) datasets, while significantly outperforming the top 5th (LOF) across (24) datasets with $d \leq 20$ as well as (38) with $d \leq 50$.

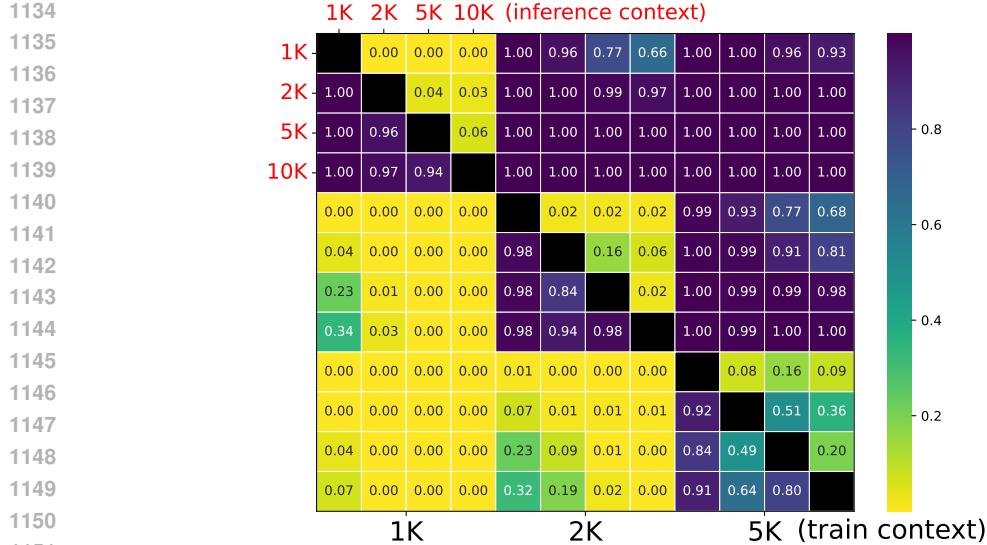


Figure 9: p -values of the pairwise Wilcoxon signed rank test between models (larger p implies col-method is better than row-method) w/ different context sizes for **training** (1K/2K/5K, 1st/2nd/3rd four grids, in **black**) and **inference** (1K/2K/5K/10K, every 1st/2nd/3rd/4th grid, in **red**): Larger training context improves overall performance, while smaller inference context is preferable.

Table 6: Ablation results on dataset reuse across epochs with varying $P \in \{1, 50, 100\}$ show stable p -values against the top-5 baselines, where FoMo-0D with $D = 20$ remains no different from the top 3rd baseline at 0.05 w.r.t. pairwise Wilcoxon signed rank test comparisons, while it continues to significantly outperform the top 5th baseline (LOF) when $d \leq 50$. (setting: $D=20$, $R=500$, context size=5K, w/out transformation T)

	$P = 1$ (#unique datasets: 8K)				$P = 50$ (#unique datasets: $8 \times 50 = 400$ K)				$P = 100$ (#unique datasets: $8 \times 100 = 800$ K)							
	DTE-NP	kNN	ICL	DTE-C	LOF	DTE-NP	kNN	ICL	DTE-C	LOF	DTE-NP	kNN	ICL	DTE-C	LOF	
All	0.001	0.019	0.062	0.128	0.460	0.001	0.019	0.089	0.159	0.394	0.001	0.015	0.072	0.121	0.290	
$d \leq 20$	0.583	0.755	0.943	0.736	0.998	0.572	0.789	0.968	0.616	0.993	0.439	0.678	0.953	0.550	0.972	
$d \leq 50$	0.415	0.750	0.869	0.962	0.999	0.347	0.794	0.893	0.946	0.997	0.293	0.697	0.890	0.924	0.994	

E.6 EFFECT OF TRANSFORMATION T FOR SYNTHESIS

How do FoMo-0D performances compare when pretrained on datasets with vs. w/out linear transformation?

Setting $P = 1$, we next study the impact of linear transformation T . Table 7 presents the results, where we compare reuse of the *same* 8K unique datasets across epochs (w/out T), versus *transforming* these datasets with T at every epoch with different parameters (w/ T). FoMo-0D performance remains stable; no different from the top 3rd model on All datasets, while significantly outperforming the top 5th across those with $d \leq 20$ and $d \leq 50$. This suggests that T can be employed without sacrificing performance to save time during pretraining.

Table 7: Ablation results on performance w/ & w/out linear transformation T show stable p -values against the top-5 baselines, where FoMo-0D with $D = 20$ remains no different from the top 3rd baseline at 0.05 w.r.t. pairwise Wilcoxon signed rank test comparisons. (setting: $D = 20$, $R = 500$, context size=5K, $P = 1$)

	w/out transformation T				w/ transformation T					
	DTE-NP	kNN	ICL	DTE-C	LOF	DTE-NP	kNN	ICL	DTE-C	LOF
All	0.001	0.019	0.062	0.128	0.460	0.002	0.015	0.226	0.210	0.280
$d \leq 20$	0.583	0.755	0.943	0.736	0.998	0.648	0.708	0.988	0.718	0.955
$d \leq 50$	0.415	0.750	0.869	0.962	0.999	0.264	0.382	0.971	0.900	0.963

1188
1189E.7 SPEED UP BY T 1190
1191
1192**What is the time saving on data synthesis with linear transformation?**1193
1194
1195
1196
1197
1198

Figure 10 shows the distribution of pretraining running-time per epoch with and w/out data transformation. Specifically, we compare (left) generating 8K unique datasets/epoch on-the-fly and (right) first generating 500 unique datasets on-the-fly and then transforming each one 15 times using T with different parameters to reach 8K datasets at each epoch.

1199
1200
1201
1202
1203
1204
1205
1206
1207
1208

Notice that pretraining with T takes about 450 sec./epoch on average, while without T it requires 1200 sec./epoch to generate 8K unique datasets and gradient descent across 1000 steps. Different from other ablation results, which are based on the $D = 20$ model, here we report the running times for our $D = 100$ model. Overall, our final FoMo-0D took ≈ 25 hours for pre-training (450 sec. \times 200 epochs). Importantly, this is a one-time cost that amortizes across many downstream tasks with as low as **7.7 ms inference time** per test sample (see Table 3 and Appendix Figure 6).

1209
1210

E.8 EFFECT OF DATA DIVERSITY AND PROLONGED TRAINING

1211
1212**How does FoMo-0D’s performance change by increasing pre-training data diversity and number of training epochs?**1213
1214
1215
1216
1217

Originally we have trained FoMo-0D w/ $D = 100$ using 0.5K unique + 7.5K transformed datasets over 200 epochs. As mentioned earlier, learning in higher dimensions tends to incur a larger loss in general but also specifically here, as subspace outliers are harder to detect in high dimensions.

1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232

Toward reducing the loss further, we resume the pretraining for another 100 epochs. Further, to simplify the tasks and thereby increase data diversity, we also decrease the inlier/outlier labeling percentile threshold from 90% to 80% during on-the-fly data generation in the last 100 epochs. In Figure 12, we present the training loss of FoMo-0D ($D = 100$) trained with 0.5K unique + 7.5K transformed datasets/epoch over 200 epochs (90th percentile as labeling threshold) and then 100 additional epochs (80th percentile as the threshold) to show how data diversity and amount affect model performance. Figure 11 compares FoMo-0D’s performance (w/ $D = 100$) to top-5 baselines w.r.t. p -values of the paired Wilcoxon signed rank test on datasets with $d \leq 100$, after the first 200 epochs versus after 300 epochs. The increase in all the p -values showcases the benefit of additional training.

1233
1234

E.9 EFFECT OF APPLYING QUANTILE TRANSFORM ON BENCHMARK DATASETS

What is the impact of quantile data transform preceding inference on performance?1235
1236
1237
1238
1239

We pretrain FoMo-0D on synthetic datasets from a simple data prior based on GMMs. The real-world benchmark datasets, on the other hand, may exhibit features with distributions different from Gaussians. To close the gap, we apply a quantile transform (denoted QT) on the benchmark datasets prior to feeding them to FoMo-0D for inference, which transforms the features to exhibit a more Gaussian-like probability distribution.

1240
1241

Figure 13 compares the performance of three FoMo-0D w/ $D = 100$ variants with and w/out QT against the top-5 baselines w.r.t. the p -values of the paired Wilcoxon signed rank test. FoMo-0D tends to perform better as suggested by larger p -values when QT is applied.

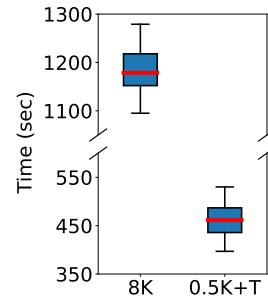


Figure 10: Runtime/epoch dist.n over 100 epochs for FoMo-0D ($D=100$) with (left) $P=100$, i.e. 8K unique datasets/epoch vs. (right) 0.5K unique+7.5K transformed datasets/epoch.

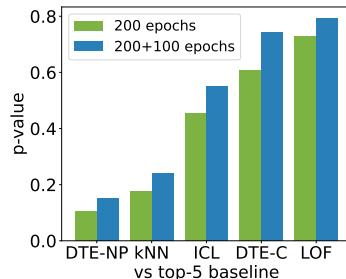
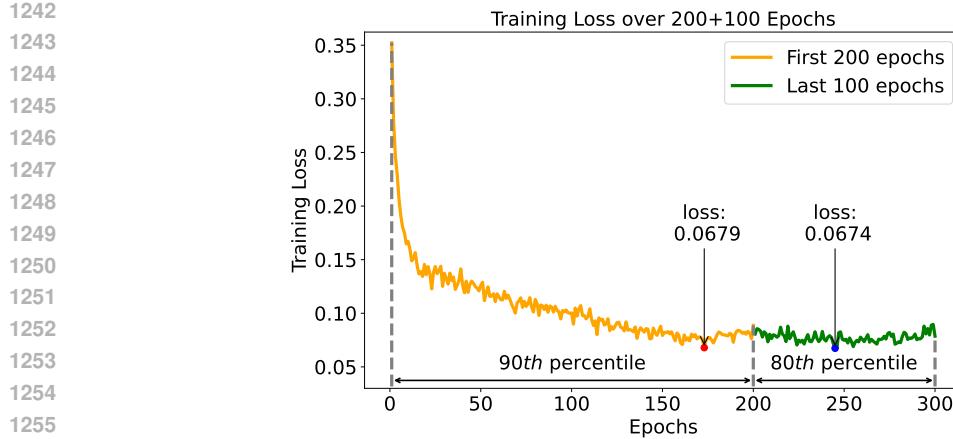


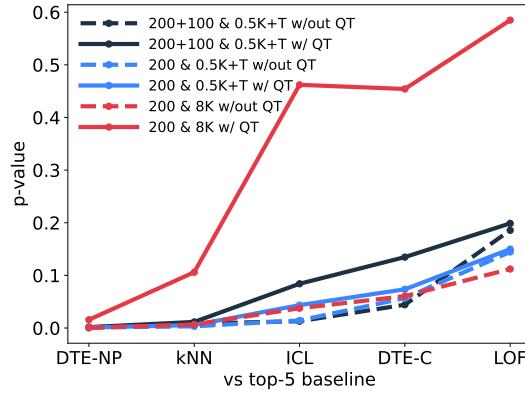
Figure 11: p -values increase with additional 100 epochs of pretraining, i.e. FoMo-0D w/ $D = 100$ performs better against top-5 baselines w.r.t. p -values of the paired Wilcoxon signed rank test on datasets with $d \leq 100$.

Figure 11 compares FoMo-0D’s performance (w/ $D = 100$) to top-5 baselines w.r.t. p -values of the paired Wilcoxon signed rank test on datasets with $d \leq 100$, after the first 200 epochs versus after 300 epochs. The increase in all the p -values showcases the benefit of additional training.



1256
1257
1258
1259

Figure 12: (best in color) Training loss of FoMo-0D ($D = 100$) with 0.5K unique + 7.5K transformed datasets/epoch for 200 epochs (in orange), followed with additional 100 epochs of training (in green). For the first 200 epochs we train with 90th percentile as the inlier/outlier threshold, which we reduce to 80th in the subsequent 100 epochs.



1274
1275
1276

Figure 13: p -values increase, i.e. FoMo-0D performance improves, against top-5 baselines with quantile transform (QT) preceding inference, for 3 different settings of FoMo-0D w/ $D = 100$.

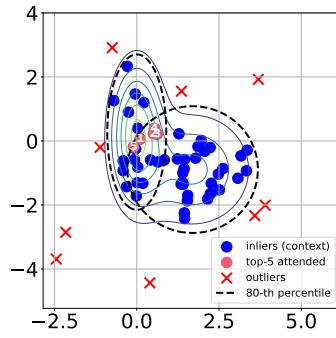
1277
1278
1279
1280

Besides the ablation studies, we provide a qualitative case study of sample-to-sample attention in Appendix F, showing that an outlier attends to the points in context that are within a short distance significantly more than random points, suggesting that PFNs tend to mimic non-parametrics.

F QUALITATIVE ANALYSIS ON SAMPLE-TO-SAMPLE ATTENTION

1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

We sample 50 inliers as context and 100 outliers from a 2-d GMM using the 80th percentile as the labeling threshold, and visualize the top 5 inliers most attended by the 100 outliers based on the average (cross) attention weights over 4 heads from the last layer of FoMo-0D ($D = 100$), which accurately labeled all the 100 outliers. In Figure 14, the most frequently attended inliers are close to either the center of a Gaussian (e.g., 1st, 5th) or the criterion (e.g., 3rd, 4th), suggesting FoMo-0D tends to learn decision boundaries that reflect the prior data generation process. For each outlier, we compute the sum of L2 distances to its top-5 attended inliers (att), the sum of L2 distances to 5 randomly chosen inliers (rcdm), and the sum of L2 distances to top-5 inliers with highest likelihood under the GMM (prob). We perform



24

Figure 14: Top-5 attended inliers (all 50 inliers and only part of the outliers are shown for better visualization).

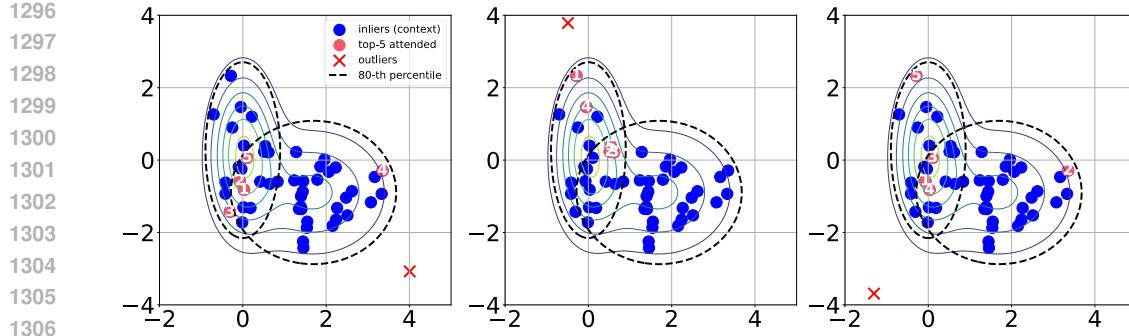


Figure 15: Top-5 attended inliers of 3 outliers at different position of the GMM

Wilcoxon signed rank test between `att` and `rdm` (alternative “less”), `att` and `prob` (alternative “greater”) over all the outliers, with a p -value of 4.4×10^{-4} and 0.99, respectively, suggesting the distances based on attention weights are significantly less than the random distances, and **not** significantly greater than the distances to inliers in high probability region.

We visualize the top-5 attended inliers for 3 outliers at different position of the 2-d GMM in Figure 15. For a specific outlier, there is a similar trend of attending to the center of a Gaussian (as shown in Figure 14), besides, inliers that reflect the criterion boundary or are close to the outlier are actively attended (e.g., 3rd, 4th in the left, 1st in the middle, 2nd, 5th in the right), suggesting FoMo-0D is incorporating both boundary and nearest neighbor information dynamically for each outlier.

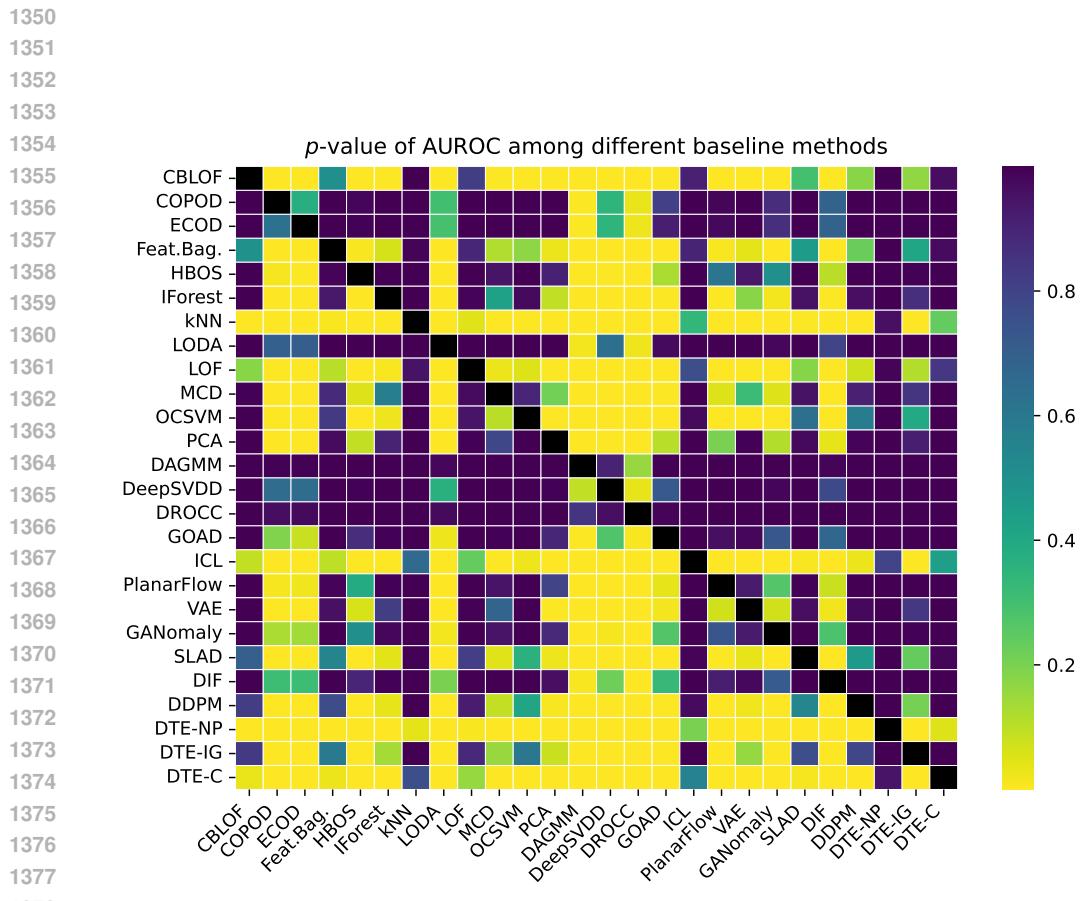
G FULL RESULTS

Tables 9.1& 9.2, 10.1 & 10.2, and 11.1 & 11.2 respectively show the AUROC, AUPR and F1 scores of the top-4 baselines, DTE-NP, k NN, ICL, and DTE-C as well as their corresponding ^{avg} model with the average performance across HPs, as listed in Table 4.

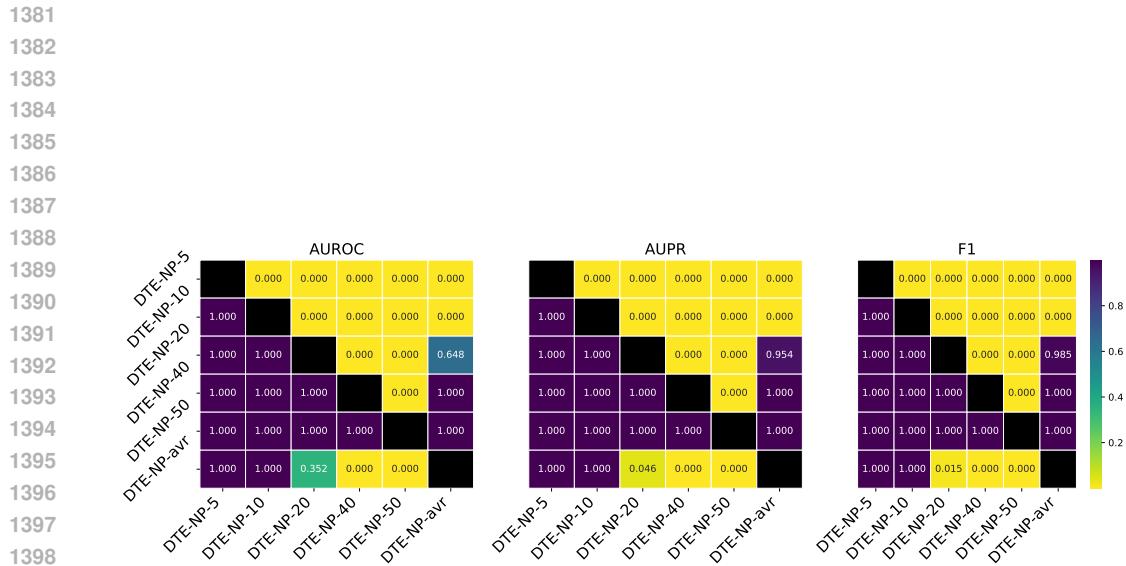
Tables 12.1&12.2, 13.1&13.2, and 14.1&14.2 respectively show the AUROC, AUPR and F1 scores of all methods across all benchmark datasets. In all these tables, the last four rows show the avg_rank of methods across datasets, and p -values of the Wilcoxon signed rank test comparing FoMo-0D w/ $D = 100$ with other baselines. The preceding four rows are the same for FoMo-0D w/ $D = 20$, when ranking 31 models (26 baselines + 4 ^{avg} variants of top-4 baselines + FoMo-0D w/ $D = 20$).

Table 8: Comparison of methods across datasets. (top row) Rank w.r.t. AUROC performance avg.’ed over 57 datasets is presented for FoMo-0D (with $D = 100$), **top-10 baselines** with default HPs, and **top-4⁵** baselines with performance avg.’ed over varying HPs (denoted w/ ^{avg}); followed by p -values of the pairwise Wilcoxon signed rank test, comparing FoMo-0D to each baseline (from top to bottom) over All (57) datasets, those (24) w/ $d \leq 20$, (38) w/ $d \leq 50$, (42) w/ $d \leq 100$ and (46) datasets w/ $d \leq 500$ dimensions. FoMo-0D performs as well as (**i.e., statistically no different from**) the **2nd best model** (k NN, w/ $p = 0.106$) across All datasets, while it is **comparable to** ($p > 0.05$) or **better than** ($p > 0.95$) **all baselines** over datasets w/ $d \leq 100$ (aligned w/ pretraining where $D = 100$) and $d \leq 500$ (generalizing beyond pretraining).

	FoMo-0D	DTE-NP	k NN	ICL	DTE-C	LOF	CBLOF	Feat.Bag.	SLAD	DDPM	OCSVM	DTE-NP ^{avg}	k NN ^{avg}	ICL ^{avg}	DTE-C ^{avg}
Rank(avg)	11.886	7.553	9.018	10.851	11.36	12.316	13.342	13.386	12.982	14.061	13.851	9.079	11.105	12.991	22.263
All	-	0.016	0.106	0.462	0.454	0.585	0.750	0.823	0.759	0.901	0.895	0.112	0.315	0.670	1.000
$d \leq 20$	-	0.428	0.665	0.987	0.727	0.911	0.940	0.987	0.868	0.758	0.968	0.781	0.868	0.990	1.000
$d \leq 50$	-	0.734	0.923	0.992	0.973	0.989	0.987	0.999	0.948	0.985	0.986	0.948	0.967	0.989	1.000
$d \leq 100$	-	0.415	0.700	0.949	0.953	0.970	0.971	0.996	0.876	0.980	0.978	0.752	0.860	0.958	1.000
$d \leq 200$	-	0.315	0.605	0.923	0.919	0.944	0.977	0.990	0.904	0.970	0.983	0.663	0.789	0.937	1.000
$d \leq 500$	-	0.220	0.569	0.827	0.894	0.960	0.968	0.994	0.910	0.960	0.979	0.607	0.756	0.846	1.000



1379 Figure 16: Pairwise *p*-values among baseline methods based on the Wilcoxon signed rank test w.r.t.
1380 AUROC performances across datasets.



1399 Figure 17: *p*-values w.r.t. AUROC/AUPR/F1 among different HP configurations of DTE-NP (i.e.,
1400 $k \in \{5, 10, 20, 40, 50\}$), along with the avg model with the average performance across HPs.

1401
1402
1403

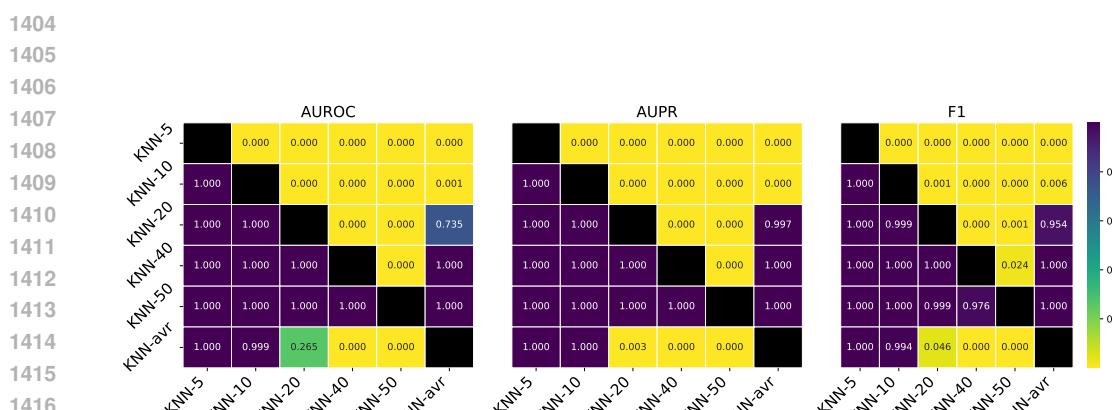


Figure 18: p -values w.r.t. AUROC/AUPR/F1 among different HP configurations of $k\text{NN}$ (i.e., $k \in \{5, 10, 20, 40, 50\}$), along with the $^{\text{avg}}$ model with the average performance across HPs.

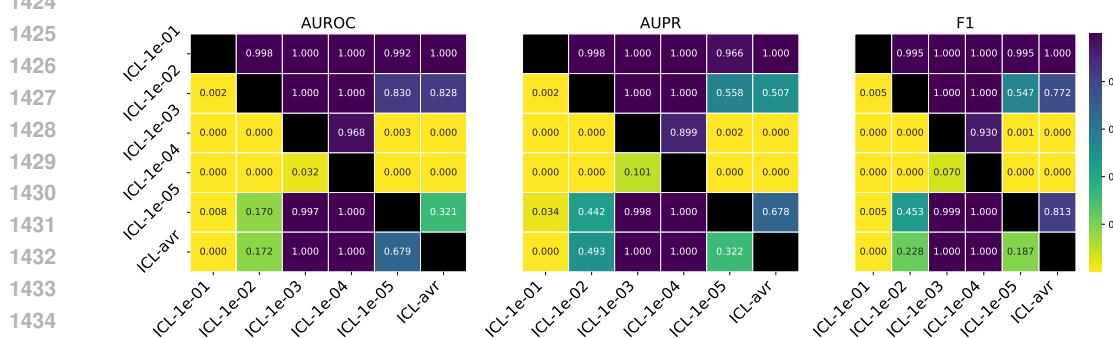


Figure 19: p -values w.r.t. AUROC/AUPR/F1 among different HP configurations of ICL (i.e., $\text{learning_rate} \in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$), along with the $^{\text{avg}}$ model with the average performance across HPs.

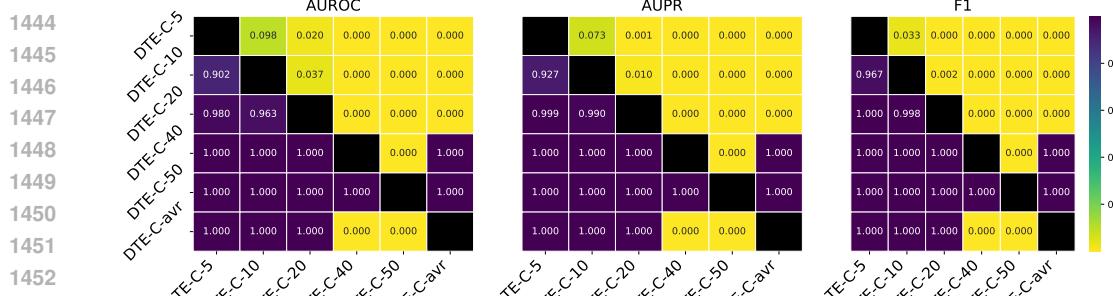


Figure 20: p -values w.r.t. AUROC/AUPR/F1 among different HP configurations of DTE-C (i.e., $k \in \{5, 10, 20, 40, 50\}$), along with the $^{\text{avg}}$ model with the average performance across HPs.

1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

Table 9.1: Average AUROC \pm standard dev. over five seeds for the semi-supervised setting of DTE-NP, k NN with varying hyperparameter (HP) values; $k \in \{5, 10, 20, 40, 50\}$. Also reported is the avg model. We use **bold** and underline respectively to mark the **best** and the **worst** performance of each model to showcase the variability of performance across different HP settings.

dataset	DTE-NP-5	DTE-NP-10	DTE-NP-20	DTE-NP-40	DTE-NP-50	DTE-NP-avr	<u>DTE-NP-avr</u>	KNN-5	KNN-10	KNN-20	KNN-40	KNN-50	KNN-avr	
alo	50.69 \pm 0.00	51.02 \pm 0.00	51.26 \pm 0.00	51.58 \pm 0.00	51.69 \pm 0.00	51.25 \pm 0.00	51.04 \pm 0.00	51.33 \pm 0.00	51.63 \pm 0.00	51.97 \pm 0.00	52.08 \pm 0.00	51.61 \pm 0.00	51.61 \pm 0.00	
amazon	60.76 \pm 0.00	60.69 \pm 0.00	60.53 \pm 0.00	60.47 \pm 0.00	60.47 \pm 0.00	60.58 \pm 0.00	60.52 \pm 0.00	60.62 \pm 0.00	60.73 \pm 0.00	60.91 \pm 0.00	60.25 \pm 0.00	60.25 \pm 0.00	60.25 \pm 0.00	
amnthyroid	93.01 \pm 0.00	92.89 \pm 0.00	92.66 \pm 0.00	92.38 \pm 0.00	92.26 \pm 0.00	92.64 \pm 0.00	92.81 \pm 0.00	92.60 \pm 0.00	92.34 \pm 0.00	91.98 \pm 0.00	91.29 \pm 0.00	91.00 \pm 0.00	91.00 \pm 0.00	
backdoor	94.48 \pm 0.42	93.72 \pm 0.46	92.67 \pm 0.46	91.20 \pm 0.45	90.68 \pm 0.45	92.55 \pm 0.44	93.71 \pm 0.46	92.58 \pm 0.46	91.14 \pm 0.46	89.18 \pm 0.47	88.39 \pm 0.51	91.00 \pm 0.47	91.00 \pm 0.47	
breastw	99.49 \pm 0.28	98.91 \pm 0.35	98.59 \pm 0.34	98.40 \pm 0.35	98.36 \pm 0.38	98.67 \pm 0.35	99.09 \pm 0.24	99.11 \pm 0.27	99.16 \pm 0.22	99.21 \pm 0.18	99.21 \pm 0.17	99.16 \pm 0.21	99.16 \pm 0.21	99.16 \pm 0.21
campaign	78.34 \pm 0.00	78.71 \pm 0.00	78.91 \pm 0.00	78.93 \pm 0.00	78.90 \pm 0.00	78.76 \pm 0.00	78.48 \pm 0.00	78.74 \pm 0.00	78.65 \pm 0.00	78.66 \pm 0.00	78.66 \pm 0.00	78.66 \pm 0.00	78.66 \pm 0.00	78.66 \pm 0.00
cardio	62.53 \pm 0.00	62.03 \pm 0.00	62.14 \pm 0.00	62.14 \pm 0.00	62.14 \pm 0.00	62.46 \pm 0.00	62.28 \pm 0.00	62.47 \pm 0.00	62.44 \pm 0.00	62.49 \pm 0.00	62.07 \pm 0.00	62.07 \pm 0.00	62.07 \pm 0.00	62.07 \pm 0.00
cardiotocography	70.39 \pm 0.33	72.58 \pm 0.26	74.81 \pm 0.34	76.87 \pm 0.38	77.47 \pm 0.37	75.24 \pm 0.29	75.29 \pm 0.37	75.24 \pm 0.40	75.24 \pm 0.40	75.14 \pm 0.38	75.68 \pm 0.35	75.68 \pm 0.35	75.68 \pm 0.35	75.68 \pm 0.35
census	72.18 \pm 0.34	72.34 \pm 0.17	72.28 \pm 0.10	71.93 \pm 0.17	71.80 \pm 0.17	71.93 \pm 0.17	72.23 \pm 0.29	72.23 \pm 0.29	72.36 \pm 0.12	71.94 \pm 0.19	71.37 \pm 0.16	71.28 \pm 0.16	71.28 \pm 0.15	71.28 \pm 0.15
cover	97.90 \pm 0.17	97.72 \pm 0.14	97.40 \pm 0.18	96.84 \pm 0.24	97.37 \pm 0.19	97.51 \pm 0.15	97.51 \pm 0.15	97.19 \pm 0.15	97.40 \pm 0.15	98.85 \pm 0.10	98.20 \pm 0.14	98.73 \pm 0.09	98.73 \pm 0.09	98.73 \pm 0.09
donors	58.34 \pm 0.00	58.61 \pm 0.03	99.43 \pm 0.06	99.14 \pm 0.09	99.02 \pm 0.10	99.38 \pm 0.06	99.51 \pm 0.06	99.24 \pm 0.08	99.52 \pm 0.08	99.52 \pm 0.08	99.52 \pm 0.08	99.52 \pm 0.08	99.52 \pm 0.08	99.52 \pm 0.08
fraud	95.70 \pm 0.90	95.67 \pm 0.93	95.64 \pm 0.93	95.60 \pm 0.92	95.60 \pm 0.92	95.64 \pm 0.92	95.59 \pm 0.97	95.55 \pm 0.99	95.55 \pm 0.99	95.54 \pm 0.92	95.52 \pm 0.88	95.52 \pm 0.88	95.52 \pm 0.88	95.52 \pm 0.88
glass	96.08 \pm 0.39	93.04 \pm 1.06	89.82 \pm 1.12	87.89 \pm 1.10	87.31 \pm 1.40	90.82 \pm 0.91	92.13 \pm 0.94	88.67 \pm 0.98	87.24 \pm 1.18	84.95 \pm 2.22	83.55 \pm 2.61	87.30 \pm 1.59	87.30 \pm 1.59	87.30 \pm 1.59
hepatitis	99.34 \pm 0.20	93.27 \pm 0.51	98.89 \pm 0.96	93.15 \pm 0.99	93.15 \pm 0.99	93.15 \pm 0.99	93.77 \pm 1.47	86.88 \pm 2.21	86.88 \pm 2.21	85.46 \pm 2.34	84.88 \pm 2.09	87.90 \pm 1.75	87.90 \pm 1.75	87.90 \pm 1.75
hip	99.99 \pm 0.00	99.98 \pm 0.01	99.95 \pm 0.00	99.95 \pm 0.01	99.95 \pm 0.01	100.00 \pm 0.01	100.00 \pm 0.01	99.99 \pm 0.02	99.95 \pm 0.01	99.95 \pm 0.01	99.95 \pm 0.01	99.95 \pm 0.01	99.95 \pm 0.01	99.95 \pm 0.01
imdb	50.48 \pm 0.00	50.38 \pm 0.00	50.32 \pm 0.00	50.28 \pm 0.00	50.27 \pm 0.00	50.32 \pm 0.00	50.39 \pm 0.00	50.30 \pm 0.00	50.30 \pm 0.00	50.30 \pm 0.00	50.23 \pm 0.00	50.18 \pm 0.00	50.18 \pm 0.00	50.18 \pm 0.00
interneads	70.96 \pm 0.00	68.65 \pm 0.00	68.86 \pm 0.00	68.97 \pm 0.00	68.97 \pm 0.00	68.97 \pm 0.00	69.52 \pm 0.00	67.65 \pm 0.00	68.08 \pm 0.00	65.48 \pm 0.00	65.04 \pm 0.00	65.04 \pm 0.00	65.04 \pm 0.00	65.04 \pm 0.00
ionosphere	98.48 \pm 0.60	98.13 \pm 0.74	97.84 \pm 0.64	97.83 \pm 0.71	97.83 \pm 0.71	97.83 \pm 0.71	97.83 \pm 0.85	97.62 \pm 0.81	96.33 \pm 0.76	95.12 \pm 0.64	95.12 \pm 0.64	95.12 \pm 0.64	95.12 \pm 0.64	95.12 \pm 0.64
landsat	68.99 \pm 0.60	68.62 \pm 0.60	68.46 \pm 0.60	68.46 \pm 0.60	68.46 \pm 0.60	68.46 \pm 0.60	66.47 \pm 0.00	66.47 \pm 0.00	66.48 \pm 0.00	64.36 \pm 0.00	64.29 \pm 0.00	64.29 \pm 0.00	64.29 \pm 0.00	64.29 \pm 0.00
letter	36.12 \pm 0.00	35.66 \pm 0.00	34.78 \pm 0.00	33.72 \pm 0.00	33.72 \pm 0.00	33.72 \pm 0.00	34.17 \pm 0.00	34.53 \pm 0.00	34.54 \pm 0.00	34.54 \pm 0.00	34.54 \pm 0.00	34.54 \pm 0.00	34.54 \pm 0.00	34.54 \pm 0.00
lymphography	99.88 \pm 0.25	99.79 \pm 0.32	99.76 \pm 0.31	99.76 \pm 0.31	99.76 \pm 0.31	99.80 \pm 0.30	99.87 \pm 0.10	99.88 \pm 0.05	99.85 \pm 0.05	99.85 \pm 0.05	99.88 \pm 0.08	99.88 \pm 0.08	99.88 \pm 0.08	99.88 \pm 0.08
magic gamma	83.91 \pm 0.00	82.87 \pm 0.00	82.87 \pm 0.00	82.87 \pm 0.00	82.87 \pm 0.00	82.73 \pm 0.00	82.81 \pm 0.00	82.73 \pm 0.00	82.21 \pm 0.00	81.85 \pm 0.00	81.85 \pm 0.00	81.85 \pm 0.00	81.85 \pm 0.00	81.85 \pm 0.00
mammography	87.65 \pm 0.00	87.73 \pm 0.00	87.68 \pm 0.00	87.42 \pm 0.00	87.29 \pm 0.00	87.55 \pm 0.00	87.58 \pm 0.00	87.38 \pm 0.00	86.97 \pm 0.00	86.78 \pm 0.00	86.78 \pm 0.00	86.78 \pm 0.00	86.78 \pm 0.00	86.78 \pm 0.00
mnist	94.22 \pm 0.00	93.93 \pm 0.00	93.57 \pm 0.00	93.20 \pm 0.00	93.08 \pm 0.00	93.60 \pm 0.00	93.85 \pm 0.00	93.45 \pm 0.00	93.00 \pm 0.00	92.55 \pm 0.00	92.46 \pm 0.00	92.46 \pm 0.00	92.46 \pm 0.00	92.46 \pm 0.00
musk	95.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00
optdigits	99.90 \pm 0.00	88.94 \pm 0.00	88.94 \pm 0.00	89.37 \pm 0.00	89.51 \pm 0.00	90.38 \pm 0.00	89.37 \pm 0.00	89.70 \pm 0.00	90.97 \pm 0.00	91.21 \pm 0.00	91.21 \pm 0.00	91.21 \pm 0.00	91.21 \pm 0.00	91.21 \pm 0.00
pendigits	82.21 \pm 1.82	79.74 \pm 1.61	77.98 \pm 1.38	77.44 \pm 1.35	77.14 \pm 1.33	77.87 \pm 1.33	77.87\pm1.43	77.44 \pm 1.07	77.44 \pm 1.07	77.44 \pm 1.07	77.44 \pm 1.07	77.44 \pm 1.07	77.44 \pm 1.07	77.44 \pm 1.07
satellite	82.40 \pm 0.00	82.09 \pm 0.00	81.55 \pm 0.00	80.71 \pm 0.00	80.38 \pm 0.00	81.43 \pm 0.00	82.24 \pm 0.00	81.56 \pm 0.00	80.85 \pm 0.00	80.56 \pm 0.00	80.56 \pm 0.00	80.56 \pm 0.00	80.56 \pm 0.00	80.56 \pm 0.00
simage-2	99.94 \pm 0.00	99.82 \pm 0.00	99.79 \pm 0.00	99.91 \pm 0.00	99.91 \pm 0.00	99.91 \pm 0.00	99.82 \pm 0.00	99.77 \pm 0.00	99.77 \pm 0.00	99.77 \pm 0.00	99.77 \pm 0.00	99.77 \pm 0.00	99.77 \pm 0.00	99.77 \pm 0.00
shuttle	99.66 \pm 0.05	99.52 \pm 0.04	99.47 \pm 0.04	99.54 \pm 0.09	99.54 \pm 0.09	99.51 \pm 0.06	99.51 \pm 0.06	99.15 \pm 0.05	99.15 \pm 0.05	97.43 \pm 0.05	97.43 \pm 0.05	97.43 \pm 0.05	97.43 \pm 0.05	97.43 \pm 0.05
skin	92.94 \pm 2.55	92.84 \pm 2.56	93.14 \pm 2.20	93.14 \pm 2.20	93.14 \pm 2.20	93.04 \pm 2.20	92.91 \pm 2.48	92.91 \pm 2.48	93.11 \pm 2.25	93.11 \pm 2.25	93.11 \pm 2.25	93.11 \pm 2.25	93.11 \pm 2.25	93.11 \pm 2.25
splice	94.46 \pm 0.00	85.94 \pm 0.00	82.97 \pm 0.00	82.97 \pm 0.00	82.97 \pm 0.00	83.07 \pm 0.00	83.36 \pm 0.00	82.68 \pm 0.00	82.68 \pm 0.00	82.68 \pm 0.00	82.68 \pm 0.00	82.68 \pm 0.00	82.68 \pm 0.00	82.68 \pm 0.00
stamps	97.98 \pm 0.33	97.04 \pm 0.33	95.60 \pm 0.78	94.59 \pm 1.07	94.29 \pm 1.08	94.29 \pm 1.08	95.88 \pm 0.66	96.19 \pm 1.24	94.63 \pm 1.35	93.67 \pm 1.09	93.44<			

Table 9.2: Average AUROC \pm standard dev. over five seeds for the semi-supervised setting of ICL and DTE-C baselines with varying hyperparameter (HP) values; For ICL, the learning rate $\in \{0.1, 0.02, 0.001, 0.0001, 1e-05\}$, for DTE-C, $k \in \{5, 10, 20, 40, 50\}$. Also reported is the ^{avg} model. We use **bold** and underline respectively to mark the **best** and the **worst** performance of each model to showcase the variability of performance across different HP settings.

dataset	ICL-0.1	ICL-0.01	ICL-0.001	ICL-0.0001	ICL-Je-45	ICL-Je-05	ICL-C-avr	DTE-C-5	DTE-C-10	DTE-C-20	DTE-C-40	DTE-C-50	DTE-C-avr	
ai0i	47.74 \pm 0.28	47.12 \pm 0.51	46.81 \pm 0.47	48.42 \pm 0.24	48.06 \pm 0.23	47.63 \pm 0.15	47.63 \pm 0.15	50.20 \pm 0.21	50.84 \pm 0.10	50.16 \pm 0.34	50.26 \pm 0.33	50.00 \pm 0.00	50.29 \pm 0.00	
amazon	53.07 \pm 0.19	53.44 \pm 0.19	52.75 \pm 0.68	53.18 \pm 0.15	53.20 \pm 0.21	52.75 \pm 0.19	52.75 \pm 0.19	52.75 \pm 0.16	52.75 \pm 0.16	52.75 \pm 0.16	52.75 \pm 0.16	52.75 \pm 0.16	53.48 \pm 0.33	
amnthyroid	84.02 \pm 0.46	75.72 \pm 0.69	72.68 \pm 3.79	87.29 \pm 1.26	88.84 \pm 2.35	88.52 \pm 1.40	84.27 \pm 2.31	97.47 \pm 0.10	97.65 \pm 0.11	97.75 \pm 0.15	97.40 \pm 0.22	97.40 \pm 0.22	97.40 \pm 0.22	
backdoor	93.03 \pm 0.66	92.91 \pm 0.79	93.30 \pm 0.44	93.93 \pm 0.58	93.32 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	93.30 \pm 0.48	
breast	98.96 \pm 0.47	98.87 \pm 0.20	99.19 \pm 0.34	99.11 \pm 0.18	97.61 \pm 0.65	98.75 \pm 0.18	93.45 \pm 1.04	94.34 \pm 1.25	96.71 \pm 0.91	98.85 \pm 0.45	99.26 \pm 0.07	96.52 \pm 0.53	96.52 \pm 0.53	
campaign	76.07 \pm 0.87	74.61 \pm 1.07	78.88 \pm 0.42	70.79 \pm 0.91	82.36 \pm 0.38	78.33 \pm 0.52	79.18 \pm 1.09	78.24 \pm 2.09	78.49 \pm 1.13	50.00 \pm 0.00	50.00 \pm 0.00	67.18 \pm 0.74	67.18 \pm 0.74	
cardio	48.99 \pm 0.72	84.98 \pm 4.75	82.30 \pm 0.36	78.68 \pm 2.18	88.59 \pm 0.64	77.17 \pm 1.76	88.07 \pm 0.51	87.54 \pm 0.59	87.66 \pm 0.63	50.00 \pm 0.00	50.00 \pm 0.00	72.66 \pm 0.21	72.66 \pm 0.21	
cardiotocography	79.38 \pm 2.17	54.18 \pm 2.77	53.24 \pm 3.25	50.67 \pm 2.55	47.20 \pm 1.33	50.68 \pm 1.24	60.09 \pm 2.19	60.36 \pm 1.54	59.05 \pm 1.34	50.00 \pm 0.00	50.00 \pm 0.00	55.90 \pm 0.23	55.90 \pm 0.23	
census	70.41 \pm 0.77	67.52 \pm 1.79	73.93 \pm 0.78	75.03 \pm 0.46	74.37 \pm 0.53	72.25 \pm 1.63	76.95 \pm 0.91	76.95 \pm 0.91	76.95 \pm 0.91	68.05 \pm 0.22	68.05 \pm 0.22	68.05 \pm 0.22	68.05 \pm 0.22	
cover	86.20 \pm 0.29	83.52 \pm 9.14	91.92 \pm 4.58	94.27 \pm 2.91	98.64 \pm 0.57	99.37 \pm 0.30	95.51 \pm 0.14	96.30 \pm 1.95	97.57 \pm 0.86	97.81 \pm 1.75	97.57 \pm 0.63	97.57 \pm 0.63	97.57 \pm 0.63	
donors	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	6.33 \pm 1.53	
fault	52.16 \pm 0.31	51.88 \pm 0.40	52.28 \pm 0.15	52.35 \pm 0.12	52.35 \pm 0.15	52.35 \pm 0.12	52.35 \pm 0.12	47.68 \pm 1.90	47.68 \pm 1.90	50.00 \pm 0.00	50.00 \pm 0.00	49.27 \pm 1.12	49.27 \pm 1.12	
fraud	9.32 \pm 1.45	63.04 \pm 1.95	61.65 \pm 0.61	61.65 \pm 0.61	95.84 \pm 0.87	93.90 \pm 1.05	94.97 \pm 0.71	94.49 \pm 0.87	94.49 \pm 0.87	93.08 \pm 2.20	93.08 \pm 2.20	93.08 \pm 2.20	93.08 \pm 2.20	
glass	84.02 \pm 8.38	92.66 \pm 3.83	99.95 \pm 0.25	99.95 \pm 0.25	99.95 \pm 0.25	99.95 \pm 0.25	99.95 \pm 0.25	98.27 \pm 1.45	98.27 \pm 1.45	89.00 \pm 2.70	84.89 \pm 2.89	66.42 \pm 8.60	66.42 \pm 8.60	
hepatitis	99.95 \pm 0.11	98.76 \pm 1.84	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	99.93 \pm 0.14	
hup	70.99 \pm 0.07	99.98 \pm 0.01	100.00 \pm 0.00	100.00 \pm 0.00	99.98 \pm 0.02	99.98 \pm 0.02	99.95 \pm 0.08	99.64 \pm 0.28	99.59 \pm 0.06	99.39 \pm 0.07	50.00 \pm 0.00	50.00 \pm 0.00	89.39 \pm 0.09	89.39 \pm 0.09
imdb	72.61 \pm 1.19	73.97 \pm 0.44	74.09 \pm 1.75	73.47 \pm 0.45	68.51 \pm 0.74	72.53 \pm 0.65	72.93 \pm 1.49	77.71 \pm 1.82	77.72 \pm 2.48	50.00 \pm 0.00	50.00 \pm 0.00	66.79 \pm 1.73	66.79 \pm 1.73	
internets	75.53 \pm 1.43	63.04 \pm 1.95	62.32 \pm 1.25	64.94 \pm 0.71	95.84 \pm 0.87	93.90 \pm 1.05	94.97 \pm 0.71	94.67 \pm 1.52	94.67 \pm 1.52	87.57 \pm 1.37	87.57 \pm 1.37	87.57 \pm 1.37	87.57 \pm 1.37	
ionosphere	65.71 \pm 2.22	96.13 \pm 3.45	76.66 \pm 1.79	65.95 \pm 0.76	67.85 \pm 0.68	61.82 \pm 1.06	64.39 \pm 0.59	53.80 \pm 0.94	53.80 \pm 0.94	48.67 \pm 0.62	48.67 \pm 0.62	48.67 \pm 0.62	48.67 \pm 0.62	
landsat	9.61 \pm 2.22	9.61 \pm 2.22	9.61 \pm 2.22	9.61 \pm 2.22	9.61 \pm 2.22	9.61 \pm 2.22	9.61 \pm 2.22	47.20 \pm 1.97	47.20 \pm 1.97	36.75 \pm 1.19	36.75 \pm 1.19	36.75 \pm 1.19	36.75 \pm 1.19	
letter	48.14 \pm 3.53	42.74 \pm 3.41	40.00 \pm 2.13	40.00 \pm 2.13	40.00 \pm 2.13	40.00 \pm 2.13	40.00 \pm 2.13	40.82 \pm 1.97	40.82 \pm 1.97	37.72 \pm 1.31	37.72 \pm 1.31	37.72 \pm 1.31	37.72 \pm 1.31	
lymphography	69.73 \pm 0.00	10.00 \pm 0.00	10.00 \pm 0.00	10.00 \pm 0.00	10.00 \pm 0.00	10.00 \pm 0.00	10.00 \pm 0.00	98.97 \pm 0.31	98.97 \pm 0.31	99.40 \pm 0.22	98.78 \pm 0.66	93.82 \pm 8.49	93.82 \pm 8.49	
magic gamma	79.77 \pm 3.27	76.71 \pm 0.78	78.80 \pm 0.88	77.61 \pm 0.88	78.36 \pm 3.34	77.60 \pm 1.17	78.07 \pm 1.17	80.42 \pm 0.59	87.42 \pm 0.35	87.42 \pm 0.35	87.42 \pm 0.35	87.42 \pm 0.35	87.42 \pm 0.35	
mammography	75.90 \pm 5.19	85.08 \pm 3.05	79.26 \pm 2.14	80.75 \pm 1.40	80.75 \pm 1.40	80.75 \pm 1.40	80.75 \pm 1.40	80.47 \pm 1.62	83.00 \pm 3.62	84.88 \pm 0.62	84.88 \pm 0.62	84.88 \pm 0.62	84.88 \pm 0.62	
mnist	75.53 \pm 1.43	76.05 \pm 3.81	70.13 \pm 1.20	87.05 \pm 1.12	89.01 \pm 0.67	81.37 \pm 0.93	80.21 \pm 0.70	86.88 \pm 1.70	86.88 \pm 1.70	86.88 \pm 1.70	86.88 \pm 1.70	86.88 \pm 1.70	86.88 \pm 1.70	
musik	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	
optdigits	87.01 \pm 1.10	94.97 \pm 1.79	96.45 \pm 2.40	96.17 \pm 2.68	98.70 \pm 0.71	98.88 \pm 0.77	95.10 \pm 0.89	94.60 \pm 2.36	94.60 \pm 2.36	96.89 \pm 0.39	96.89 \pm 0.39	96.89 \pm 0.39	96.89 \pm 0.39	
pima	74.62 \pm 3.67	83.23 \pm 2.89	85.43 \pm 0.52	89.24\pm0.57	85.35 \pm 0.48	85.37 \pm 0.48	85.37 \pm 0.48	80.21 \pm 1.57	80.21 \pm 1.57	67.12 \pm 1.89	67.12 \pm 1.89	67.12 \pm 1.89	67.12 \pm 1.89	
satellite	94.41 \pm 2.46	92.04 \pm 1.26	92.04 \pm 1.26	92.04 \pm 1.26	99.80 \pm 0.06	99.80 \pm 0.06	99.80 \pm 0.06	99.80 \pm 0.06	99.80 \pm 0.06	98.75 \pm 0.22	98.75 \pm 0.22	98.75 \pm 0.22	98.75 \pm 0.22	
satimage-2	99.43 \pm 0.50	99.88 \pm 0.05	99.99 \pm 0.00	99.99 \pm 0.00	99.99 \pm 0.00	99.99 \pm 0.00	99.99 \pm 0.00	99.85 \pm 0.11	99.85 \pm 0.11	99.74 \pm 0.01	99.74 \pm 0.01	99.74 \pm 0.01	99.74 \pm 0.01	
shuttle	53.71 \pm 12.40	86.42 \pm 5.89	86.94 \pm 2.22	70.72 \pm 15.74	73.81 \pm 10.45	91.69 \pm 0.28	92.12 \pm 0.27	87.25 \pm 1.51	87.25 \pm 1.51	86.62 \pm 0.72	86.62 \pm 0.72	86.62 \pm 0.72	86.62 \pm 0.72	
skin	50.12 \pm 0.57	89.53 \pm 4.57	89.53 \pm 4.57	89.53 \pm 4.57	89.53 \pm 4.57	89.53 \pm 4.57	89.53 \pm 4.57	88.37 \pm 1.54	88.37 \pm 1.54	86.65 \pm 1.22	86.65 \pm 1.22	86.65 \pm 1.22	86.65 \pm 1.22	
spambase	79.20 \pm 2.89	79.32 \pm 5.07	90.06 \pm 2.48	84.39 \pm 3.40	83.37 \pm 0.34	83.37 \pm 0.34	83.37 \pm 0.34	82.92 \pm 0.25	82.92 \pm 0.25	82.92 \pm 0.25	82.92 \pm 0.25	82.92 \pm 0.25	82.92 \pm 0.25	
speech	77.62 \pm 0.13	84.33 \pm 3.05	97.29 \pm 0.48	96.70 \pm 0.85	94.64 \pm 3.19	90.01 \pm 2.02	93.01 \pm 1.38	97.12 \pm 1.57	97.12 \pm 1.57	86.48 \pm 2.22	86.48 \pm 2.22	86.48 \pm 2.22	86.48 \pm 2.22	
stamps	97.69 \pm 1.60	97.24 \pm 1.04	96.20 \pm 1.04	96.20 \pm 1.04	96.20 \pm 1.04	96.20 \pm 1.04	96.20 \pm 1.04	95.93 \pm 0.60	95.93 \pm 0.60	95.93 \pm 0.60	95.93 \pm 0.60	95.93 \pm 0.60	95.93 \pm 0.60	
thyroid	53.80 \pm 6.94	58.76 \pm 6.73	75.60 \pm 6.18	82.96 \pm 2.86	81.92 \pm 1.33	70.61 \pm 1.11	87.25 \pm 1.57	69.39 \pm 4.15	69.39 \pm 4.15	70.67 \pm 2.73	69.39 \pm 4.15	69.39 \pm 4.15	69.39 \pm 4.15	
vowels	73.59 \pm 6.94	79.11 \pm 2.66	84.90 \pm 3.49	84.90 \pm 3.49	84.90 \pm 3.49	83.99 \pm 1.93	81.22 \pm 1.57	86.47 \pm 0.72	86.47 \pm 0.72	86.47 \pm 0.72				

Table 10.1: Average AUPR \pm standard dev. over five seeds for the semi-supervised setting of DTE-NP, k NN baselines with varying hyperparameter (HP) values; $k \in \{5, 10, 20, 40, 50\}$. Also reported is the avg model. We use **bold** and underline respectively to mark the **best** and the **worst** performance of each model to showcase the variability of performance across different HP settings.

dataset	DTE-NP-5	DTE-NP-10	DTE-NP-20	DTE-NP-40	DTE-NP-50	DTE-NP-avg	<u>KNN-5</u>	KNN-10	KNN-20	KNN-40	KNN-50	KNN-avg
aloj	5.92 \pm 0.00	5.99 \pm 0.00	6.02 \pm 0.00	6.06 \pm 0.00	6.07 \pm 0.00	6.02 \pm 0.00	6.02 \pm 0.00	6.07 \pm 0.00	6.09 \pm 0.00	6.13 \pm 0.00	6.15 \pm 0.00	6.09 \pm 0.00
amazon	11.68 \pm 0.00	11.68 \pm 0.00	11.68 \pm 0.00	11.61 \pm 0.00	11.62 \pm 0.00	11.65 \pm 0.00	11.69 \pm 0.00	11.70 \pm 0.00	11.65 \pm 0.00	11.69 \pm 0.00	11.65 \pm 0.00	11.65 \pm 0.00
amnthyroid	67.49 \pm 0.00	66.73 \pm 0.00	66.04 \pm 0.00	65.39 \pm 0.00	64.87 \pm 0.00	66.11 \pm 0.00	68.97 \pm 0.00	67.90 \pm 0.00	67.26 \pm 0.00	66.27 \pm 0.00	65.79 \pm 0.00	67.06 \pm 0.00
backdoor	55.90 \pm 0.99	47.16 \pm 1.15	44.55 \pm 1.47	38.31 \pm 0.92	37.36 \pm 0.37	40.48 \pm 0.81	46.70 \pm 0.22	29.58 \pm 1.35	27.36 \pm 1.35	24.34 \pm 0.41	22.34 \pm 0.53	32.06 \pm 0.76
breastw	98.51 \pm 0.56	98.19 \pm 0.58	97.56 \pm 0.51	97.15 \pm 0.62	97.05 \pm 0.45	97.69 \pm 0.40	98.97 \pm 0.28	99.08 \pm 0.31	99.08 \pm 0.23	99.16 \pm 0.16	99.08 \pm 0.22	99.08 \pm 0.22
campaign	28.18 \pm 0.00	49.05 \pm 0.00	49.77 \pm 0.00	49.77 \pm 0.00	49.51 \pm 0.00	49.31 \pm 0.00	59.04 \pm 0.00	49.89 \pm 0.00	49.47 \pm 0.00	49.33 \pm 0.00	49.64 \pm 0.00	49.64 \pm 0.00
cardio	26.90 \pm 0.00	77.73 \pm 0.00	79.19 \pm 0.00	78.35 \pm 0.00	79.13 \pm 0.00	78.35 \pm 0.00	79.22 \pm 0.00	78.35 \pm 0.00	78.35 \pm 0.00	79.30 \pm 0.00	79.30 \pm 0.00	79.30 \pm 0.00
cardiotocography	56.55 \pm 0.00	57.18 \pm 0.00	58.19 \pm 0.00	59.42 \pm 0.00	59.95 \pm 0.00	58.26 \pm 0.00	57.43 \pm 0.00	58.37 \pm 0.00	59.44 \pm 0.00	61.41 \pm 0.00	62.19 \pm 0.00	59.77 \pm 0.00
celuba	10.56 \pm 0.44	11.63 \pm 0.49	13.92 \pm 0.38	14.30 \pm 0.59	12.63 \pm 0.59	12.63 \pm 0.59	11.99 \pm 0.57	12.66 \pm 0.61	14.50 \pm 0.58	15.70 \pm 0.65	14.31 \pm 0.60	14.31 \pm 0.60
census	21.14 \pm 0.39	21.38 \pm 0.54	21.16 \pm 0.43	20.67 \pm 0.41	20.52 \pm 0.42	21.36 \pm 0.76	21.22 \pm 0.39	20.59 \pm 0.33	20.00 \pm 0.42	19.94 \pm 0.44	20.62 \pm 0.44	20.62 \pm 0.44
cover	63.67 \pm 3.21	57.88 \pm 3.52	44.55 \pm 2.49	42.11 \pm 2.29	42.21 \pm 2.29	51.95 \pm 2.15	48.77 \pm 2.84	48.67 \pm 2.84	41.44 \pm 1.04	37.72 \pm 1.51	31.69 \pm 1.35	42.14 \pm 2.20
donors	6.33 \pm 0.80	9.15 \pm 0.72	88.17 \pm 0.95	83.92 \pm 1.29	82.34 \pm 1.32	87.78 \pm 0.99	89.44 \pm 0.96	85.33 \pm 1.15	80.15 \pm 1.33	73.68 \pm 1.32	71.00 \pm 1.30	94.72 \pm 1.13
fault	6.20 \pm 0.00	61.58 \pm 0.00	61.29 \pm 0.00	62.31 \pm 0.00	61.84 \pm 0.00	61.98 \pm 0.00	61.60 \pm 0.00	62.35 \pm 0.61	62.35 \pm 0.61	62.56 \pm 0.00	64.06 \pm 0.00	62.56 \pm 0.00
fraud	40.60 \pm 0.67	43.77 \pm 5.38	43.03 \pm 4.92	39.91 \pm 4.75	38.80 \pm 4.94	41.22 \pm 5.02	42.35 \pm 6.61	44.96 \pm 3.90	41.19 \pm 4.73	37.33 \pm 4.15	36.42 \pm 4.12	40.45 \pm 4.07
glass	60.15 \pm 6.89	47.75 \pm 5.62	37.27 \pm 4.92	31.23 \pm 3.12	30.08 \pm 2.85	41.38 \pm 4.38	44.05 \pm 4.38	32.96 \pm 3.74	29.73 \pm 3.75	26.63 \pm 2.65	26.64 \pm 3.61	31.91 \pm 3.73
hepatitis	9.94 \pm 0.73	9.78 \pm 1.45	9.71 \pm 2.26	81.63 \pm 3.68	81.63 \pm 3.68	91.63 \pm 4.41	69.25 \pm 3.59	69.25 \pm 3.59	64.29 \pm 3.58	64.29 \pm 3.58	70.35 \pm 4.16	70.35 \pm 4.16
hipp	98.52 \pm 0.37	95.38 \pm 0.26	88.66 \pm 1.01	88.43 \pm 2.65	80.40 \pm 4.55	89.48 \pm 0.90	100.00 \pm 0.90	98.04 \pm 3.98	91.24 \pm 1.42	91.44 \pm 1.25	91.28 \pm 1.36	94.39 \pm 1.31
imdb	9.11 \pm 0.00	9.09 \pm 0.00	9.06 \pm 0.00	9.07 \pm 0.00	9.06 \pm 0.00	9.08 \pm 0.00	8.94 \pm 0.00	8.92 \pm 0.00	8.99 \pm 0.00	8.99 \pm 0.00	8.96 \pm 0.00	8.96 \pm 0.00
internets	52.20 \pm 0.00	49.76 \pm 0.00	48.19 \pm 0.00	47.56 \pm 0.00	47.45 \pm 0.00	49.03 \pm 0.00	47.22 \pm 0.00	47.29 \pm 0.00	46.93 \pm 0.00	46.93 \pm 0.00	46.94 \pm 0.00	47.47 \pm 0.00
ionosphere	98.46 \pm 0.54	98.46 \pm 0.54	97.27 \pm 0.42	97.44 \pm 0.50	97.44 \pm 0.50	97.96 \pm 0.46	97.86 \pm 0.60	97.12 \pm 0.52	97.04 \pm 0.60	97.12 \pm 0.45	97.90 \pm 1.65	96.01 \pm 0.78
landsat	56.14 \pm 0.00	54.22 \pm 0.00	50.75 \pm 0.00	46.43 \pm 0.00	45.17 \pm 0.00	50.55 \pm 0.00	54.85 \pm 0.00	50.62 \pm 0.00	45.18 \pm 0.00	41.32 \pm 0.00	40.50 \pm 0.00	44.44 \pm 0.00
letter	8.86 \pm 0.00	8.78 \pm 0.00	8.67 \pm 0.00	8.54 \pm 0.00	8.54 \pm 0.00	8.67 \pm 0.00	8.67 \pm 0.00	8.67 \pm 0.00	8.41 \pm 0.00	8.41 \pm 0.00	8.44 \pm 0.00	8.44 \pm 0.00
lymphography	97.27 \pm 4.45	97.28 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	96.07 \pm 6.79	98.45 \pm 0.65
magic gamma	86.30 \pm 0.00	85.28 \pm 0.00	85.28 \pm 0.00	84.50 \pm 0.00	84.26 \pm 0.00	84.26 \pm 0.00	85.86 \pm 0.00	85.26 \pm 0.00	84.51 \pm 0.00	83.61 \pm 0.00	83.25 \pm 0.00	84.50 \pm 0.00
mammography	42.14 \pm 0.00	41.51 \pm 0.00	40.67 \pm 0.00	40.37 \pm 0.00	40.50 \pm 0.00	41.04 \pm 0.00	41.27 \pm 0.00	40.55 \pm 0.00	40.24 \pm 0.00	38.97 \pm 0.00	39.82 \pm 0.00	39.82 \pm 0.00
mnist	74.43 \pm 0.00	73.09 \pm 0.00	71.84 \pm 0.00	70.69 \pm 0.00	70.36 \pm 0.00	72.08 \pm 0.00	71.40 \pm 0.00	70.09 \pm 0.00	69.02 \pm 0.00	69.02 \pm 0.00	70.56 \pm 0.00	70.56 \pm 0.00
musik	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00
optdigits	34.44 \pm 0.00	30.53 \pm 0.00	26.28 \pm 0.00	24.67 \pm 0.00	24.49 \pm 0.00	21.16 \pm 0.00	21.11 \pm 0.00	24.76 \pm 0.00	24.76 \pm 0.00	24.13 \pm 0.00	23.87 \pm 0.00	24.92 \pm 0.00
pendigits	97.68 \pm 0.00	62.30 \pm 0.00	62.20 \pm 0.00	61.62 \pm 0.00	61.74 \pm 0.00	61.76 \pm 0.00	67.44 \pm 0.00	67.44 \pm 0.00	67.44 \pm 0.00	67.15 \pm 0.00	67.15 \pm 0.00	67.15 \pm 0.00
pima	80.27 \pm 1.65	78.05 \pm 2.13	75.87 \pm 2.40	74.73 \pm 2.71	74.49 \pm 2.71	76.68 \pm 2.22	75.66 \pm 2.91	73.62 \pm 2.59	73.62 \pm 2.59	73.42 \pm 2.79	74.01 \pm 2.75	74.01 \pm 2.75
satellite	85.98 \pm 0.00	85.74 \pm 0.00	85.15 \pm 0.00	84.15 \pm 0.00	83.72 \pm 0.00	84.95 \pm 0.00	86.01 \pm 0.00	85.31 \pm 0.00	84.02 \pm 0.00	82.19 \pm 0.00	81.56 \pm 0.00	83.82 \pm 0.00
satimage-2	96.10 \pm 0.00	97.46 \pm 0.00	97.02 \pm 0.00	97.33 \pm 0.00	97.42 \pm 0.00	96.92 \pm 0.00	96.69 \pm 0.00	97.21 \pm 0.00	97.42 \pm 0.00	97.22 \pm 0.00	97.38 \pm 0.00	97.38 \pm 0.00
shuttle	99.16 \pm 0.00	98.76 \pm 0.00	98.22 \pm 0.00	98.78 \pm 0.00	98.77 \pm 0.00	98.84 \pm 0.00	97.86 \pm 0.00	97.34 \pm 0.00	97.28 \pm 0.00	97.27 \pm 0.00	97.27 \pm 0.00	97.27 \pm 0.00
skin	98.92 \pm 0.23	98.31 \pm 0.30	96.81 \pm 0.30	94.52 \pm 0.43	93.30 \pm 0.53	96.37 \pm 0.25	98.31 \pm 0.34	98.31 \pm 0.30	94.09 \pm 0.51	88.90 \pm 0.08	88.13 \pm 0.46	92.71 \pm 0.16
smp	56.70 \pm 7.16	54.77 \pm 7.80	54.75 \pm 7.81	54.76 \pm 7.81	48.74 \pm 0.23	53.94 \pm 7.83	50.26 \pm 5.73	50.18 \pm 5.73	50.18 \pm 5.73	50.18 \pm 5.73	50.18 \pm 5.73	50.18 \pm 5.73
waveform	83.93 \pm 0.00	83.03 \pm 0.00	82.79 \pm 0.00	82.63 \pm 0.00	83.15 \pm 0.00	82.63 \pm 0.00	82.79 \pm 0.00	82.79 \pm 0.00	82.41 \pm 0.00	82.11 \pm 0.00	82.55 \pm 0.00	82.55 \pm 0.00
whb	92.08 \pm 6.52	89.03 \pm 6.77	86.03 \pm 5.81	83.79 \pm 5.59	83.41 \pm 5.19	86.17 \pm 5.59	85.35 \pm 4.46	83.71 \pm 5.59	82.05 \pm 4.56	82.37 \pm 4.57	82.33 \pm 4.54	83.17 \pm 5.03
wilt	13.43 \pm 0.00	12.36 \pm 0.00	11.30 \pm 0.00	10.40 \pm 0.00	10.12 \pm 0.00	11.52 \pm 0.00	12.25 \pm 0.00	11.04 \pm 0.00	10.12 \pm 0.00	9.99 \pm 0.00	10.36 \pm 0.00	10.36 \pm 0.00
wine	99.92 \pm 0.77	98.33 \pm 0.52	96.09 \pm 1.31	93.12 \pm 1.51	91.92 \pm 2.12	95.71 \pm 0.92	95.18 \pm 1.66	88.54 \pm 3.30	88.36 \pm 3.30	85.45 \pm 3.97	85.79 \pm 1.56	88.72 \pm 2.32
yeast	48.37 \pm 0.00	47.19 \pm 0.00	47.54 \pm 0.00	47.84 \pm 0.00	47.28 \pm 0.00	47.83 \pm 0.00	47.65 \pm 0.00	47.82 \pm 0.00	47.42 \pm 0.00	47.24 \pm 0.00	46.48 \pm 0.00	47.24 \pm 0.

Table 10.2: Average AUPR \pm standard dev. over five seeds for the semi-supervised setting of ICL and DTE-C baselines with varying hyperparameter (HP) values; For ICL, the learning rate $\in \{0.1, 0.02, 0.001, 0.0001, 1e-05\}$, for DTE-C, $k \in \{5, 10, 20, 40, 50\}$. Also reported is the avg. model. We use **bold** and underline respectively to mark the **best** and the **worst** performance of each model to showcase the variability of performance across different HP settings.

dataset	ICL-0.1	ICL-0.01	ICL-0.001	ICL-0.0001	ICL-je-05	ICL-car	<u>DTE-C-5</u>	DTE-C-10	DTE-C-20	DTE-C-40	DTE-C-50	DTE-C-avg	
ai0i	5.50±0.09	5.39±0.07	5.50±0.08	5.59±0.04	5.46±0.01	5.49±0.02	5.76±0.03	5.82±0.02	5.72±0.05	5.73±0.05	5.91±0.00	5.79±0.00	
amazon	10.06±0.10	10.08±0.11	9.91±0.03	9.91±0.02	10.01±0.43	10.01±0.55	10.01±0.50	10.01±0.43	10.01±0.52	10.01±0.52	10.02±0.40	10.02±0.40	
amnthyroid	58.53±13.33	39.69±5.23	53.56±3.08	53.83±5.44	55.94±2.70	52.34±3.72	82.46±0.50	83.25±0.34	83.27±0.63	81.52±1.18	83.81±0.00	68.86±0.40	
backdoor	88.81±2.45	88.14±1.38	88.99±1.12	88.96±1.20	87.63±1.03	87.63±1.03	61.74±2.58	63.64±1.61	4.83±0.09	35.68±0.79	35.68±0.79	35.68±0.79	
breast	98.65±0.81	98.46±0.51	98.98±0.57	98.50±0.42	95.32±1.11	97.98±0.31	88.69±2.54	89.49±2.00	94.04±1.49	98.26±0.68	99.22±0.11	94.04±0.69	
campaign	47.46±0.41	42.03±2.05	47.94±0.25	42.03±2.05	49.49±0.21	48.01±0.47	49.96±0.20	46.77±1.41	48.30±1.60	20.25±0.00	37.11±0.64	37.11±0.64	
cardio	48.81±14.30	65.70±1.26	63.55±5.07	61.75±2.55	59.90±5.07	70.19±0.65	69.32±0.45	70.19±0.65	70.25±0.00	49.02±0.29	49.02±0.29	49.02±0.29	
cardiotocography	43.21±5.03	55.56±1.14	50.69±2.12	49.01±1.29	39.79±1.29	47.05±1.34	53.83±0.94	53.83±0.94	53.36±0.73	49.12±4.17	36.12±0.00	45.75±0.91	
celiba	12.89±1.17	13.90±1.49	13.09±1.15	13.30±1.73	23.32±0.70	23.68±0.70	22.37±0.59	22.37±0.59	13.97±1.22	13.97±1.59	10.05±0.49	10.05±0.49	
census	20.29±1.17	20.38±2.24	20.32±1.28	20.32±1.28	47.19±1.76	47.19±1.76	31.42±3.97	31.42±3.97	17.28±0.52	17.45±0.64	11.66±0.20	15.29±0.31	
cover	29.58±13.00	73.32±4.93	82.59±7.03	91.60±4.15	92.24±2.17	76.75±3.97	76.07±1.84	66.66±3.73	33.74±2.07	18.86±15.32	35.03±0.82	35.03±0.82	
donors	39.48±1.33	73.88±8.28	63.83±0.94	63.83±0.94	64.32±0.96	64.46±1.09	63.92±1.40	63.62±1.06	51.74±0.32	58.79±0.41	58.79±0.41	58.79±0.41	
fault	65.37±3.40	64.97±9.27	56.24±0.47	62.41±10.30	72.24±4.63	58.46±7.58	68.85±0.35	62.09±1.19	57.77±2.19	29.78±2.62	30.33±0.25	30.33±0.25	
glass	49.70±15.55	66.98±13.15	88.19±0.70	87.97±1.44	87.79±0.75	75.58±1.38	86.09±0.67	75.58±1.38	75.58±1.38	27.63±1.55	7.83±1.03	7.83±1.03	
hepatitis	99.85±0.30	97.89±2.51	99.70±0.91	98.94±1.21	99.92±0.13	99.92±0.13	99.92±0.13	99.92±0.13	99.92±0.13	54.16±1.12	74.38±1.79	74.38±1.79	
hup	94.85±0.03	95.65±3.25	97.86±1.72	99.77±0.25	99.55±0.61	99.55±0.61	99.55±0.61	99.55±0.61	99.55±0.61	50.07±0.78	48.35±0.22	48.35±0.22	
imdb	10.09±0.09	10.02±0.07	10.16±0.04	10.41±0.04	10.18±0.05	10.41±0.04	10.71±0.03	10.71±0.03	9.49±0.47	8.79±0.30	9.52±0.00	9.21±0.20	
internets	50.73±2.19	60.94±1.44	62.86±1.88	63.83±0.88	47.64±1.20	58.52±1.04	60.45±1.51	60.45±1.51	56.24±2.74	57.76±6.93	31.53±0.00	47.15±1.92	
ionosphere	51.45±12.34	64.97±9.27	56.24±0.47	62.41±10.30	72.24±4.63	58.46±7.58	68.85±0.35	62.09±1.19	57.77±2.19	29.78±2.62	30.33±0.25	30.33±0.25	
landsat	66.94±2.27	97.20±2.57	95.50±0.77	95.52±1.20	95.52±0.77	95.52±0.77	95.52±0.77	95.52±0.77	95.52±0.77	95.52±0.77	95.52±0.77	95.52±0.77	
letter	11.22±1.58	10.84±0.96	9.41±0.34	15.87±1.01	10.01±0.55	9.71±0.43	10.01±0.43	9.71±0.43	9.71±0.43	34.01±2.31	34.27±0.09	34.32±0.00	
lymphography	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	10.09±0.04	10.09±0.04	
magic gamma	81.20±2.83	82.61±2.83	82.64±0.71	82.47±0.98	83.45±0.84	80.74±1.03	87.09±0.70	87.09±0.70	93.30±2.48	83.45±0.20	72.50±16.37	68.82±5.96	
mammography	33.51±4.66	37.72±4.74	27.06±1.69	17.30±1.13	17.30±1.13	17.30±1.13	17.30±1.13	17.30±1.13	17.30±1.13	77.29±3.52	77.29±3.52	77.29±3.52	
mnist	49.22±1.39	50.49±3.02	54.46±1.34	63.49±1.45	65.20±1.41	56.56±0.65	60.64±1.13	56.56±0.65	56.56±0.65	53.17±3.05	41.15±1.11	41.15±1.11	
musk	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	6.14±0.00	62.46±0.00	62.46±0.00	
optdigits	64.80±1.36	66.26±1.49	63.86±1.28	63.86±1.24	64.31±1.19	62.39±0.52	62.09±0.68	62.09±0.68	62.09±0.68	62.82±1.63	17.71±0.02	17.71±0.02	
pendigits	39.64±0.69	66.26±1.23	63.38±8.59	73.08±7.23	50.36±3.65	68.37±1.14	64.41±1.13	64.41±1.13	64.41±1.13	64.99±4.40	54.53±1.01	54.53±1.01	
pima	73.92±1.66	81.23±2.83	82.64±0.71	82.47±0.98	83.45±0.84	80.74±1.03	87.09±0.70	87.09±0.70	93.30±2.48	83.45±0.20	72.50±16.37	68.82±5.96	
satellite	31.55±4.47	37.72±4.74	27.06±1.69	17.30±1.13	17.30±1.13	17.30±1.13	17.30±1.13	17.30±1.13	17.30±1.13	77.29±3.52	77.29±3.52	77.29±3.52	
satimage-2	97.77±1.43	99.19±1.23	99.82±0.12	99.91±0.05	99.72±0.12	99.28±0.28	99.45±0.28	99.45±0.28	99.45±0.28	93.19±0.16	46.25±2.72	46.25±2.72	
shuttle	44.77±0.20	47.14±0.35	47.39±0.71	47.39±0.71	47.39±0.71	47.39±0.71	47.39±0.71	47.39±0.71	47.39±0.71	88.77±1.78	33.35±0.00	27.84±1.88	
skin	38.06±0.67	45.88±7.84	37.40±6.89	37.40±6.89	37.40±6.89	37.40±6.89	37.40±6.89	37.40±6.89	37.40±6.89	69.80±0.53	69.80±0.53	69.80±0.53	
spambase	80.80±0.90	81.11±8.20	81.11±8.20	81.11±8.20	84.52±0.70	85.16±0.70	85.42±0.70	85.30±0.50	85.30±0.50	85.39±0.51	31.76±0.00	31.76±0.00	
speech	37.44±1.02	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	37.99±0.38	
stamps	45.19±8.30	50.55±12.46	50.56±5.89	50.60±5.56	50.60±5.56	50.60±5.56	50.60±5.56	50.60±5.56	50.60±5.56	53.33±6.09	58.82±10.61	58.82±10.61	
thyroid	72.79±3.97	89.69±1.74	88.43±0.25	90.38±0.44	86.65±0.69	85.84±2.16	88.25±0.62	87.01±0.52	87.42±0.41	55.55±14.95	71.75±8.26	71.75±8.26	
vowel	25.76±4.65	31.37±4.07	50.83±13.40	53.53±3.61	54.90±6.08	44.44±4.60	59.35±5.57	59.35±5.57	59.35±5.57	59.35±5.57	46.25±2.06	46.25±2.06	
waveform	24.00±15.67	21.57±7.57	28.45±5.52	28.58±5.64	28.58±5.64	28.58±5.64	28.58±5.64	28.58±5.64	28.58±5.64	28.95±5.65	47.17±4.01	47.17±4.01	
whale	82.98±11.29	91.66±5.01	99.04±1.73	99.16±1.35	96.58±2.18	97.22±2.30	28.91±0.57	28.91±0.57	28.91±0.57	40.09±7.55	67.56±7.06	67.56±7.06	
wilt	11.08±5.73	12.91±2.67	98.33±1.88	99.48±0.15	97.94±2.12	98.11±3.43	98.34±2.75	98.34±2.75	98.34±2.75	21.04±1.52	25.25±1.80	25.25±1.80	
wine	82.28±2.77	66.72±1.51	49.01±0.39	49.32±0.70	49.32±0.70	80.78±4.52	60.51±5.10	61.31±5.87	59.53±3.80	40.41±3.04	38.36±1.82	38.36±1.82	
yeast	9.79±0.11	9.91±0.05	9.82±0.17	9.84±0.06	9.59±0.05	48.61±0.81	49.63±0.87	50.38±1.14	47.04±1.19	49.58±0.00	62.54±0.43	62.54±0.43	
yelp	44.47±0.54	48.19±0.57	50.18±0.26	50.34±0.31	50.18±0.26	47.83±0.12	46.68±0.62	47.39±0.52	46.12±0.32	9.52±0.00	11.86±0.62	11.86±0.62	
MNIST-C	FashionMNIST	58.47±0.45	64.11±0.18	64.96±0.32	65.22±0.32	55.84±0.54	61.72±0.18	54.92±0.36	56.00±0.23	56.31±0.35	9.50±0.00	37.25±0.17	37.25±0.17
CIFAR10	CIFAR10	14.30±0.27	17.49±0.41	19.10±0.21	18.89±0.21	16.71±0.19	19.95±0.09	19.61±0.15	19.52±0.00	15.08±0.07	15.08±0.07	15.08±0.07	15.08±0.07
SVHN	SVHN	13.92±0.24	15.32±0.05	15.82±0.01	15.74±0.09	15.16±0.13	15.52±0.19	15.52±0.19	15.52±0.19	15.16±0.04	13.16±0.04	13.16±0.04	13.16±0.04
MTeC-AD	MTeC-AD	86.94±0.92	88.30±0.67	88.91±0.75	88.91±0.75	83.64±0.91	87.35±0.76	83.04±1.01	84.66±1.01	84.41±1.22	38.89±1.52	65.84±0.58	65.84±0.58
2news	2news	12.37±0.20	13.20±0.34	13.52±0.34	13.98±0.35	13.88±0.29	12.98±0.23	13.75±0.40	15.75±0.40	13.29±0.47	13.29±0.47	13.29±0.47	13.29±0.47
agnews	agnews	12.45±0.12	12.84±0.25	13.04±0.08	13.05±0.04	12.55±0.06	12.79±0.06	16.51±0.99	16.18±1.53	9.52±0.00	14.03±0.43	14.03±0.43	14.03±0.43

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

Table 11.1: Average F1 score \pm standard dev. over five seeds for the semi-supervised setting of DTE-NP, k NN baselines with varying hyperparameter (HP) values; $k \in \{5, 10, 20, 40, 50\}$. Also reported is the avg model. We use **bold** and underline respectively to mark the **best** and the **worst** performance of each model to showcase the variability of performance across different HP settings.

dataset	DTE-NP-5	DTE-NP-10	DTE-NP-20	DTE-NP-40	DTE-NP-50	DTE-NP-avr	<u>DTE-NP-avr</u>	KNN-5	KNN-10	KNN-20	KNN-40	KNN-50	KNN-avr
aloj	5.90 \pm 0.00	5.84 \pm 0.00	5.70 \pm 0.00	5.90 \pm 0.00	5.97 \pm 0.00	5.86 \pm 0.00	5.90\pm0.00	5.90 \pm 0.00	5.64 \pm 0.00	5.97 \pm 0.00	6.17 \pm 0.00	6.37\pm0.00	6.01 \pm 0.00
amazon	10.80 \pm 0.00	10.80 \pm 0.00	10.20 \pm 0.00	10.20 \pm 0.00	10.40 \pm 0.00	10.60 \pm 0.00	10.40\pm0.00	10.20 \pm 0.00	11.20 \pm 0.00	11.20 \pm 0.00	11.20 \pm 0.00	11.20\pm0.00	10.92 \pm 0.00
amnthyroid	62.55\pm0.00	61.80 \pm 0.00	60.67 \pm 0.00	59.89 \pm 0.00	58.80 \pm 0.00	60.56 \pm 0.00	61.99\pm0.00	60.49 \pm 0.00	58.43 \pm 0.00	58.24 \pm 0.00	56.74 \pm 0.00	59.18\pm0.00	59.18 \pm 0.00
backdoor	64.15\pm1.04	52.30 \pm 1.37	40.62 \pm 1.46	26.35 \pm 1.20	42.86 \pm 1.32	40.37 \pm 1.04	52.53\pm1.63	42.71 \pm 1.50	20.21 \pm 0.78	17.52 \pm 0.83	31.87 \pm 1.22	31.87\pm1.22	31.87 \pm 1.22
breastw	96.72\pm0.64	96.23 \pm 0.50	96.17 \pm 0.47	95.99 \pm 0.39	95.99 \pm 0.39	96.00 \pm 0.43	96.00\pm0.44	95.87 \pm 0.33	95.99 \pm 0.32	95.97 \pm 0.32	95.97 \pm 0.32	95.97\pm0.32	95.97 \pm 0.32
campaign	49.94 \pm 0.00	50.62 \pm 0.00	51.14 \pm 0.00	51.38 \pm 0.00	51.57 \pm 0.00	50.93 \pm 0.00	50.37\pm0.00	50.88 \pm 0.00	51.27 \pm 0.00	51.29 \pm 0.00	51.01 \pm 0.00	51.01\pm0.00	51.01 \pm 0.00
cardio	63.64 \pm 0.00	61.56 \pm 0.00	61.93 \pm 0.00	62.40 \pm 0.00	63.64 \pm 0.00	62.95 \pm 0.00	61.93\pm0.00	61.93 \pm 0.00	64.20 \pm 0.00	67.61 \pm 0.00	65.00 \pm 0.00	65.00\pm0.00	65.00 \pm 0.00
cardiotocography	44.64 \pm 0.00	45.71 \pm 0.00	47.00 \pm 0.00	48.50 \pm 0.00	49.14 \pm 0.00	47.00 \pm 0.00	46.35\pm0.00	46.35 \pm 0.00	46.78 \pm 0.00	47.85 \pm 0.00	46.87 \pm 0.00	46.87\pm0.00	46.87 \pm 0.00
celsba	15.83 \pm 0.69	17.05 \pm 0.69	19.30 \pm 0.69	17.87 \pm 0.69	19.30 \pm 0.69	17.48 \pm 0.57	17.08\pm0.58	17.84 \pm 0.68	19.30 \pm 0.81	19.11 \pm 0.61	19.11 \pm 0.61	19.11\pm0.61	19.11 \pm 0.61
census	22.22 \pm 0.54	21.93 \pm 0.52	21.46 \pm 0.25	21.38 \pm 0.48	21.12 \pm 0.29	21.62 \pm 0.14	22.23\pm0.42	21.48 \pm 0.40	21.47 \pm 0.57	21.33 \pm 0.65	21.26 \pm 0.24	21.26\pm0.24	21.26 \pm 0.24
cover	69.15\pm1.12	66.87 \pm 2.35	63.15 \pm 2.35	59.59 \pm 2.08	53.06 \pm 2.23	61.65 \pm 2.14	65.04\pm1.92	60.56 \pm 2.04	42.69 \pm 1.94	39.92 \pm 1.99	52.19 \pm 1.92	52.19\pm1.92	52.19 \pm 1.92
donors	97.27 \pm 0.36	96.20 \pm 0.45	94.49 \pm 0.55	91.70 \pm 0.00	90.57 \pm 0.86	94.05 \pm 0.62	94.98\pm0.62	92.36 \pm 0.59	88.71 \pm 0.99	80.36 \pm 1.90	76.62 \pm 1.50	86.83\pm0.98	86.83 \pm 0.98
fault	56.02 \pm 0.00	55.72 \pm 0.00	54.70 \pm 0.00	54.76 \pm 0.00	56.91 \pm 0.00	56.32 \pm 0.00	56.78\pm0.00	57.06 \pm 0.00	58.25 \pm 0.00	58.25 \pm 0.00	58.25 \pm 0.00	58.25\pm0.00	58.25 \pm 0.00
fraud	48.18 \pm 4.56	49.60 \pm 3.28	49.00 \pm 3.95	46.66 \pm 3.52	45.46 \pm 3.60	47.78 \pm 3.53	49.39\pm5.24	46.64 \pm 4.18	42.58 \pm 3.16	41.64 \pm 3.36	45.61 \pm 3.48	45.61\pm3.48	45.61 \pm 3.48
glass	47.81\pm5.78	35.14 \pm 2.75	27.98 \pm 4.21	18.37 \pm 2.40	12.81 \pm 2.98	29.42 \pm 2.61	29.87\pm4.62	22.55 \pm 5.87	18.58 \pm 1.19	17.23 \pm 3.73	17.23 \pm 3.73	17.23\pm3.73	17.23 \pm 3.73
hepatitis	98.94\pm1.41	94.65 \pm 2.13	91.88 \pm 4.18	75.90 \pm 4.21	71.99 \pm 4.14	84.54 \pm 5.50	84.93\pm4.53	76.59 \pm 6.21	66.15 \pm 7.50	67.16 \pm 7.50	66.15 \pm 7.50	66.15\pm7.50	66.15 \pm 7.50
hip	98.50\pm0.38	95.02 \pm 2.50	86.26 \pm 1.38	82.83 \pm 3.80	83.66 \pm 6.40	88.73 \pm 2.54	100.00\pm0.00	92.67 \pm 0.91	92.67 \pm 0.91	92.67 \pm 0.91	92.67 \pm 0.91	92.67\pm0.91	92.67 \pm 0.91
imdb	5.20\pm0.00	5.40 \pm 0.00	5.20 \pm 0.00	5.20 \pm 0.00	5.20 \pm 0.00	5.40 \pm 0.00	5.40\pm0.00	5.40 \pm 0.00	5.40 \pm 0.00	5.40 \pm 0.00	5.40 \pm 0.00	5.40\pm0.00	5.40 \pm 0.00
interneads	55.16\pm0.00	51.63 \pm 0.00	48.37 \pm 0.00	46.47 \pm 0.00	46.20 \pm 0.00	49.57 \pm 0.00	51.90\pm0.00	46.20 \pm 0.00	45.11 \pm 0.00	45.11 \pm 0.00	45.11 \pm 0.00	45.11\pm0.00	45.11 \pm 0.00
ionosphere	92.33\pm1.17	92.05 \pm 1.63	91.63 \pm 1.09	90.41 \pm 1.35	89.19 \pm 1.72	89.12 \pm 1.24	91.81\pm1.87	89.23 \pm 1.86	89.23 \pm 1.86	87.86 \pm 1.86	87.86 \pm 1.86	87.86\pm1.86	87.86 \pm 1.86
landsat	52.29\pm0.00	51.62 \pm 0.00	49.06 \pm 0.00	49.99 \pm 0.00	45.39 \pm 0.00	48.79 \pm 0.00	51.46\pm0.00	49.29 \pm 0.00	45.76 \pm 0.00	45.76 \pm 0.00	46.69 \pm 0.00	46.69\pm0.00	46.69 \pm 0.00
letter	1.00\pm0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00\pm0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00\pm0.00	1.00 \pm 0.00
lymphography	77.89\pm4.21	93.16 \pm 7.97	94.67 \pm 6.53	92.71 \pm 6.09	91.66 \pm 7.34	91.57 \pm 5.93	91.57\pm5.93	90.69 \pm 3.65	90.69 \pm 3.65	90.69 \pm 3.65	90.69 \pm 3.65	90.69\pm3.65	90.69 \pm 3.65
magic gamma	76.79\pm0.00	76.20 \pm 0.00	75.49 \pm 0.00	74.84 \pm 0.00	74.75 \pm 0.00	75.61 \pm 0.00	76.17\pm0.00	75.13 \pm 0.00	74.60 \pm 0.00	74.73 \pm 0.00	74.73 \pm 0.00	74.73\pm0.00	74.73 \pm 0.00
mammography	41.92 \pm 0.00	41.15 \pm 0.00	44.23\pm0.00	44.23 \pm 0.00	44.23 \pm 0.00	43.15 \pm 0.00	43.46\pm0.00	43.46 \pm 0.00	43.85 \pm 0.00	43.85 \pm 0.00	43.85 \pm 0.00	43.85\pm0.00	43.85 \pm 0.00
mnist	72.71\pm0.00	72.29 \pm 0.00	70.45 \pm 0.00	70.45 \pm 0.00	69.88 \pm 0.00	71.35 \pm 0.00	71.86\pm0.00	80.06 \pm 0.00	98.06 \pm 0.00	98.06 \pm 0.00	98.06 \pm 0.00	98.06\pm0.00	98.06 \pm 0.00
musik	100.00\pm0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00\pm0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00\pm0.00	100.00 \pm 0.00
optdigits	89.00 \pm 0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20\pm0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20 \pm 0.00	89.20\pm0.00	89.20 \pm 0.00
pendigits	94.23\pm0.00	92.31 \pm 0.00	91.60 \pm 0.00	91.60 \pm 0.00	91.60 \pm 0.00	91.60 \pm 0.00	91.60\pm0.00	91.60 \pm 0.00	91.60 \pm 0.00	91.60 \pm 0.00	91.60 \pm 0.00	91.60\pm0.00	91.60 \pm 0.00
pima	74.73\pm2.13	72.94 \pm 2.46	71.48 \pm 2.36	71.93 \pm 1.99	70.58 \pm 1.97	71.67 \pm 2.12	71.81\pm0.00	70.73 \pm 0.00	69.84 \pm 0.00	69.84 \pm 0.00	69.84 \pm 0.00	69.84\pm0.00	69.84 \pm 0.00
satimage-2	90.14 \pm 0.00	76.76 \pm 0.00	90.14 \pm 0.00	90.14 \pm 0.00	90.14 \pm 0.00	90.14 \pm 0.00	92.96\pm0.00	92.96 \pm 0.00	92.96 \pm 0.00	92.96 \pm 0.00	92.96 \pm 0.00	92.96\pm0.00	92.96 \pm 0.00
shuttle	98.35\pm0.00	98.23 \pm 0.00	98.15 \pm 0.00	98.12 \pm 0.00	98.12 \pm 0.00	98.19 \pm 0.00	98.19\pm0.00	98.06 \pm 0.00	98.06 \pm 0.00	98.06 \pm 0.00	98.06 \pm 0.00	98.06\pm0.00	98.06 \pm 0.00
skin	97.12\pm0.46	96.83 \pm 0.38	96.14 \pm 0.17	94.73 \pm 0.33	95.83 \pm 0.27	95.82 \pm 0.14	96.71\pm0.41	95.82 \pm 0.17	95.82 \pm 0.17	95.82 \pm 0.17	95.82 \pm 0.17	95.82\pm0.17	95.82 \pm 0.17
smp	68.05 \pm 5.12	68.05 \pm 5.12	69.59\pm3.95	69.59 \pm 3.95	63.31 \pm 5.12	67.75 \pm 4.49	69.59\pm3.95	69.59 \pm 3.95	69.59 \pm 3.95	69.59 \pm 3.95	69.59 \pm 3.95	69.59\pm3.95	69.59 \pm 3.95
waveform	28.00 \pm 0.00	28.00 \pm 0.00	28.00 \pm 0.00	28.00 \pm 0.00	28.00 \pm 0.00	28.00 \pm 0.00	28.00\pm0.00	27.00 \pm 0.00	27.00 \pm 0.00	27.00 \pm 0.00	27.00 \pm 0.00	27.00\pm0.00	27.00 \pm 0.00
whb	85.35\pm0.08												

1728
 1729
 1730
 1731
 1732
 1733
 1734
 1735
 1736
 1737
 1738
 1739
 1740
 1741
 1742
 1743
 1744
 1745
 1746
 1747
 1748
 1749
 1750
 1751
 1752
 1753
 1754
 1755
 1756
 1757
 1758
 1759
 1760
 1761
 1762
 1763
 1764
 1765
 1766
 1767
 1768
 1769
 1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781

Table 11.2: Average F1 score \pm standard dev. over five seeds for the semi-supervised setting of ICL and DTE-C baselines with varying hyperparameter (HP) values; For ICL, the learning rate $\in \{0.1, 0.02, 0.001, 0.0001, 1e-05\}$, for DTE-C, $k \in \{5, 10, 20, 40, 50\}$. Also reported is the ^{avg} model. We use **bold** and underline respectively to mark the **best** and the worst performance of each model to showcase the variability of performance across different HP settings.

dataset	ICL-0.1	ICL-0.01	ICL-0.001	ICL-0.0001	ICL-1e-05	ICL-avr		DTE-C-5		DTE-C-10		DTE-C-20		DTE-C-40		DTE-C-50		DTE-C-avr
aloj	4.51 \pm 0.69	4.34 \pm 0.42	5.28 \pm 0.47	4.68 \pm 0.30	4.16 \pm 0.38	4.49 \pm 0.07		4.75 \pm 0.27		4.27 \pm 0.19		4.28 \pm 0.10		4.51 \pm 0.17		4.00 \pm 0.00		3.36 \pm 0.03
amazon	9.76 \pm 0.34	9.92 \pm 0.84	9.76 \pm 0.34	9.34 \pm 0.45	5.73 \pm 3.35	52.56 \pm 2.97		11.96 \pm 0.32		11.48 \pm 0.97		11.60 \pm 1.68		11.06 \pm 0.28		7.01 \pm 0.72		7.01 \pm 0.72
amnthyroid	54.87 \pm 3.24	47.25 \pm 3.35	53.45 \pm 4.13	54.72 \pm 5.45	57.33 \pm 2.97	52.56 \pm 3.28		77.23 \pm 0.25		77.94 \pm 0.85		77.33 \pm 0.85		75.43 \pm 0.93		61.63 \pm 0.20		61.63 \pm 0.20
backdoor	87.17 \pm 0.98	87.32 \pm 0.99	87.11 \pm 1.09	86.85 \pm 0.95	85.37 \pm 1.01	86.76 \pm 1.00		86.48 \pm 0.28		86.19 \pm 0.39		83.03 \pm 2.14		84.50 \pm 0.60		0.00 \pm 0.00		0.00 \pm 0.00
breastw	95.28 \pm 0.34	96.07 \pm 0.04	96.80 \pm 0.40	97.44 \pm 0.05	96.90 \pm 0.75	96.48 \pm 0.59		96.50 \pm 0.59		96.10 \pm 0.35		92.46 \pm 1.78		95.31 \pm 0.70		96.11 \pm 0.44		92.50 \pm 0.79
campaign	48.12 \pm 0.36	46.88 \pm 1.72	46.93 \pm 1.51	46.88 \pm 1.00	40.52 \pm 4.28	53.40 \pm 0.51		50.07 \pm 0.49		51.98 \pm 0.70		52.45 \pm 1.07		52.33 \pm 1.20		0.00 \pm 0.00		0.00 \pm 0.00
cardio	49.09 \pm 1.28	61.93 \pm 5.57	58.86 \pm 1.59	57.95 \pm 4.57	40.52 \pm 4.28	53.68 \pm 2.83		57.95 \pm 4.57		58.86 \pm 4.43		58.07 \pm 0.58		57.84 \pm 0.43		34.91 \pm 0.09		34.91 \pm 0.09
cardiotocography	36.14 \pm 1.28	41.07 \pm 1.73	39.18 \pm 1.49	35.36 \pm 2.14	32.66 \pm 1.59	36.88 \pm 1.59		39.91 \pm 1.05		39.48 \pm 1.54		37.73 \pm 1.73		42.02 \pm 0.00		48.23 \pm 0.39		48.23 \pm 0.39
celsba	15.42 \pm 2.39	17.97 \pm 2.55	17.20 \pm 1.92	17.46 \pm 1.17	17.06 \pm 1.65	17.46 \pm 1.12		17.12 \pm 1.45		17.18 \pm 2.74		17.58 \pm 1.42		17.54 \pm 1.79		16.44 \pm 1.41		10.12 \pm 1.04
census	22.27 \pm 1.73	24.06 \pm 2.05	25.80 \pm 1.34	24.68 \pm 2.74	27.15 \pm 1.12	24.34 \pm 2.90		26.61 \pm 2.59		26.61 \pm 2.59		26.61 \pm 2.59		26.61 \pm 2.59		10.31 \pm 0.42		10.31 \pm 0.42
cover	26.77 \pm 1.54	15.52 \pm 3.18	42.68 \pm 1.70	33.70 \pm 16.88	44.34 \pm 2.29	36.61 \pm 2.59		36.61 \pm 2.59		36.61 \pm 2.59		36.61 \pm 2.59		36.61 \pm 2.59		38.39 \pm 0.62		38.39 \pm 0.62
donors	43.71 \pm 1.52	81.85 \pm 3.56	83.59 \pm 4.47	89.28 \pm 2.06	92.77 \pm 1.39	78.24 \pm 2.61		87.99 \pm 1.87		87.99 \pm 1.87		83.08 \pm 1.20		11.80 \pm 23.60		0.00 \pm 0.00		0.00 \pm 0.00
fault	60.33 \pm 3.36	56.52 \pm 2.09	58.57 \pm 1.11	58.57 \pm 1.36	59.13 \pm 1.51	52.30 \pm 1.36		58.37 \pm 1.51		58.37 \pm 1.51		54.25 \pm 5.80		22.55 \pm 24.73		0.00 \pm 0.00		0.00 \pm 0.00
fraud	57.54 \pm 1.13	48.97 \pm 8.02	58.90 \pm 6.77	58.90 \pm 6.77	79.18 \pm 3.21	62.30 \pm 4.91		75.61 \pm 7.76		75.61 \pm 7.76		75.23 \pm 5.70		0.00 \pm 0.00		30.48 \pm 4.66		30.48 \pm 4.66
glass	43.53 \pm 0.21	57.05 \pm 16.03	80.05 \pm 6.11	87.24 \pm 5.04	82.50 \pm 5.41	70.87 \pm 3.12		35.43 \pm 5.07		34.82 \pm 4.93		31.26 \pm 6.69		19.45 \pm 5.66		24.15 \pm 2.73		24.15 \pm 2.73
hepatitis	99.94 \pm 0.71	99.60 \pm 7.81	99.64 \pm 0.71	99.64 \pm 0.71	98.65 \pm 2.11	96.40 \pm 3.09		96.49 \pm 3.09		96.65 \pm 2.11		96.65 \pm 2.11		78.91 \pm 7.98		78.91 \pm 7.98		78.91 \pm 7.98
hip	9.39 \pm 0.35	9.67 \pm 3.07	9.69 \pm 1.51	9.84 \pm 0.54	10.40 \pm 0.34	9.77 \pm 1.10		9.52 \pm 0.34		9.52 \pm 0.34		9.52 \pm 0.34		16.73 \pm 12.06		16.73 \pm 12.06		16.73 \pm 12.06
imdb	10.52 \pm 0.65	10.44 \pm 0.54	10.44 \pm 0.54	10.44 \pm 0.54	10.40 \pm 0.53	10.40 \pm 0.53		10.19 \pm 0.53		10.19 \pm 0.53		10.19 \pm 0.53		0.00 \pm 0.00		4.47 \pm 0.43		4.47 \pm 0.43
interneads	55.92 \pm 2.66	57.45 \pm 0.61	57.77 \pm 1.00	58.26 \pm 0.50	49.02 \pm 1.45	56.68 \pm 0.94		64.64 \pm 0.89		64.64 \pm 0.89		64.95 \pm 0.24		41.85 \pm 0.00		56.50 \pm 0.87		56.50 \pm 0.87
ionosphere	92.60 \pm 4.66	91.41 \pm 4.67	93.86 \pm 2.13	91.41 \pm 4.67	94.48 \pm 2.13	94.48 \pm 1.56		89.67 \pm 2.21		89.67 \pm 1.44		89.38 \pm 2.21		79.44 \pm 2.07		79.23 \pm 4.14		79.23 \pm 4.14
landsat	6.29 \pm 1.24	6.21 \pm 1.24	6.21 \pm 1.24	6.21 \pm 1.24	54.51 \pm 0.52	54.25 \pm 0.74		54.25 \pm 0.74		54.25 \pm 0.74		54.25 \pm 0.74		54.25 \pm 0.74		54.25 \pm 0.74		54.25 \pm 0.74
letter	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00		77.11 \pm 6.79		79.84 \pm 1.51		79.84 \pm 1.51
lymphography	64.88 \pm 2.98	69.99 \pm 2.98	69.99 \pm 2.98	70.70 \pm 1.87	70.70 \pm 1.87	68.83 \pm 2.00		71.64 \pm 0.93		69.34 \pm 1.08		69.34 \pm 1.08		74.75 \pm 7.97		76.96 \pm 0.00		76.96 \pm 0.00
magic gamma	36.83 \pm 4.76	38.15 \pm 8.41	27.69 \pm 1.80	29.08 \pm 2.63	18.51 \pm 1.23	29.94 \pm 2.59		30.62 \pm 2.59		30.62 \pm 2.59		32.69 \pm 2.41		30.62 \pm 2.59		38.36 \pm 1.69		38.36 \pm 1.69
mammography	55.37 \pm 1.84	46.20 \pm 3.18	50.54 \pm 1.26	50.54 \pm 1.26	52.66 \pm 2.09	52.79 \pm 0.77		53.76 \pm 2.21		53.76 \pm 2.21		53.76 \pm 2.21		53.76 \pm 2.21		28.42 \pm 1.40		28.42 \pm 1.40
mnist	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00	100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00		100.00 \pm 0.00
musik	70.28 \pm 1.68	41.00 \pm 3.49	44.67 \pm 5.56	44.67 \pm 5.56	41.72 \pm 5.53	46.00 \pm 5.11		46.69 \pm 2.00		46.69 \pm 2.00		46.69 \pm 2.00		10.00 \pm 0.00		60.00 \pm 0.00		60.00 \pm 0.00
optdigits	46.03 \pm 1.84	62.31 \pm 2.96	62.31 \pm 1.65	63.03 \pm 1.15	63.03 \pm 1.34	60.50 \pm 0.62		62.93 \pm 0.29		62.93 \pm 0.29		62.93 \pm 0.29		62.70 \pm 0.40		49.33 \pm 0.37		49.33 \pm 0.37
pendigits	76.77 \pm 1.23	76.83 \pm 1.81	63.36 \pm 1.23	76.83 \pm 1.81	76.83 \pm 1.81	76.83 \pm 1.81		71.75 \pm 1.27		71.75 \pm 1.27		71.75 \pm 1.27		71.75 \pm 1.27		65.74 \pm 2.10		65.74 \pm 2.10
satellite	65.94 \pm 0.44	73.29 \pm 2.03	76.27 \pm 0.81	76.27 \pm 0.81	76.93 \pm 3.26	76.93 \pm 0.81		77.94 \pm 0.28		77.94 \pm 0.28		77.94 \pm 0.28		77.94 \pm 0.28		81.04 \pm 1.96		81.04 \pm 1.96
satimage-2	38.03 \pm 4.59	72.77 \pm 3.06	98.83 \pm 0.14	98.83 \pm 0.14	98.91 \pm 0.56	98.91 \pm 0.12		98.93 \pm 0.21		98.93 \pm 0.21		98.93 \pm 0.21		98.93 \pm 0.21		42.31 \pm 7.14		42.31 \pm 7.14
shuttle	97.17 \pm 1.11	98.22 \pm 0.10	97.17 \pm 1.11	98.22 \pm 0.10	97.99 \pm 0.12	98.38 \pm 0.12		98.31 \pm 0.23		98.31 \pm 0.23		98.31 \pm 0.23		92.83 \pm 2.84		34.87 \pm 1.48		34.87 \pm 1.48
skin	38.03 \pm 2.70	38.03 \pm 2.70	67.99 \pm 9.92	54.29 \pm 12.30	54.49 \pm 13.81	54.49 \pm 13.81		56.43 \pm 9.55		56.43 \pm 9.55		56.43 \pm 9.55		81.66 \pm 0.57		76.25 \pm 3.32		76.25 \pm 3.32

Table 12.1: Average AUROC \pm standard dev. over five seeds for the semi-supervised setting on ADBench. Rank of each model among 32 models (26 baselines + 4 avg variants of top-4 baselines + 2 FoMo-0D variants w/ $D = 100$ and $D = 20$) per dataset is provided (in parentheses) (the lower, the better). We use blue and green respectively to mark the top-1 and the top-2 method. Last four rows show avg_rank of methods across datasets, and p -values of the Wilcoxon signed rank test comparing FoMo-0D ($D = 100$) with other baselines. The previous four rows are the same for FoMo-0D ($D = 20$), when ranking 31 models (26 baselines + 4 avg variants of top-4 baselines + FoMo-0D w/ $D = 20$).

Table 1.2.2: Average AUROC \pm standard dev. over five seeds for the semi-supervised setting on ADBench. Rank of each model per dataset is provided (in parentheses) (the lower, the better). We use blue and green respectively to mark the top-1 and the top-2 method.

Table 13.1: Average AUPR \pm standard dev. over five seeds for the semi-supervised setting on ADBench. Rank of each model per dataset is provided (in parentheses) (the lower, the better). We use blue and green respectively to mark the top-1 and the top-2 method.

Table 13.2: Average AUPR \pm standard dev. over five seeds for the semi-supervised setting on ADBench. Results (the lower, the better). We use blue and green respectively to mark the top-1 and the top-2 method.

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051

Table 14.1: Average F1 score \pm standard dev. over five seeds for the semi-supervised setting on AD-Bench. Rank of each model per dataset is provided (in parentheses) (the lower, the better). We use blue and green respectively to mark the top-1 and the top-2 method.

Dataset	FMDAD@100	FMDAD@200	DTL-SP	KNN	ICL	DT-C	LOF	CBLOF	SLAD	DPBM	FeatureBagging	SLAD	DT-EG	OCSVM	IForest	MCD
adult	782.43(14)	6.49±0.71(2)	5.82±0.07(17)	4.91±0.56(22)	4.2±0.26(25)	8.16±0.03(5)	67.94±0.08(10)	8.93±0.57(2)	5.32±0.11(19)	6.76±0.19(9)	7.29±0.08(8)	5.12±0.08(21)	4.2±0.26(26.5)	3.41±0.43(1)		
amazon	117.23±2.24(4)	10.7±2.1(20.6)	5.15±0.2(28)	1.8±0.1(5)	10.64±0.2(23)	11.28±0.64(10)	11.28±0.4(11)	11.28±0.2(22.1)	10.0±0.4(11)	11.3±0.2(39)	10.0±0.4(10)	11.28±0.64(10)	11.28±0.2(37.20)	10.37±0.2(39)		
ambryoid	49.25±3.0(24)	56.18±1.9(16)	61.99±0.9(3)	49.44±3.8(28.3)	7.7±2.5(4)	49.63±0.6(22)	56.7±3.3(21)	56.67±2.5(76.8)	65.99±0.6(62)	57.23±2.5(61.0)	53.56±0.6(15)	48.8±0.4(25)	48.8±0.4(25)	48.8±0.4(25)	48.8±0.4(25)	
babuor	48.12±1.1(9.1)	12.8±1.5(17.1)	9.68±1.1(12.6)	9.57±1.1(19.7)	9.51±0.9(18.7)	8.88±1.1(18.0)	72.42±2.5(1.8)	77.44±1.1(20.5)	8.53±1.1(53)	8.66±1.1(20.5)	8.66±1.1(15.5)	9.8±0.4(16.4)	9.8±0.4(16.4)	9.8±0.4(16.4)	9.8±0.4(16.4)	
blastew	59.93±2.1(1.9)	40.6±1.4(9.4)	40.6±1.4(9.2)	40.6±1.4(9.2)	40.6±1.4(9.2)	40.6±1.4(9.2)	42.2±0.6(26)	85.43±0.3(24.7)	49.83±0.3(24.0)	49.73±0.3(24.0)	49.73±0.3(24.0)	48.73±0.3(19.6)	48.73±0.3(19.6)	48.73±0.3(19.6)	48.73±0.3(19.6)	
campaign	39.72±1.6(2.4)	39.72±1.6(2.4)	39.72±1.6(2.4)	39.72±1.6(2.4)	39.72±1.6(2.4)	39.72±1.6(2.4)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	42.2±0.6(26)	
cartographicraphy	56.22±3.0(19.5)	46.35±1.0(19.5)	46.35±1.0(19.5)	46.35±1.0(19.5)	46.35±1.0(19.5)	46.35±1.0(19.5)	48.25±0.6(23)	38.37±2.3(25)	48.25±0.6(23)	48.25±0.6(23)	48.25±0.6(23)	48.25±0.6(23)	48.25±0.6(23)	48.25±0.6(23)	48.25±0.6(23)	
celiba	84.1±0.8(8.2)	48.4±1.1(8.1)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	17.19±1.4(8.2)	
census	69.92±2.7(24.6)	69.53±1.0(8.0)	22.2±0.5(14.5)	23.96±0.5(14.6)	23.96±0.5(14.6)	23.96±0.5(14.6)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	21.46±0.2(23)	
cover	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	60.57±1.1(2.7)	
donors	57.51±0.0(8)	53.3±2.3(22.0)	53.3±2.3(22.0)	53.3±2.3(22.0)	53.3±2.3(22.0)	53.3±2.3(22.0)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	56.4±0.3(52)	
fnau	61.34±1.8(6)	44.6±1.8(4.6)	44.6±1.8(4.6)	44.6±1.8(4.6)	44.6±1.8(4.6)	44.6±1.8(4.6)	48.36±5.4(11.3)	45.22±4.7(5.8)	45.22±4.7(5.8)	45.22±4.7(5.8)	45.22±4.7(5.8)	45.22±4.7(5.8)	45.22±4.7(5.8)	45.22±4.7(5.8)	45.22±4.7(5.8)	
fraud	77.6±6.0(7.3)	59.9±6.0(10.6)	56.2±1.1(12.0)	56.2±1.1(12.0)	56.2±1.1(12.0)	56.2±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	57.59±1.1(12.0)	
glass	99.64±0.7(1.2)	97.43±1.3(3.7)	97.43±1.3(3.7)	97.43±1.3(3.7)	97.43±1.3(3.7)	97.43±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	98.46±1.3(3.7)	
hepatitis	62.0±0.0(1.0)	60.7±0.0(1.0)	60.7±0.0(1.0)	60.7±0.0(1.0)	60.7±0.0(1.0)	60.7±0.0(1.0)	69.64±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	69.63±0.0(2.2)	
hipp	89.24±5.7(9.1)	60.1±1.9(4.6)	60.1±1.9(4.6)	60.1±1.9(4.6)	60.1±1.9(4.6)	60.1±1.9(4.6)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	70.96±0.0(11.0)	
imdb	5.4±0.0(2.6)	10.65±0.0(4.2)	7.24±1.6(12)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	6.4±0.0(18.0)	
intentsards	40.83±2.7(48.5)	33.83±1.3(72.0)	65.1±2.8(59.0)	70.93±1.1(32.0)	70.93±1.1(32.0)	70.93±1.1(32.0)	71.04±0.9(46.8)	7.41±1.0(6.2)	13.99±1.0(6.2)	48.46±1.0(44.1)	56.1±1.4(41)	56.86±1.0(44.1)	56.86±1.0(44.1)	56.86±1.0(44.1)	56.86±1.0(44.1)	56.86±1.0(44.1)
ionosphere	97.98±4.0(6.6)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	92.2±2.8(8.9)	
landesk	57.51±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	56.4±0.0(19)	
letter	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	96.02±0.0(1.0)	
lymphography	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	96.32±0.5(1.4)	
mammography	77.83±0.9(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	76.46±0.6(2.6)	
mnist	62.0±0.0(14)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	42.84±0.0(8.5)	
mask	93.2±3.2(30.8)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	91.21±0.7(17.0)	
origami	37.33±0.0(7.7)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	22.47±1.2(27.1)	
pageblocks	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	61.18±0.8(0.0)	
pendigits	69.23±2.0(8.2)	80.1±2.0(6.0)	80.1±2.0(6.0)	80.1±2.0(6.0)	80.1±2.0(6.0)	80.1±2.0(6.0)	74.69±1.8(8.2)	70.56±2.2(29.8)	73.54±2.3(23.0)	66.73±3.4(25.2)	66.73±3.4(25.2)	66.73±3.4(25.2)	66.73±3.4(25.2)	66.73±3.4(25.2)	66.73±3.4(25.2)	
pima	7.65±1.8(8.2)	7.42±1.6(6.5)	7.21±1.1(6.5)	7.19±0.9(6.6)	7.19±0.9(6.6)	7.19±0.9(6.6)	74.99±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	72.33±0.0(46.5)	
satimage-2	88.83±0.0(1.5)	88.77±0.0(1.5)	88.77±0.0(1.5)	88.77±0.0(1.5)	88.77±0.0(1.5)	88.77±0.0(1.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	88.45±0.0(7.5)	
shuttle	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	77.73±0.0(2.9)	
spambase	67.97±2.0(2.6)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	68.94±4.4(9.2)	
stamps	8.89±1.7(7.2)	8.89±1.7(7.2)	8.89±1.7(7.2)	8.89±1.7(7.2)	8.89±1.7(7.2)	8.89±1.7(7.2)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	77.47±0.9(45.8)	
thyroid	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	59.14±0.0(2.5)	
venetian	16.21±0.5(1.6)	16.21±0.5(1.6)	16.21±0.5(1.6)	16.21±0.5(1.6)	16.21±0.5(1.6)	16.21±0.5(1.6)	18.88±0.0(1.5)	28.82±0.2(19.4)	28.82±0.2(19.4)	28.82±0.2(19.4)	28.82±0.2(19.4)	28.82±0.2(19.4)	28.82±0.2(19.4)	28.82±0.2(19.4)	28.82±0.2(19.4)	
vowels	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	26.03±0.0(15.5)	
waveform	67.65±2.8(2.0)	67.65±2.8(2.0)	67.65±2.8(2.0)	67.65±2.8(2.0)	67.65±2.8(2.0)	67.65±2.8(2.0)	68.95±1.1(12.3)	5.65±1.1(5.2)	68.95±1.1(5.2)	68.95±1.1(5.2)	68.95±1.1(5.2)	68.95±1.1(5.2)	68.95±1.1(5.2)	68.95±1.1(5.2)	68.95±1.1(5.2)	
webspam	96.53±3.2(1.0)	87.93±2.0(6.0)	87.93±2.0(6.													

Table 14.2: Average F1 score \pm standard dev. over five seeds for the semi-supervised setting on ADBench. Rank of each model per dataset is provided (in parentheses) (the lower, the better). We use blue and green respectively to mark the top-1 and the top-2 method.

2106 **H BENCHMARK OD DATASETS**

2107

2108 Table 15: Description of all datasets in ADBench Livernoche et al. (2024).

2109

Dataset Name	# Samples	# Features	# Anomaly	% Anomaly	Category
ALOI	49534	27	1508	3.04	Image
annthyroid	7200	6	534	7.42	Healthcare
backdoor	95329	196	2329	2.44	Network
breastw	683	9	239	34.99	Healthcare
campaign	41188	62	4640	11.27	Finance
cardio	1831	21	176	9.61	Healthcare
Cardiotocography	2114	21	466	22.04	Healthcare
celeba	202599	39	4547	2.24	Image
census	299285	500	18568	6.20	Sociology
cover	286048	10	2747	0.96	Botany
donors	619326	10	36710	5.93	Sociology
fault	1941	27	673	34.67	Physical
fraud	284807	29	492	0.17	Finance
glass	214	7	9	4.21	Forensic
Hepatitis	80	19	13	16.25	Healthcare
http	567498	3	2211	0.39	Web
InternetAds	1966	1555	368	18.72	Image
Ionosphere	351	32	126	35.90	Oryctognosy
landsat	6435	36	1333	20.71	Astronautics
letter	1600	32	100	6.25	Image
Lymphography	148	18	6	4.05	Healthcare
magic.gamma	19020	10	6688	35.16	Physical
mammography	11183	6	260	2.32	Healthcare
mnist	7603	100	700	9.21	Image
musk	3062	166	97	3.17	Chemistry
optdigits	5216	64	150	2.88	Image
PageBlocks	5393	10	510	9.46	Document
pendigits	6870	16	156	2.27	Image
Pima	768	8	268	34.90	Healthcare
satellite	6435	36	2036	31.64	Astronautics
satimage-2	5803	36	71	1.22	Astronautics
shuttle	49097	9	3511	7.15	Astronautics
skin	245057	3	50859	20.75	Image
smtp	95156	3	30	0.03	Web
SpamBase	4207	57	1679	39.91	Document
speech	3686	400	61	1.65	Linguistics
Stamps	340	9	31	9.12	Document
thyroid	3772	6	93	2.47	Healthcare
vertebral	240	6	30	12.50	Biology
vowels	1456	12	50	3.43	Linguistics
Waveform	3443	21	100	2.90	Physics
WBC	223	9	10	4.48	Healthcare
WDBC	367	30	10	2.72	Healthcare
Wilt	4819	5	257	5.33	Botany
wine	129	13	10	7.75	Chemistry
WPBC	198	33	47	23.74	Healthcare
yeast	1484	8	507	34.16	Biology
CIFAR10	5263	512	263	5.00	Image
FashionMNIST	6315	512	315	5.00	Image
MNIST-C	10000	512	500	5.00	Image
MVTec-AD	5354	512	1258	23.50	Image
SVHN	5208	512	260	5.00	Image
Agnews	10000	768	500	5.00	NLP
Amazon	10000	768	500	5.00	NLP
Imdb	10000	768	500	5.00	NLP
Yelp	10000	768	500	5.00	NLP
20newsgroups	11905	768	591	4.96	NLP

2156
2157
2158
2159

2160 I DIFFERENCES TO PRIOR WORK ON PFNs FOR TABULAR DATA 2161

2162 There exist applications of PFNs (originally developed by Müller et al. (2022)) that pre-date our
2163 proposed FoMo-OD, namely, TabPFN (Hollmann et al., 2023) for supervised classification, LC-PFN
2164 (Adriaensen et al., 2024) for learning curve extrapolation, PFN4BO (Müller et al., 2023) for Bayesian
2165 optimization, and ForecastPFN (Dooley et al., 2023) for time series forecasting.

2166 Here we highlight the differences of our proposed FoMo-OD from these existing PFNs.
2167

- 2168 1. **First PFN4OD:** We employ prior-data fitted networks (PFNs) for outlier detection (OD) for
2169 the first time.
- 2170 2. **First large-scale pretrained OD model:** FoMo-OD is the first model for zero-shot OD that
2171 is pretrained at large scale on a large collection of (synthetic) datasets, due to the minuscule
2172 nature of existing real-world OD benchmark datasets.
- 2173 3. **New data prior:** Thanks to PFN’s reliance on synthetically generated datasets, we establish
2174 a new data prior for OD, specifically for outlier synthesis.
- 2175 4. **Data transformation for scale:** While drawing samples from a data prior may be relatively
2176 fast, pretraining a large foundation model requires many such draws for every step of each
2177 epoch. To speed up data synthesis on-the-fly, we are the first to leverage a linear transformation.
- 2178 5. **Router-based attention for scale:** PFNs ingest the entire training dataset as context for in-
2179 context learning at inference time. To accommodate larger datasets at both training (for better
2180 generalization) and inference (for large-scale real-world datasets), we leveraged a “bottleneck”
2181 architecture for scalable self-attention, and in turn, larger context size.

2182 J RELATED WORK 2183

2184 **Outlier Detection (OD):** Thanks to diverse applications in numerous fields, such as security, finance,
2185 manufacturing, to name a few, OD on tabular (or point-cloud) datasets has a vast literature with a
2186 long list of techniques. For earlier, shallow approaches preceding the advances in deep learning, we
2187 refer to the books by Aggarwal (2013) and Aggarwal and Sathe (2017). The modern, deep learning
2188 based techniques are surveyed in Chalapathy and Chawla (2019); Pang et al. (2021); Ruff et al.
2189 (2021). Most recent deep OD techniques take advantage of newly emerging paradigms, including
2190 self-supervised learning (Hojjati et al., 2022; Yoo et al., 2023) as well as the most recently popularized
2191 diffusion-based models (Yoon et al., 2023; Livernoche et al., 2024; Du et al., 2024; He et al., 2024).
2192

2193 **Unsupervised Model Selection for OD:** It is typical of models to exhibit various hyperparameters
2194 (HPs) that play a role in the bias-variance trade-off and hence the generalization performance, and OD
2195 models are no exception. Many earlier work on OD have showcased the sensitivity of classical (i.e.
2196 shallow) OD methods to the choice of their HP(s) (Aggarwal and Sathe, 2015; Campos et al., 2016;
2197 Goldstein and Uchida, 2016). Similarly, sensitivity to HPs has also been shown for deep OD models
2198 more recently (Zhao et al., 2021; Ding et al., 2022), as well as for those relying on self-supervised
2199 learning/data augmentation (Yoo et al., 2023).

2200 While critical, work on unsupervised outlier model selection (UOMS) is slim as compared to the vast
2201 literature on detection methods. A handful of existing, mostly heuristic strategies has been studied by
2202 Ma et al. (2023) reporting discouraging results; they have shown that existing heuristics are either not
2203 significantly different from random selection, or do not outperform iForest (Liu et al., 2008) with its
2204 default HPs (an extremely fast ensemble of randomized trees).

2205 More recent, state-of-the-art (SOTA) UOMS approaches go beyond heuristic measures and instead
2206 design scalable hyperensembles (Ding et al., 2022; 2024), as well as take advantage of meta-learning
2207 on historical real-world OD datasets (Zhao et al., 2021; 2022; Zhao and Akoglu, 2024). These
2208 SOTA approaches demonstrate the value of learning from many other OD datasets, and transfer these
2209 learnings to a new dataset. While sharing the same spirit on learning from a large collection of (in
2210 our case, simulated) datasets, our FoMo-OD differs from these prior art in a key aspect; FoMo-OD is
2211 *not* a model selection technique, but rather, a foundation model that abolishes model training and
2212 selection altogether and unlocks zero-shot inference on a new dataset.

2213 **Prior-data Fitted Networks:** Based on the seminal work by Müller et al. (2022), Prior-data-fitted
2214 Networks (PFNs) establish a new paradigm for machine learning, where a PFN is pretrained on

synthetic datasets generated from a data prior, and the pretrained PFN can then infer the posterior predictive distribution (PPD) for test points in a new dataset in a single forward pass, through in-context learning (Xie et al., 2021; Garg et al., 2022). It is shown that PFNs provably approximate Bayesian inference (Müller et al., 2022). Follow-up TabPFN (Hollmann et al., 2023) achieved SOTA classification performance on small tabular datasets of size up to 1024. Other subsequent works designed LC-PFN (Adriaensen et al., 2024) and ForecastPFN (Dooley et al., 2023), respectively zero-shot learning curve extrapolation and zero-shot time-series forecasting models, trained purely on synthetic data. PFN4BO (Müller et al., 2023) employed PFNs for Bayesian optimization, while Nagler (2023) studied the statistical foundations of PFNs. As training data is passed as context to PFN, others proposed scaling solutions to enable training on larger pretraining datasets for better generalization (Ma et al., 2024; Feuer et al., 2023; 2024).

Our proposed FoMo-OD differs from these in being the first PFN for OD, using a novel inlier/outlier data prior, employing linear transform for fast data synthesis, and incorporating the “router” attention mechanism for linear-time scalability w.r.t. context size. See Appendix I for additional details.

Zero-Shot Outlier Detection: Foundation models pretrained on massive text and image corpora, such as large language and/or vision models (L(V)LMs) like OpenAI’s GPT-series (Achiam et al., 2023), DALL-E (Ramesh et al., 2021) and Flamingo (Alayrac et al., 2022), CLIP (Radford et al., 2021), and LLaVA (Liu et al., 2024) to name a few, have demonstrated remarkable success on several zero-shot tasks in CV and NLP. Follow-up work extended these models for zero-shot out-of-distribution detection (Esmailpour et al., 2022), zero-shot image OD (Liznerski et al., 2022; Jeong et al., 2023) as well as dialogue-based industrial image anomaly detection (Gu et al., 2024).

Foundation models, however, do not exist for tabular data which is widespread across OD applications in the real world, such as detecting credit card fraud, network intrusion, medical anomalies, and any sensor measurement abnormalities, to name a few. The recent ACR model by Li et al. (2023) on zero-shot OD does *not* rely on a pretrained foundation model, but rather is meta-trained on each specific domain using inlier-only datasets from the *same domain*. Concurrent to our work, Li et al. (2024) apply pretrained LLMs for prompt-based OD on tabular data which they serialize to text. Similar to our work, they also use *simulated* labeled OD datasets to fine-tune several existing LLMs to improve their performance. Their work, however, is quite preliminary in several fronts; a key limitation is that they assume independent features and query the LLM one-feature-at-a-time to reach an outlier score. Further, they fine-tune using only 5,000 data batches with up to 100 samples each, subsample 150 points and the first 10 columns of each dataset for evaluation (due to GPU memory constraint), and their testbed includes only two baseline methods. In contrast, FoMo-OD employs and pretrains PFNs at a much larger scale with rigorous evaluation on a much larger testbed.

K DISCUSSION

Summary: We introduced FoMo-OD, **the first foundation model for outlier detection (OD)** on tabular data. FoMo-OD is a prior-data fitted network (PFN), pretrained on a large number of *synthetic* datasets generated from a new data prior for OD, which can infer the posterior predictive distribution for test points in a new dataset in a **zero-shot** fashion where the training data is input as context, capitalizing on *in-context learning*.

Zero-shot OD implies no more OD model (parameter) training and **no more model selection**, given a new OD task. That is a revolution for OD (!), for which algorithm and hyperparameter selection are notoriously-hard *without any labeled data*, and also computationally taxing especially for today’s modern deep OD models with numerous parameters *and* a long list of hyperparameters. What is more, FoMo-OD provides **extremely fast inference** thanks to a mere *single forward pass*, making it amenable for OD on data streams.

Building on the PFN paradigm (Müller et al., 2022), FoMo-OD breaks new ground not only conceptually by abolishing the burden of model training and selection, but also empirically: Against **26** different (both classical and modern) baselines on **57** public benchmark datasets from diverse domains, FoMo-OD performs on par with the top *2nd* baseline, while significantly outperforming the majority of the baselines. Without the need to train any, let alone multiple models for HP tuning, FoMo-OD takes a mere **7.7 ms** per test sample for inference only.

2268
 2269 **Limitations and Future Directions:** FoMo-OD employs a simple straightforward data prior based
 2270 on GMMs. While it is remarkable to see how far one can go with synthetic data from such a simple
 2271 prior, future work can design more comprehensive data priors, inclusive of discrete features as well
 2272 as other possible outlier types. We have also pretrained FoMo-OD solely on synthetic datasets, while
 2273 future work can augment both synthetic and real-world datasets for pretraining.

2274 Besides the lack of massive real-world datasets for tabular OD, a motivation for a data prior to pretrain
 2275 purely on synthetic datasets comes from neural scaling laws (Kaplan et al., 2020; Zhai et al., 2022).
 2276 Interestingly, the scaling laws for large Transformer models have shown that their generalization
 2277 error tends to drop as a power law with the amount of training data (also, with number of parameters
 2278 and amount of compute), but the power law exponent is very small—suggesting that acquiring more
 2279 colossal real-world datasets would be a slow, if not expensive approach to advancing ML/AI. Others
 2280 have proposed ways to subset-select smaller, non-redundant “foundation datasets” (Sorscher et al.,
 2281 2022; Paul et al., 2021), and emphasized the importance of task/dataset diversity in pretraining
 2282 (Raventós et al., 2024). Arguably, synthetic data from a complex and diverse data prior is a potential
 2283 gateway to obtaining non-redundant and diverse datasets for pretraining large foundation models like
 2284 FoMo-OD. On the other hand, designing such a data prior requires a level of domain/prior knowledge.
 2285

2286 Another improvement could be scaling up to even larger context (i.e. dataset) size and dimensionality.
 2287 While FoMo-OD generalizes beyond pretrained context sizes and dimensionality, it is limited to and
 2288 performs particularly well on downstream datasets of similar nature as our experiments showed. A
 2289 promising direction for size generalization is using PFNs as extremely fast ensemble components
 2290 at inference; since “*PFNs are quick enough to be used as ensemble members. The size constraints*
 2291 *could therefore be overcome by boosting and bagging techniques*” (Nagler, 2023).

2292 Further, our work focused on semi-supervised OD with clean/inlier-only training data. Future work
 2293 can study the unsupervised OD setting and pretraining with mixed/“contaminated” data in this
 2294 transductive setting, where the unlabeled test data is the same as training data. In addition, we
 2295 performed offline evaluation of FoMo-OD on static datasets, while its fast inference lends itself to
 2296 streaming OD, which future work can explore. Technically, both extensions (unsupervised OD and
 2297 streaming OD) are straightforward from the implementation perspective.

2298 Our current work is limited to OD for tabular (or point-cloud) data. Our ideas can be extended to
 2299 other data modalities, such as image, graph, and text outliers, to comprise other domains with critical
 2300 OD applications such as video surveillance, fraud detection and LLM hallucination detection. To
 2301 that end, the design of novel inlier/outlier priors would be an open direction. A promising approach
 2302 here could be the use of pretrained generative models to draw synthesized image/text/etc. datasets for
 2303 pretraining the PFN, in place of manually-designed data priors.

2304 Finally, our quest here has been mainly experimental. Theoretically understanding why these models
 2305 work as well as they do and investigating their failure cases are important yet open questions.

2306 As the first foundation model for OD, FoMo-OD inspires many promising directions for future research
 2307 that could lead to fruition for additional practical applications.

2308 L BROADER IMPACT STATEMENT

2309
 2310 FoMo-OD is zero-shot, abolishing not only parameter training but also model selection given a new
 2311 dataset. This is a radical paradigm shift for OD literature, which historically focused on designing
 2312 new models and recently also effective ways for unsupervised model selection. Obviating the need
 2313 for either, we expect FoMo-OD to route attention of the community from new OD model design and
 2314 selection to designing better data priors and gathering datasets for PFN pretraining, along with better
 2315 and more scalable architectures for PFN.

2316 From the applied perspective, a zero-shot OD model like FoMo-OD is a game-changer for practitioners!
 2317 Given the plethora of OD algorithms to choose from, which often come with a list of hyperparameters
 2318 to set, and not having the tools for effective and efficient model selection, the practitioners are
 2319 burdened with a “choice paralysis”. With FoMo-OD, practitioners can not only bypass such dilemmas
 2320 on one dataset, but thanks to the “train once, use many times” nature of pretrained models, they can
 2321 do so for any dataset such as those arriving over time. In fact, provided its lightening-fast inference

2322 via a single forward pass, FoMo-OD is amenable to deploy in real time on streaming datasets, such
2323 that each (test) sample over a stream can be inferred with the preceding samples passed as context.
2324

2325 M REPRODUCIBILITY STATEMENT

2327 We expect that the disruptive nature of FoMo-OD will trigger future innovations in the OD literature,
2328 as well as a widespread adoption by practitioners thanks to its key desirable properties. To
2329 foster future research and accessibility in practice, we make all resources (our codebase used for
2330 prior data synthesis, data transformation, and pretraining as well as our pretrained model check-
2331 points) publicly available at <https://anonymous.4open.science/r/PFN40D>. Further,
2332 full implementation details are provided in Appendix C.

2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375