TIME-LAPSE: LEARNING TO SAY "I DON'T KNOW" THROUGH SPATIO-TEMPORAL UNCERTAINTY SCORING

Anonymous authors
Paper under double-blind review

ABSTRACT

Safe deployment of trained ML models requires determining when input samples go out-of-distribution (OOD) and refraining from making uncertain predictions on them. Existing approaches inspect test samples in isolation to estimate their corresponding predictive uncertainty. However, in the real-world, deployed models typically see test inputs consecutively and predict labels continuously over time during inference. In this work, we propose TIME-LAPSE, a spatio-temporal framework for uncertainty scoring that examines the sequence of predictions prior to the current sample to determine its predictive uncertainty. Our key insight is that in-distribution samples will be more "similar" to each other compared to OOD samples, not just over the encoding latent-space but also across time. Specifically, (a) our spatial uncertainty score estimates how different OOD latent-space representations are from those of an in-distribution set using metrics such as Mahalanobis distance and cosine similarity and (b) our temporal uncertainty score determines deviations in correlations over time using representations of past inputs in a non-parametric, sliding-window based algorithm. We evaluate TIME-LAPSE on both audio and vision tasks using public datasets and further benchmark our approach on a challenging, real-world, electroencephalograms (EEG) dataset for seizure detection. We achieve state-of-the-art results for OOD detection in the audio and EEG domain and observe considerable gains in semantically corrected vision benchmarks. We show that TIME-LAPSE is more driven by semantic content compared to other methods, i.e., it is more robust to dataset statistics. We also propose a sequential OOD detection evaluation framework to emulate real-life drift settings and show that TIME-LAPSE outperforms spatial methods significantly.

1 INTRODUCTION

Modern machine learning (ML) has seen tremendous success in various tasks across multiple domains (Bojarski et al., 2016; Hinton, 2018; Kreinovich & Kosheleva, 2020; van den Oord et al., 2016; Schrittwieser et al., 2020), surpassing human performance in many benchmarks (Esteva et al., 2017; Yala et al., 2019; Krizhevsky et al., 2012). However, deep learning models have been known to fail silently and catastrophically with highly confident predictions (Nguyen et al., 2015; Goodfellow et al., 2015; Guo et al., 2017). Such models assume a closed world scenario, i.e., they assume that variability encountered when deployed in the real-world will be similar to the variability present in their training data. In practice, they encounter an open world where incoming samples can come from shifted distributions or completely new distributions (Liu et al., 2020).

This behaviour can have severe consequences in mission-critical domains such as healthcare and autonomous driving, where errors can have a high cost resulting in injury or even death (Amodei et al., 2016). Most widely cited models do not come with a mechanism to say "I don't know" or to abstain when uncertain (Kompa et al., 2021). As predictions from deployed machine learning models are used to inform real-world decisions, it is important that they have the ability to both understand when they are out of their depth (when faced with distributional shifts or new semantic content) and abstain from prediction under such uncertainty, allowing for safer deployment.

The task of identifying when incoming inputs are drawn far from the training distribution is called out-of-distribution (OOD) detection. At its core, it is a binary classification problem and many methods approach it as such, requiring labelled examples of OOD inputs during training (Liang

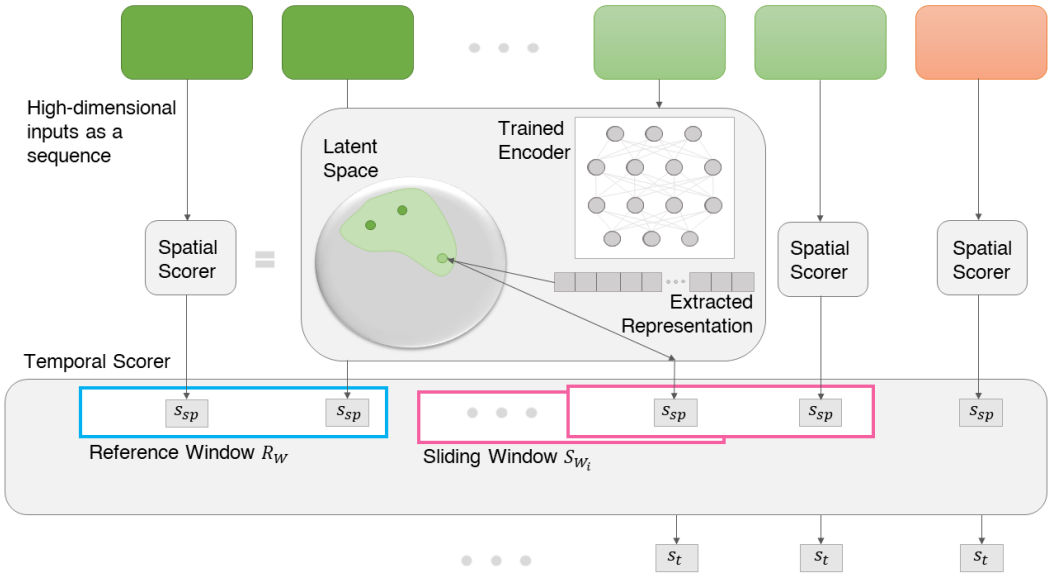

Figure 1: TIME-LAPSE: a spatio-temporal framework for uncertainty scoring. High-dimensional inputs (images, audio, speech, EEG signals, etc) are passed in sequentially during inference. Dark green indicates samples from the training set. Light green and light orange represent potential InD and OOD inputs respectively. Each input is projected on to a lower-dimensional latent space using a well-trained encoder and its spatial uncertainty score s_{sp} is extracted by the spatial scorer. The temporal scorer utilizes the sequence of spatial scores to give the final temporal uncertainty scores s_t.

et al., 2018; DeVries & Taylor, 2018). However, real-world OOD samples can be drawn from widely varying distributions and need not be restricted to the distributions specified by labelled OOD samples. Moreover, current OOD evaluation schemes hold different datasets to be OOD sources, without considering if the distinctions are semantic in nature (Ahmed & Courville, 2020) leading to methods that are insensitive to semantic overlap and biased to dataset statistics.

In this paper, we propose TIME-LAPSE, a spatio-temporal framework for the joint task of predictive uncertainty quantification and subsequent OOD detection. A key insight in our approach is that most real-world scenarios (such as electroencephalogram (EEG) seizure analysis, healthcare decision making, autonomous driving, etc.) involve consecutive inputs to the model that are likely to be correlated over time. For example, an obstacle detector deployed in a self-driving car will see images correlated over time. A seizure detector installed in a neurology clinic will process hours of time-correlated EEG signals. Even low-risk image classification models deployed in search engines over the cloud will see correlated inputs over time when grouped by user ID or location. Thus, sequentially occurring samples share meaningful semantic correlations.

In TIME-LAPSE (Fig. 1), we compute spatial uncertainty scores based on the following hypotheses: an encoder with enough inductive bias likely (a) maps OOD inputs “far" from other InD samples in the latent space under the Mahalanobis distance metric (Lee et al., 2018) and (b) produces dissimilarities between their representations as captured by the cosine similarity metric (Jones & Furnas, 1987). We compute temporal uncertainty scores by exploiting correlations between lower dimensional representations of consecutive samples over time when examined as a sequence. A sequence of InD inputs will show more correlation than a sequence with both InD and OOD inputs or a sequence with just OOD inputs.

Recent works on OOD detection limit themselves to evaluations on highly curated image benchmarks, but do not extend their tests to more diverse and realistic scenarios required for practical applications of uncertainty scoring (Ren et al., 2019; Ahmed & Courville, 2020). Such OOD image benchmarks consider different datasets to be OOD sources, neglecting any semantic overlap between their OOD test sets and InD train sets, leading to narrow model capabilities (Yang et al., 2021). We show that TIME-LAPSE is more driven by semantic content compared to other techniques, i.e. it is more

robust to dataset statistics. Moreover, current evaluation schemes do not consider different dataset drift settings that could occur in the real-world. We put forth an additional sequential evaluation framework for OOD detection to evaluate methods under realistic conditions of data drifts and show that TIME-LAPSE outperforms other methods significantly. We demonstrate the efficacy of our proposed method by comparing against multiple baselines over different domains (vision, audio and clinical EEGs). We achieve state-of-the-art (SOTA) on both audio and EEG datasets while improving performance on semantically corrected image benchmarks. To the best of our knowledge, we are the first to use spatio-temporal modelling for uncertainty quantification and subsequent OOD detection. We are also the first to use deep-learning based uncertainty scoring techniques for EEG analyses.

To summarize, our key contributions are:

- We propose TIME-LAPSE, a spatio-temporal framework for uncertainty quantification and subsequent OOD detection that uses samples' spatial similarity (in the latent space) and their temporal similarity amongst sequential inputs to determine their predictive uncertainty.
- We show that scores from TIME-LAPSE are driven by semantic content and are more robust to dataset statistics whereas popular baselines, benchmarked on standard image datasets, are susceptible to dataset statistics and overlook semantic overlap.
- We benchmark TIME-LAPSE on diverse domains: audio speech classification, seizure detection using clinical EEGs and image classification. TIME-LAPSE achieves state-of-the-art (SOTA) in both audio and EEG domains and improves on semantically corrected image benchmarks.
- We propose a sequential evaluation framework for OOD detection to evaluate methods under realistic conditions of data drifts and show that TIME-LAPSE outperforms other methods significantly. We hope this evaluation scheme will be adopted for routinely characterizing OOD detection performance in the wild.

2 Related Work

OOD detection (sometimes called outlier detection or anomaly detection) is a rich field with a long history. Methods developed for OOD detection are highly context dependent. Classical techniques include density estimation (Breunig et al., 2000), one-class SVMs (Schölkopf et al., 1999), tree-isolation forests (Liu et al., 2008), etc for one-dimensional data such as time-series. There is a vast amount of literature on change point detection in time-series data (Aminikhanghahi & Cook, 2017; Kifer et al., 2004), where simple streams of signals can undergo non-stationary or piece-wise stationary changes and the goal is to detect these changes. However, such techniques scale badly with input dimensionality (Rabanser et al., 2019). Moreno-Torres et al. (2012) give a complete overview of closely related topic of distributional shift including covariate shift, label shift and concept drift (Gama et al., 2014). More recently, OOD detection techniques with deep neural networks have become popular for high dimensional data.

Several works consider semi-supervised techniques for OOD such as outlier exposure methods (Hendrycks et al., 2019a; Ruff et al., 2020), while many assume the supervised setting where the target OOD distribution is known and propose alternative training strategies and auxillary OOD tasks (DeVries & Taylor, 2018; Hendrycks et al., 2019b; Mohseni et al., 2020; Shalev et al., 2018) in a multi-task setting. In this paper, we consider the more general setting of unsupervised OOD detection where no OOD examples are labelled as such and the encoder has access to only InD training labels.

Other methods typically use representations from classification networks to infer OOD inputs. Hendrycks & Gimpel (2017) propose using the maximum softmax probabilites (MSP) to detect OOD data. Liang et al. (2018) improve this by introducing a temperature parameter to the softmax equation. Lee et al. (2018) fit class-conditional Gaussians to intermediate activations and use the Mahalanobis distance to identify OOD samples. Other works use contrastive or other self-supervised techniques to get better representations from the classification heads for further downstream OOD detection (Tack et al., 2020; Winkens et al., 2020; Sehwag et al., 2021). Several methods directly use input likelihoods as the OOD detection score. However, studies have shown that generative techniques may be overconfident on complex inputs (Nalisnick et al., 2019). Modified likelihood scores have been proposed since then (Serrà et al., 2020; Xiao et al., 2020; Choi et al., 2019) including energy-based models (Liu et al., 2020; Du & Mordatch, 2019; Grathwohl et al., 2020) and likelihood ratios (Ren et al., 2019). Reconstruction-based OOD methods use reconstruction loss as the uncertainty score

and primarily use auto-encoders (Zong et al., 2018; Pidhorskyi et al., 2018) or GANs (Schlegl et al., 2017; Deecke et al., 2019; Perera et al., 2019). Bayesian approaches such as Blundell et al. (2015); Malinin & Gales (2018); Chen et al. (2019) learn distributions over model weights, while Gal & Ghahramani (2016) use Monte-Carlo dropout sampling to derive uncertainty estimates on OOD data. Note that Monte-Carlo test-time dropout requires k forward passes for each instance during inference time and is computationally more expensive.

3 TIME-LAPSE: SPATIO-TEMPORAL UNCERTAINTY SCORING FRAMEWORK

We propose TIME-LAPSE (**la**tent-**s**pace **e**mbeddings over **time**) for the joint task of uncertainty scoring and OOD detection. We start by defining the problem setup and background notation. We consider a multi-class classification setting here, though our framework can be extended to other scenarios such as regression, segmentation, etc.

3.1 PROBLEM SETUP

Let $\mathcal{X}$ represent our high-dimensional input space, $\mathcal{X} \subseteq \mathbb{R}^n$. Let $\mathcal{Y} = \{0, 1, 2, ..., C-1\}$ denote the label space where C is the number of classes. In standard multi-class classification, we learn a classifier $f : \mathcal{X} \mapsto \mathcal{Y}$ using a dataset D_{train} (assumed to be sampled from an underlying distribution p^*) such that $f(x) = p(Y = y_i|x)$, where $x \in \mathcal{X}$ and $y_i \in \mathcal{Y} \; \forall \; i$. The final prediction for an unseen input x^* is given by $\hat{y} = \arg\max_{y_i} p(y_i|x^*)$. When the deployed classifier encounters inputs that are generated from a different distribution with respect to its estimated distribution during inference, its predictions cannot be trusted. The goal of uncertainty scoring and OOD detection is to give high uncertainty scores to such inputs and filter them out using a threshold.

Given a sequence of unseen inputs, x_t (where t denotes time) to a classifier, our goal is to output a score $s(x_t)$ and a selective function $g(x_t) \; \forall \; x_t$ such that

$$(g, s)(x_t) := \begin{cases} \hat{y} = \arg\max_{y_i} p(y_i|x_t), & \text{if } \; s(x_t) \leq \text{th} \\ \text{ABSTAIN or FLAG AS OOD}, & \text{else} \end{cases}$$

We call the score $s(x_t)$ the associated uncertainty score of input sample x_t

3.2 TIME-LAPSE FRAMEWORK

We use the learnt classifier f to derive an encoder $h : \mathcal{X} \mapsto \mathcal{U}$ that maps high-dimensional inputs $x \in \mathcal{X} \subseteq \mathbb{R}^n$ onto a latent space $\mathcal{U} \subseteq \mathbb{R}^d$ through its intermediate representations. Note that h is inherently learnt when f is trained using the InD training dataset D_{train}. TIME-LAPSE is agnostic to the training method, which can be supervised, self-supervised or unsupervised depending on how the label information in D_{train} is utilized.

3.2.1 LEARNING SPATIAL UNCERTAINTY SCORES

To extract the spatial uncertainty score of an input test sample's prediction, we apply two fundamental hypotheses: (a) a well-trained encoder maps InD samples onto a dense region in the latent space but maps OOD inputs outside this region under a distance metric d with high likelihood, and (b) the OOD inputs that get mapped to the latent space do not share similarity with InD inputs under a similarity metric, sim.

To identify spatially distinguishable OOD inputs, we first model the extent of InD region $\mathfrak{D}$ within the fixed-dimensional latent space, using the InD samples. We extract a coreset by sampling from our InD training data, $\mathcal{D}_{train}$.

$$\text{coreset} = \{h(x_i) \mid x_i \sim \mathcal{D}_{train}\}; \quad |\mathcal{D}_{train}| \geq |\text{coreset}|$$

We compute the distance score s_{dist} and the similarity score s_{sim} of the unseen test sample x by comparing its latent representation to that of samples in the coreset using a distance metric d and a similarity metric sim such that:

$$s_{\text{dist}}(x) = \min_{h(x_i) \in \text{ coreset}} d(h(x_i), h(x)) \; ; \quad s_{\text{sim}}(x) = \max_{h(x_i) \in \text{ coreset}} \text{sim}(h(x_i), h(x))$$

Finally, the combined spatial uncertainty score is given by

$$s_{\text{spatial}}(x) = s_{\text{dist}}(x) \; . \; s_{\text{sim}}(x)$$

We choose the distance metric d to be the Mahalanobis distance obtained using class-conditional Gaussians (Lee et al., 2018) without label smoothing (Winkens et al., 2020) where class-wise means μ_c and covariance matries Σ_c are estimated from the coreset.

$$d(h(x)) = \min_c \, (h(x) - \mu_c)^T \, \Sigma_c^{-1} \, (h(x) - \mu_c)$$

As similarity metric sim, we adopt the cosine similarity between the test input's representation and the coreset.

$$\text{sim}(h(x), h(x_i)) = \frac{h(x) \, . \, h(x_i)}{||h(x)|| \; ||h(x_i)||}; \quad h(x_i) \in \text{coreset}$$

We note that TIME-LAPSE will accept any d and sim that is most suited for the target domain.

3.2.2 Learning Temporal Uncertainty Scores

As we describe earlier, most real-life scenarios for deployment involve consecutive inputs to the model that are likely to be correlated over time. Thus, there is value in evaluating samples in the context of other samples instead of viewing them independently. In our framework, we consider high-dimensional inputs to the model as a sequence (Fig. 1). We use the representational capabilities of modern encoders (Rabanser et al., 2019) and examine them over time instead of directly modeling the incoming high-dimensional, multivariate data stream or treating each sample independently, i.e., we detect when input samples go out of distribution by comparing their reduced-dimension scores (s_{spatial}) with those of other inputs over time. We note that when these scores are scalar (as in our case), the problem reduces to change-point detection for a one-dimensional sequence.

We put forth an unsupervised, non-parametric, sliding-window based algorithm that builds on the work done by Kifer et al. (2004) to generate our temporal uncertainty scores s_{temporal}. We consider a sequence of r samples in the input data stream denoted by $\{x_1, x_2, ..., x_t, ..., x_r\}$. We obtain the scalar spatial scores for the r samples, $\{s_{\text{spatial}}(x_1), s_{\text{spatial}}(x_2), ..., s_{\text{spatial}}(x_t), ..., s_{\text{spatial}}(x_r)\}$ from their representations $\{h(x_1), h(x_2), ..., h(x_t), ..., h(x_r)\}$ as explained in Section 3.2.1.

Given two window sizes w_A and w_B ($w_A, w_B < r$), we define a reference window R_W and a sliding window $(S_W)_t$ to be

$$R_W = [s_{\text{spatial}}(x_1), ..., s_{\text{spatial}}(x_{w_A})]$$

$$(S_W)_t = [s_{\text{spatial}}(x_{(w_A + t)}), ..., s_{\text{spatial}}(x_{(w_A + t + w_B)})]$$

where $t \in \{1, ..., (r - w_A - w_B)\}$ as visualized in Fig. 1. On choosing an appropriate statistic or distance measure $\mathcal{F}$ (e.g. probability odds ratio, Kolmogorov-Smirnov, Kullback-Leibler, Wilcoxon, Mann-Whitney, etc.), we compute the temporal uncertainty score between reference window R_W and the tth sliding window $(S_W)_t$ for the tth sample

$$s_{\text{temporal}}(x_t) = \mathcal{F}(R_W, (S_W)_t)$$

The null hypothesis $\mathcal{H}_0$ for our setting here is that the two windows (reference and sliding) are similar and come from the same distribution. We hope to reject $\mathcal{H}_0$ in favour of the alternate hypothesis $\mathcal{H}_1$ at a significance level of 0.05 whenever the inputs go out-of-distribution. This approach (as in (Kifer et al., 2004)) assumes the data points are generated sequentially by some underlying probability distribution, but otherwise makes no assumptions on the nature of the generating distribution nor does it assume the samples are identically distributed.

4 Experimental Setup

We perform experiments across multiple domains (vision, audio and clinical EEG signals) to evaluate our approach. We choose the task of image classification for the visual domain, spoken word classification from audio clips for the audio domain and the challenging, real-world clinical task of

seizure detection from EEG signals framed as an EEG clip binary classification task for the healthcare domain. Our code will be made publicly available after publication.

We use a variety of datasets and models in our experiments. For the audio domain, we use the Free Spoken Digits Dataset (FSDD) (Jackson et al., 2017) and Google Speech Commands (GSC) dataset (Warden, 2018). We train an M5 encoder (Dai et al., 2016) with raw audio InD data from spoken digits 0-9 from GSC and evaluate it on unseen InD samples from GSC 0-9 and FSDD 0-9 along with OOD samples from the non-digit spoken words from GSC. For our seizure tasks, we use EEG data from Stanford Health Care (SHC) (InD: patient population 20-60 years), Lucile Packard Children's Hospital (LPCH) (OOD: patient population <20 years) and the public Temple University Hospital seizure corpus (TUH) (Shah et al., 2018; Obeid & Picone, 2016) (OOD: different patient demographics) along with a Dense-Inception encoder as described by Saab et al. (2020). For our vision tasks, we use MNIST as InD and MNIST-like data as OOD and CIFAR10 as InD, SVHN as OOD and CIFAR100 as a semantic OOD evaluation set up. We describe more details about the datasets, model training and temporal framework settings in Appendix B.1 and B.3.

We compare against six baselines: Maximum Softmax Probability (MSP) (Hendrycks & Gimpel, 2017), Predictive Entropy (Ren et al., 2019), KL Divergence with Uniform Distribution (Hendrycks et al., 2019a), ODIN (Liang et al., 2018), Vanilla Mahalanobis distance (Lee et al., 2018) and Test-time Dropout (Gal & Ghahramani, 2016). We provide baseline details in Appendix B.2

4.1 Evaluation

For all experiments, we evaluate our approach only using samples unseen during training. Our test sets include unseen samples drawn from the same distribution as the training data along with samples from unseen classes and significantly different datasets. We also follow the setting recommended by Ahmed & Courville (2020) to avoid dataset bias by holding out few classes from a dataset during training and using the held-out classes as new semantic content for evaluation. Since OOD detection can be viewed as a binary classification problem, we evaluate its performance using standard binary classification metrics: area under the receiver operating characteristic (AUROC↑) , area under the precision-recall curve (AUPR↑) and the false positive rate at 80% true positive rate (FPR80↓) as commonly prescribed (Ren et al., 2019).

In practice, it is important to note that what constitutes InD or OOD is driven by context, i.e., it depends on the target task, the encoder capacity, what our model needs to be robust to and what it should detect as outliers. Standard OOD benchmarks treat different datasets to be OOD resulting in models with narrow capabilities susceptible to dataset bias (Ahmed & Courville, 2020; Yang et al., 2021). Moreover, they fail to capture any semantic overlap between different datasets. For example, if a model is trained to classify cats vs dogs and is given a different breed of a dog (high semantic overlap) from a different dataset at test time, it should still be considered InD whereas a muffin would be OOD since it has no semantic overlap with the training data. In this work, we evaluate OOD detection performance with and without semantic overlap and show that TIME-LAPSE is driven more by semantic content than other baselines.

We further propose a sequential evaluation framework to evaluate OOD detection under realistic conditions of data drifts. We use the following setup: we evaluate performance metrics such as time to detect the change point, $\%$ error across trials, detection accuracy and false positive rate over $N = 1000$ trials. For each trial, we generate a stream of test data from our test sets with change points inserted according to a drift rate parameter r and fraction of OOD points given by the parameter λ. We vary r and λ to emulate various drift conditions including rapidly changing streams and slowly shifting streams. We also generate control streams by setting $r = 0$ to estimate the false positive rate.

5 Results

5.1 Performance on Audio, Clinical EEG and Image datasets

We achieve the state-of-the-art (SOTA) with TIME-LAPSE scores on OOD detection in both our audio and EEG-based tasks of spoken digits classification and seizure detection compared to our baselines (Table 1) across all metrics: AUROC, AUPR and FPR80. The AUROC values for TIME-LAPSE scores are 75% and 77.8%, 10 points and 13 points higher than the strongest baseline, Test-time

Table 1: OOD performance: AUROC↑ / AUPR↑ / FPR80↓. The best scores are indicated in bold.

	Audio	EEG Data	Vision		
Task	Speech Classification	Seizure Detection	Image Classification / Digit Classification		
OOD Sets	Other spoken words	Other institutions	SVHN	CIFAR100	x-MNIST
MSP	0.623 / 0.523 / 0.513	0.354 / 0.419 / 0.757	0.763 / 0.789 / 0.398	0.856 / 0.986 / 0.226	0.907 / 0.982 / 0.133
Predictive Entropy	0.612 / 0.511 / 0.516	0.391 / 0.497 / 0.738	0.765 / 0.793 / 0.396	0.858 / 0.987 / 0.224	0.910 / 0.983 / 0.132
KL_U	0.555 / 0.475 / 0.571	0.387 / 0.478 / 0.712	0.797 / 0.822 / 0.353	0.872 / 0.988 / 0.199	0.909 / 0.983 / 0.144
ODIN	0.465 / 0.454 / 0.766	0.320 / 0.397 / 0.788	0.737 / 0.783 / 0.477	0.844 / 0.986 / 0.246	0.906 / 0.982 / 0.130
Vanilla Mahalanobis	0.695 / 0.645 / 0.513	0.635 / 0.654 / 0.521	0.782 / 0.815 / 0.389	0.695 / 0.968 / 0.524	0.919 / 09.83 / 0.119
Test-Time Dropout	0.650 / 0.617 / 0.526	0.652 / 0.617 / 0.585	0.679 / 0.716 / 0.546	**0.923** / 0.986 / **0.050**	**0.976** / 0.986 / **0.016**
TIME-LAPSE (ours)	**0.750 / 0.714 / 0.438**	**0.778 / 0.703 / 0.331**	**0.838 / 0.852 / 0.285**	0.875 / **0.989** / 0.288	0.975 / **0.996** / 0.106

Table 2: Effect of Semantic Affinity on OOD performance for Audio tasks (AUROC↑ / AUPR↑)

Audio	MSP	Predictive Entropy	KL_U	ODIN	Vanilla Mahalanobis	Test-Time Dropout	TIME-LAPSE (ours)
Semantic OOD	0.623 / 0.523	0.612 / 0.511	0.555 / 0.475	0.465 / 0.454	0.695 / 0.645	0.650 / 0.617	**0.750 / 0.714**
Dataset Driven	0.823 / 0.894	**0.826 / 0.897**	0.815 / 0.894	0.679 / 0.811	0.493 / 0.706	0.821 / 0.892	0.506 / 0.736

Dropout, for both tasks respectively. In our MNIST experiments, we see TIME-LAPSE performs comparably with Test-Time Dropout. It has the best AUPR and comparable AUROC and FPR80 values. Over CIFAR10 experiments with the SVHN dataset (no semantic overlap with CIFAR10) as OOD, TIME-LAPSE outperforms other baselines significantly. When CIFAR100 classes are added (considered OOD as in standard benchmarks) to the evaluation, TIME-LAPSE shows the best AUPR score and the second best AUROC and FPR80. We examine the effect of semantic overlap present in CIFAR100 classes with the InD classes and show that TIME-LAPSE is more driven by semantic content in Section 5.2. We study the specific contribution of our temporal scores in Section 5.3 using our sequential evaluation framework.

5.2 Affinity for Semantic Content and Robustness to Dataset Statistics

We study the effect of semantic overlap in OOD classes and characterize how TIME-LAPSE compares against other methods in its robustness to dataset statistics over multiple domains.

Audio In our audio experiments, we train our encoder only on spoken digits 0-9 (forming the InD classes) from the Google Speech Commands (GSC) dataset. The other 25 classes (Yes, No, Right, Left, Bird, etc) in GSC have never been encountered by the model and have drastic semantic shifts in their content when compared to the InD classes 0-9 and are considered OOD. During evaluation, we show the model unseen samples from each of the classes in GSC along with samples from the Free Spoken Digits Dataset (FSDD) which contains classes 0-9 (same semantic content as InD data but different dataset) and measure OOD detection performance (Table 2). As an ablation, we examine the performance when we flip the FSDD classes to be OOD data in our evaluation. We see that AUROC, AUPR and FPR80 values for most baselines (notably Test-time Dropout and Predictive Entropy) increase significantly, whereas they drop for TIME-LAPSE and vanilla Mahalanobis (Table 2). This shows that our strongest baseline methods have narrow capabilities that assume only data belonging to InD classes *and* having the *same dataset statistics* as the training data will be considered as InD. All other data will be rejected as OOD, leading to very high FPRs in practice. In contrast, TIME-LAPSE is able to identify true InD samples despite differing dataset statistics, i.e. TIME-LAPSE detects OOD samples based on semantic content than dataset statistics when compared to the other methods.

Vision (CIFAR10) For our CIFAR10 experiments, we train our encoder only on 7 classes of CIFAR10, leaving out classes Airplane, Bird and Dog. SVHN data consists of images of numbers and has no semantic overlap with any of the 7 InD classes (Automobile, Cat, Deer, Frog, Horse, Ship, Truck). To study the effect of semantic overlap, we consider the 100 classes from CIFAR100 along with the 3 left out classes from CIFAR10. We compare the performance metrics when naively considering all of the above as OOD ignoring any semantic overlap. We see that baselines perform better while TIME-LAPSE drops (Table 1). We analyse class-wise performance and we see that TIME-LAPSE's performance strongly correlates to semantic content compared to other methods. For example, consider Class Pickup-Truck and Streetcar from CIFAR100. While they are disjoint from the InD

Table 3: Pickup Truck and Street Car share high semantic overlap with InD classes Truck and Automobile. Naively considering them OOD should result in AUROC↓ / AUPR↓. Airplane and Fish do not show any semantic overlap with InD classes and should result in AUROC↑ / AUPR↑

CIFAR	MSP	Predictive Entropy	KL_U	ODIN	Vanilla Mahalanobis	Test-Time Dropout	TIME-LAPSE (ours)
Pickup Truck	0.793 / 0.292	0.791 / 0.276	0.698 / 0.181	0.815 / 0.182	0.691 / 0.207	0.942 / 0.506	0.791 / 0.276
Street Car	0.767 / 0.255	0.768 / 0.256	0.766 / 0.239	0.785 / 0.255	0.656 / 0.181	0.838 / 0.337	0.757 / 0.256
Airplane	0.891 / 0.488	0.894 / 0.519	0.913 / 0.571	0.893 / 0.549	0.658 / 0.192	0.955 / 0.547	0.902 / 0.575
Aquarium Fish	0.893 / 0.492	0.897 / 0.529	0.913 / 0.589	0.888 / 0.526	0.705 / 0.209	0.955 / 0.542	0.911 / 0.577

classes Automobile and Truck in CIFAR10, they shares a high degree of semantic content and should be considered InD. When models encounter samples from classes that have high semantic overlap with InD, they need to consider them InD and not OOD. Instead, if such samples are considered OOD, only methods that have narrow model capacities will show high detection performance (Table 3) indicating their dependence on data statistics instead of semantic content.

Vision (MNIST) We train two different encoders, one trained on all digits 0-9 while the other was trained on $\{0, 1, 4, 6, 7, 8, 9\}$ leaving digits 2, 3 and 5. Amongst our evaluations, only certain classes from the e-MNIST data shared semantic overlap – letters 'o','l','i','z','y','s' and 'q' share structural similarity with classes 0, 1, 2, 4, 5 and 9 respectively. We show results for both encoders with and without the above classes in the OOD evaluations (Table 4, Appendix).

5.3 Sequential Evaluation Results

To evaluate the performance under our sequential framework, we define the detector to have erred if it wrongly concludes an incoming test sample $\mathbf{x}_t$ to be OOD when it is actually InD or vice-versa. If the framework detects the change point within its fixed window size, we do not consider it as an error. We generate data streams of length 10,000 samples changing the distribution every k samples in each trial, determined by the rate r and fraction of change-points λ, emulating a data stream. The distribution changes are simulated by first randomly choosing InD or OOD (Bernoulli random variable with probability p) and then randomly drawing the k samples from the chosen distribution. We calculate the TIME-LAPSE scores using the two-sided Mann-Whitney test (Mann & Whitney, 1947) at a significance level of 0.05 for each sample in the stream and run it through our sequential detector with window sizes $w_A = w_B = 25$ for vision experiments and $w_A = w_B = 50$ for audio and EEG experiments. We show example plots of datastreams for each domain in the Appendix, Figs. 2, 3, 4, 5, 6. We perform a 2-cluster KMeans on the generated Mann-Whitney scores to assign InD and OOD predictions. We then calculate the cumulative error by adding the errors of all samples to calculate detection error. We repeat this experiment N times ($N = 1000$ here) to estimate the detector's error distribution (EEG distribution given in the Fig. 2, Appendix). Over 73% of the cumulative errors in our EEG task (over 85% for our audio task) lie within 10% and over 93% of the errors in our EEG task (over 97% for our audio task) lie within 20% indicating our detector's high performance in comparison to other spatial detectors. For our vision tasks, the error distribution is much tighter and we are able to get near 99% detection accuracy for over 95% of the time.

5.4 Ablations: Effect of Encoder capacity

The proposed TIME-LAPSE framework depends on the encoder being well-trained and having good representational capability. To study the effects of encoder capacity, we consider our audio task. We compare two encoders, trained on the same InD data and evaluated with the same OOD data. The lower capacity encoder, Kymatio, is a two-layer network with a log-scattering layer that extracts scattering coefficients from the inputs as lower-dimensional latent space representations followed by a log-softmax layer for classification with 0.75 classification accuracy for InD test samples. The higher capacity encoder is a CNN-based M5 model with 0.85 classification accuracy for InD test samples. As shown in Fig. 7 (Appendix), Kymatio is not able to separate out any of the classes based on the TIME-LAPSE scores whereas M5 is able to correctly separate out semantic InD samples from OOD samples well. Thus, the success of TIME-LAPSE is highly dependent on the representational capacity of the model.

To further investigate the effects of the lower capacity model, we train Kymatio on the much smaller FSDD dataset but with the same InD semantic classes 0-9. The model reaches 0.97 classification accuracy, showing that it is able to express high-dimensional audio signals well. However, on comparing its TIME-LAPSE scores with those of Kymatio trained on GSC 0-9 (Fig 8, Appendix), we see that the former separates out FSDD InD from GSC InD and GSC OOD indicating that it is overfitting to the dataset statistics whereas the latter is not able to separate out any of them. Thus, TIME-LAPSE scores can help in determining both OOD detection capability of an encoder as well as generalization capability over OOD samples.

6 DISCUSSION

In this work, we develop a unique spatio-temporal framework to identify when incoming inputs fall out-of-distribution with respect to trained encoders and subsequently enable models to abstain from prediction on such inputs. Our theoretical formulation encompasses all representation-based OOD detection methods and provides an opportunity to leverage temporal signals in combination. We show that longitudinal information across time can be leveraged for the challenging task of identifying distributional shifts. Through our experiments, we verify that using temporal uncertainty scores and spatial uncertainty scores is complementary and combining them gives additional information than when used separately. We proposed a simple sliding-window hypothesis-testing based approach to model temporal uncertainty of the reduced dimension representations and their spatial uncertainty scores. But TIME-LAPSE can accept any type of spatial scoring and temporal scoring. We believe that there is much scope for developing more sophisticated methods that leverage the temporal aspect and leave it open for future work.

Furthermore, alongside works by Ahmed & Courville (2020) and Yang et al. (2021), we motivate the necessity to rethink standard OOD image benchmarks that show semantic overlaps through our experiments. We point out that the task of identifying OOD inputs and subsequent selective prediction is highly dependent on context. For a classifier designed to distinguish between cats and dogs, a tree is out-of-distribution whereas all of them would be in-distribution for an object recognition system in a self-driving car. We believe that showing generalizability of OOD detection methods across domains is key to ensure robust performance in the wild. It is not sufficient to show better performance over just image OOD benchmarks. We show that TIME-LAPSE generalizes across 3 widely different domains, vision, audio and healthcare EEG, through our extensive experiments and analyses. We also propose a sequential evaluation framework to view OOD detection from the lens of data and model drift detection. To enable safe deployment and periodic monitoring of trained models, we believe that such a sequence-based evaluation scheme is essential. We hope this evaluation scheme will be adopted for routinely characterizing OOD detection performance in the wild.

Through our ablation study, we see that our method depends strongly on the capacity of our trained encoders, i.e., how well the encoders can represent the InD data. In our experiments, we train encoders with labelled InD data under the supervised learning paradigm. Our framework will support and only benefit if more robust representations are learnt through self-supervised learning or semi-supervised learning. We also note that both the audio and the seizure analyses tasks generate much higher-dimensional inputs when compared to image classification tasks. We hypothesize that model performance is more strongly dependent on semantic content for such high-dimensional inputs when compared to dataset statistics and thus, gains from TIME-LAPSE are more obvious.

In conclusion, we provide a framework that utilizes both latent-space and time to determine predictive uncertainty for deep learning models. Our method is end-to-end trainable and network agnostic: it (a) does not require exposure to outliers (though can easily be extended to accommodate them if present), (b) does not require modifications to existing network architectures, and (c) does not require computationally heavy dropout techniques or ensemble passes. It can thus can be easily adopted to high-risk settings such as clinical workflows and autonomous driving.

REFERENCES

Faruk Ahmed and Aaron Courville. Detecting semantic anomalies. *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 3154–3162, 2020.

Samaneh Aminikhanghahi and Diane J. Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017.

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.

Mathieu Andreux, Tomás Angles, Georgios Exarchakis, Roberto Leonarduzzi, Gaspar Rochette, Louis Thiry, John Zarka, Stéphane Mallat, Joakim Andén, Eugene Belilovsky, Joan Bruna, Vincent Lostanlen, Matthew J. Hirn, Edouard Oyallon, Sixin Zhang, Carmine Cella, and Michael Eickenberg. Kymatio: Scattering transforms in python. *arXiv Preprint arXiv:1812.11214*, 2018.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *International Conference on Machine Learning (ICML)*, 2015.

Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. *arXiv Preprint arXiv:1604.07316*, 2016.

Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. Lof: Identifying density-based local outliers. *Association for Computing Machinery*, 29(2):93–104, 2000.

Wenhu Chen, Yilin Shen, Hongxia Jin, and William Wang. A variational dirichlet framework for out-of-distribution detection. *arXiv preprint arXiv:1810.01392*, 2019.

Hyunsun Choi, Eric Jang, and Alexander A. Alemi. Waic, but why? generative ensembles for robust anomaly detection, 2019.

Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre van Schaik. Emnist: an extension of mnist to handwritten letters. *arXiv preprint arXiv:1702.05373*, 2017.

Wei Dai, Chia Dai, Shuhui Qu, Juncheng Li, and Samarjit Das. Very deep convolutional neural networks for raw waveforms. *arXiv preprint arXiv:1610.00087*, 2016.

Lucas Deecke, Robert Vandermeulen, Lukas Ruff, Stephan Mandt, and Marius Kloft. Image anomaly detection with generative adversarial networks. *Machine Learning and Knowledge Discovery in Databases*, pp. 3–17, 2019.

Terrance DeVries and Graham W. Taylor. Learning confidence for out-of-distribution detection in neural networks. *arXiv Preprint arXiv:1802.04865*, 2018.

Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 2017.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *International Conference on Machine Learning (ICML)*, 2016.

João Gama, Indr Žliobait, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 2014.

Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)*, 2015.

Will Grathwohl, Kuan-Chieh Wang, Joern-Henrik Jacobsen, David Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. *International Conference on Learning Representations (ICLR)*, 2020.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. *International Conference on Machine Learning (ICML)*, 2017.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *International Conference on Learning Representations (ICLR)*, 2017.

Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep anomaly detection with outlier exposure. *International Conference on Learning Representations (ICLR)*, 2019a.

Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019b.

Geoffrey Hinton. Deep learning: A technology with the potential to transform health care. *JAMA*, 320(11):1101–1102, 2018.

Zohar Jackson, César Souza, Jason Flaks, Yuxin Pan, Hereman Nicolas, and Adhish Thite. Free spoken digits dataset. *https://github.com/Jakobovski/free-spoken-digit-dataset*, 2017.

William P. Jones and George W. Furnas. Pictures of relevance: A geometric analysis of similarity measures. *Journal of the American Society for Information Science*, 1987.

Daniel Kifer, Shai Ben-David, and Johannes Gehrke. Detecting change in data streams. *International Conference on Very Large Data Bases (VLDB)*, 2004.

Benjamin Kompa, Jasper Snoek, and Andrew L. Beam. Second opinion needed: Communicating uncertainty in medical machine learning. *npj Digital Medicine*, 2021.

Vladik Kreinovich and Olga Kosheleva. Deep learning (partly) demystified. *International Conference on Intelligent Systems, Metaheuristics and Swarm Intelligence*, 2020.

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Tech Report*, 2009.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2012.

Yann Lecun, Leon Bottou, Yoshua Bengio, and Parick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.

Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. *International Conference on Learning Representations (ICLR)*, 2018.

Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. *International Conference on Data Mining (ICDM)*, 2008.

Weitang Liu, Xiaoyun Wang, John D. Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

J. MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematics, Statistics and Probability*, 1967.

Andrey Malinin and Mark Gales. Predictive uncertainty estimation via prior networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

H.B. Mann and D.R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 1947.

Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

Sina Mohseni, Mandar Pitale, JBS Yadawa, and Zhangyang Wang. Self-supervised learning for generalizable out-of-distribution detection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.

Jose G. Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V. Chawla, and Francisco Herrera. A unifying view on dataset shift in classification. *Pattern Recognition*, 45(1):521–530, 2012.

Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don't know? *International Conference on Learning Representations (ICLR)*, 2019.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Ng. Reading digits in natural images with unsupervised feature learning. *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.

Anh Mai Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

Iyad Obeid and Joseph Picone. The temple university hospital eeg data corpus. *Frontiers in neuroscience*, 2016.

Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. Ocgan: One-class novelty detection using gans with constrained latent representations. *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.

Stanislav Pidhorskyi, Ranya Almohsen, Donald A Adjeroh, and Gianfranco Doretto. Generative probabilistic novelty detection with adversarial autoencoders. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Stephan Rabanser, Stephan Günnemann, and Zachary C. Lipton. Failing loudly: An empirical study of methods for detecting dataset shift. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Jie Ren, Peter J. Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Lukas Ruff, Robert A. Vandermeulen, Nico Görnitz, Alexander Binder, Emmanuel Müller, Klaus-Robert Müller, and Marius Kloft. Deep semi-supervised anomaly detection. *International Conference on Learning Representations (ICLR)*, 2020.

Khaled Saab, Jared Dunnmon, Christopher Ré, Daniel Rubin, and Christopher Lee-Messer. Weak supervision as an efficient approach for automated seizure detection in electroencephalography. *npj Digital Medicine*, 2020.

Thomas Schlegl, Philipp Seeböck, Sebastian M. Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. *Information Processing in Medical Imaging*, pp. 146–157, 2017.

Bernhard Schölkopf, Robert Williamson, Alex Smola, John Shawe-Taylor, and John Platt. Support vector method for novelty detection. *Advances in Neural Information Processing Systems (NeurIPS)*, 1999.

Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 2020.

Vikash Sehwag, Mung Chiang, and Prateek Mittal. Ssd: A unified framework for self-supervised outlier detection. *International Conference on Learning Representations (ICLR)*, 2021.

Joan Serrà, David Álvarez, Vicenç Gómez, Olga Slizovskaia, José F. Núñez, and Jordi Luque. Input complexity and out-of-distribution detection with likelihood-based generative models. *International Conference on Learning Representations (ICLR)*, 2020.

Vinit Shah, Eva von Weltin, Silvia Lopez, James Riley McHugh, Lillian Veloso, Meysam Golmohammadi, Iyad Obeid, and Joseph Picone. The temple university hospital seizure detection corpus. *Frontiers in neuroinformatics*, 2018.

Gabi Shalev, Yossi Adi, and Joseph Keshet. Out-of-distribution detection using multiple semantic label representations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Clanuwat Tarin, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*, 2018.

Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alexander Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv Preprint arXiv:1609.03499*, 2016.

Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *https://arxiv.org/pdf/1804.03209.pdf*, 2018.

Jim Winkens, Rudy Bunel, Abhijit Guha Roy, Robert Stanforth, Vivek Natarajan, Joseph R. Ledsam, Patricia MacWilliams, Pushmeet Kohli, Alan Karthikesalingam, Simon Kohl, Taylan Cemgil, S. M. Ali Eslami, and Olaf Ronneberger. Contrastive training for improved out-of-distribution detection. *arXiv Preprint arXiv:2007.05566*, 2020.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Zhisheng Xiao, Qing Yan, and Yali Amit. Likelihood regret: An out-of-distribution detection score for variational auto-encoder. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Min Xu, Ling-Yu Duan, Jianfei Cai, Liang-Tien Chia, Changsheng Xu, and Qi Tian. Hmm-based audio keyword generation. *Advances in Multimedia Information Processing*, 2005.

Adam Yala, Constance Lehman, Tal Schuster, Tally Portnoi, and Regina Barzilay. A deep learning mammography-based model for improved breast cancer risk prediction. *Radiology*, 2019.

Jingkang Yang, Haoqi Wang, Litong Feng, Xiaopeng Yan, Huabin Zheng, Wayne Zhang, and Ziwei Liu. Semantically coherent out-of-distribution detection. *International Conference on Computer Vision (ICCV)*, 2021.

Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. *International Conference on Learning Representations (ICLR)*, 2018.

APPENDIX

A SUPPLEMENTARY TABLES

MNIST Semantic Overlap Correction We train two different encoders, one trained on all digits 0-9 while the other was trained on $\{0, 1, 4, 6, 7, 8, 9\}$ leaving digits 2, 3 and 5. Amongst our evaluations, only certain classes from the e-MNIST data shared semantic overlap – letters ‘o’,‘l’,‘i’,‘z’,‘y’,‘s’ and ‘q’ share structural similarity with classes 0, 1, 2, 4, 5 and 9 respectively. We show results for both encoders with and without the above classes in the OOD evaluations (Table 4, Appendix).

Table 4: MNIST models. 0to9 denotes encoder trained on classes 0-9 digits. M235 indicates encoder trained on classes {0 to 9} - {2,3,5}, thus trained on {0,1,4,6,7,8,9}. Semantic indicates classes from e-MNIST that have ambiguous semantic overlap with InD classes such as {l,i,o,..} are removed from OOD evaluation

MNIST	MSP	Predictive Entropy	KL_U	ODIN	Vanilla Mahalanobis	Test-Time Dropout	TIME-LAPSE (ours)
0to9	0.895 / 0.968	0.899 / 0.971	0.903 / 0.972	0.896 / 0.969	0.903 / 0.973	0.963 / 0.973	0.961 / 0.990
0to9 Semantic	0.921 / 0.970	0.926 / 0.973	0.928 / 0.975	0.922 / 0.971	0.937 / 0.976	0.976 / 0.970	**0.986 / 0.995**
M235	0.907 / 0.982	0.910 / 0.983	0.909 / 0.983	0.906 / 0.982	0.919 / 0.985	0.976 / 0.986	0.975 / 0.996
M235 Semantic	0.926 / 0.983	0.929 / 0.984	0.925 / 0.983	0.925 / 0.983	0.941 / 0.986	0.984 / 0.986	**0.990 / 998**

B Detailed Experimental Setup

B.1 Datasets

Audio For the task of spoken word classification of audio recordings, we use the Free Spoken Digits Dataset (FSDD) (Jackson et al., 2017), consisting of .wav recordings of English digits 0-9, and the Google Speech Commands (GSC) dataset (Warden, 2018) that contain .wav recordings of 35 English keywords including the digits 0-9 and non-digits such as "Yes", "No", "Stop", etc. The two datasets provide a unique setup where digits 0-9 occur in different datasets but belong to the same semantic classes.

EEG Signals For our seizure analyses tasks, we use high-dimensional EEG data from different institutions and age groups: (a) Stanford Healthcare (SHC) with patients in the age group $\sim$20-60 years, (b) Lucile Packard Children's Hospital (LPCH) with patients under 18 years and (c) publicly available Temple University Hospital seizure corpus (TUH) (Shah et al., 2018; Obeid & Picone, 2016). 12 second or 60 second EEG clips with 19 channels, each sampled at 200Hz, are used for seizure detection.

Vision We use CIFAR-10 (InD), CIFAR-100 (Krizhevsky & Hinton, 2009) and SVHN (Netzer et al., 2011) datasets along with MNIST (Lecun et al., 1998) (InD), FashionMNIST (Xiao et al., 2017), eMNIST (Cohen et al., 2017) and kMNIST (Tarin et al., 2018) (collectively referred to as x-MNIST) datasets for image classification as is typically reported in literature (Hendrycks & Gimpel, 2017; Lee et al., 2018; Ren et al., 2019).

B.2 Baselines

We use the following baselines to compare against our framework. We don't compare with methods that require exposure to outliers to match our framework settings.

MSP: Hendrycks & Gimpel (2017) use the maximum softmax probability (MSP) $p(\hat{y}|x) = \max_k p(y = k|x)$ as the uncertainty score. OOD inputs tend to have lower MSP than InD data.

Predictive Entropy: Ren et al. (2019) show that using high entropy of the predicted class distribution $-\sum_k p(y = k|x) \log p(y = k|x)$ as an indicator for OOD inputs is a strong baseline.

KL-divergence with Uniform distribution: Hendrycks et al. (2019a) use the KL divergence of the softmax predictions to the uniform distribution U, $\mathrm{KL}(U||p(y|x))$ to identify OOD inputs.

ODIN: Liang et al. (2018) use temperature-scaling and input perturbations to increase the gap between OOD and InD data using MSP. We fix $T = 1000, \epsilon = 1.4\mathrm{e}^{-3}$, following their most common setting instead of tuning the parameters with outlier exposure.

Vanilla Mahalanobis: Lee et al. (2018) use the Mahalanobis distance from the nearest class-conditional Gaussian with shared covariance as the uncertainty score. This approach directly fits in with our spatial scoring method (Section 3.2.1) though we use class-wise covariance matrices.

Test-time Dropout: Gal & Ghahramani (2016) show that Monte Carlo (test-time) dropout to estimate the prediction distribution give good uncertainty estimates. Note that it is computationally intensive requiring k forward passes of the classifier, with k-fold increase in runtime. We use $k = 10$.

B.3 MODEL TRAINING & HYPERPARAMETERS

We carry out our vision experiments using CNN-based classifiers, i.e. the LeNet architecture (Lecun et al., 1998) when working with x-MNIST data and the ResNet18 architecture (He et al., 2016) when working with CIFAR10/100/SVHN data. For our audio experiments, we use the M5 classifier (Dai et al., 2016) as our encoder. We also train a simple two-layer network – a static, normalized, log-scattering layer that extracts scattering coefficients from audio signals (Andreux et al., 2018) followed by a log-softmax layer that generates output probabilities to compare encoder capabilities (Section 5.4). Both models are trained only on classes 0-9 with the rest of the classes as OOD inputs. In all cases, we learn embeddings directly from raw, one second audio clips (resampled to 8kHz) without making use of any spectrogram features like MFCC (Xu et al., 2005). For our seizure detection task, we use the Dense-Inception based models (Saab et al., 2020) as encoders. The models are trained on 19 channel, 12 second or 60 second raw EEG clips from InD datasets, resampled to 200Hz. We train all encoders on InD training data using standard weight initializations, SGD/Adam optimizer and ReLU non-linearities. We train and tune hyperparameters for our encoders using only InD data samples. We don't expose any outlier data to our encoders. We extract activations from a fully connected (FC) hidden layer before the logits layer to form latent space representations. We normalize our distance scores and similarity scores before forming the spatial scores in cases where one of them is tightly bound while the other is not. For our temporal uncertainty scoring, we set window sizes w_A, $w_B = 25$ for vision tasks and w_A, $w_B = 50$ for audio and EEG tasks. We choose the non-parametric, two-sided Mann-Whitney test (Mann & Whitney, 1947) at a significance level of 0.05 as our measure $\mathcal{F}$. We perform two-centered kMeans clustering (MacQueen, 1967) on the temporal score sequence to identify uncertain model predictions and reject corresponding test samples as OOD.

C SEQUENTIAL EVALUATION: VISUALIZATIONS

Some examples of data streams generated from sampling EEG, audio and vision datasets are shown below. The X axis indicates the time sequence. And the Y axis gives the Mann Whitney Scores. Using a two cluster KMeans algorithm, change points were detected and individial samples were accepted or rejected as InD.

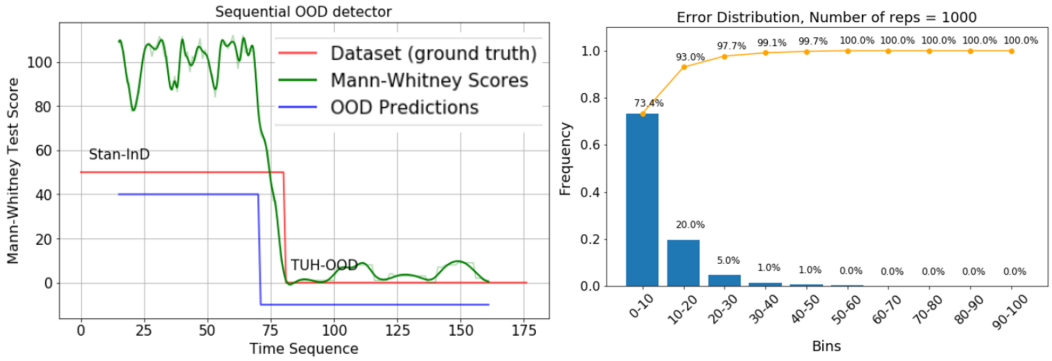

Figure 2: (left) Snapshot of the TIME-LAPSE temporal scores. Within a window of 15 samples, it is able to detect the change point for the seizure detection task. (right) % Error distribution over 1000 trials.

D SUPPLEMENTARY FIGURES: EFFECT OF ENCODER CAPACITY

We visualize the distributions of TIME-LAPSE spatial scores generated by two encoder architectures (Kymatio and M5) with varying modelling capacity. We see from Fig. 7 that more modelling capacity an encoder has, the better it is able to separate InD and OOD samples in its latent space. Note that in Fig. 7, the left plot shows the distribution of scores for InD samples from different datasets in blue and orange and that of all the OOD samples together in green. In the right plot, however, blue

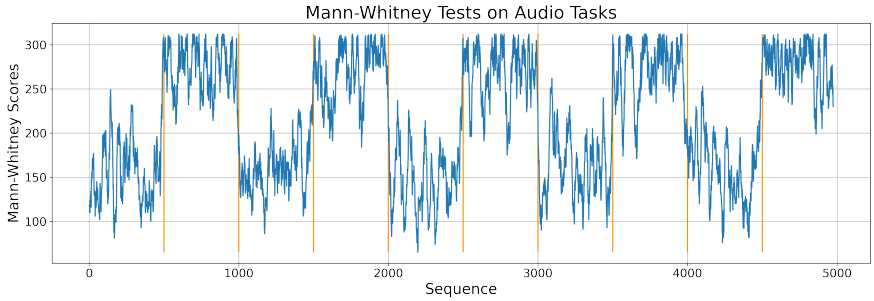

Figure 3: Sequential Evaluation: Data stream with distribution shifts occurring every 500 steps. The orange lines indicate true change points. Blue plot represents the Mann-Whitney scores generated from the spatial scores of audio clips.

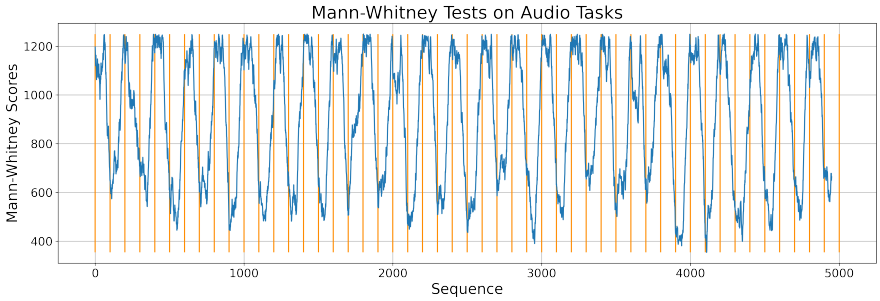

Figure 4: Sequential Evaluation: Data stream with distribution shifts occurring every 100 steps. The orange lines indicate true change points. Blue plot represents the Mann-Whitney scores generated from the spatial scores of audio clips.

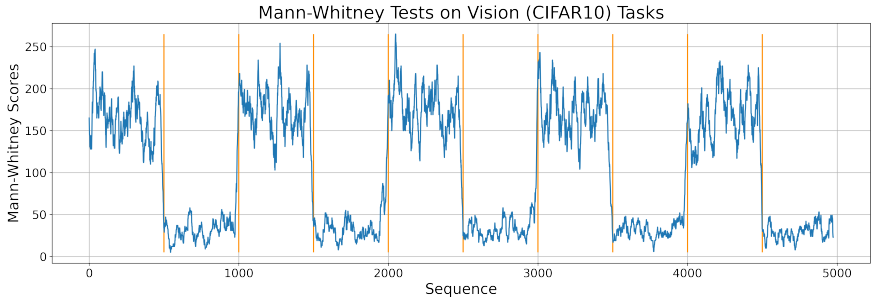

Figure 5: Sequential Evaluation: Data stream with distribution shifts occurring every 500 steps. The orange lines indicate true change points. Blue plot represents the Mann-Whitney scores generated from the spatial scores of CIFAR + OOD images.

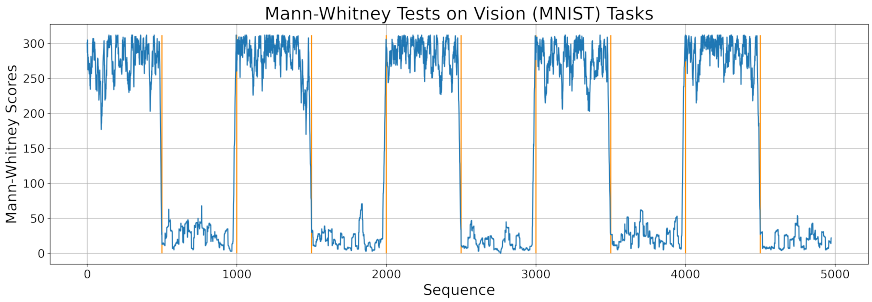

Figure 6: Sequential Evaluation: Data stream with distribution shifts occurring every 500 steps. The orange lines indicate true change points. Blue plot represents the Mann-Whitney scores generated from the spatial scores of MNIST + OOD images.

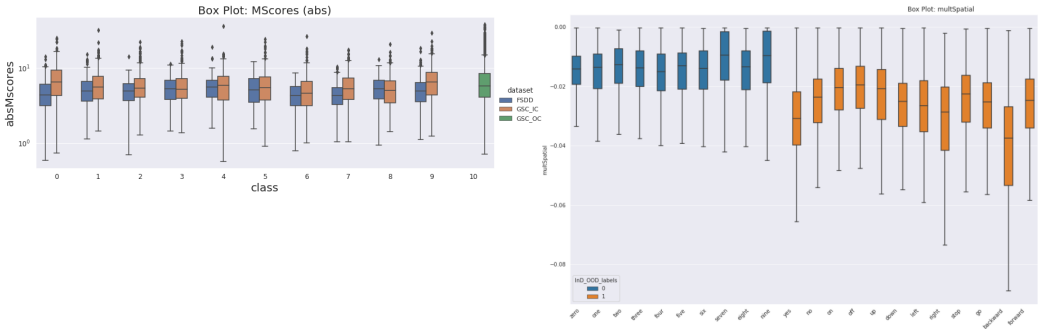

Figure 7: Distribution of TIME-LAPSE scores across semantic InD (GSC 0-9, FSDD 0-9) and OOD classes with (left) Kymatio encoder, (right) M5 encoder, class-wise OOD scores are shown. GSC 0-9 and FSDD 0-9 are merged here and only part of the full list of OOD classes shown. All GSC OOD classes are merged together as green in the left figure.

indicate the merged InD samples from both datasets while orange represents class-wise distributions across OOD samples.

In Fig. 8, we keep the architecture of the encoder fixed, i.e., Kymatio, and vary how well it is trained. We can two encoders of the same architecture, train one on GSC 0-9 classes and another on FSDD 0-9 classes. Both the encoders are thus trained on the same InD classes but different datasets. The model trained on FSDD is able to reach high accuracy on the FSDD test sets whereas the one trained on GSC cannot. From the plots, we see that the model trained on FSDD actually overfits to the data since it is able to distinguish between InD classes from different datasets (Fig. 8, right).

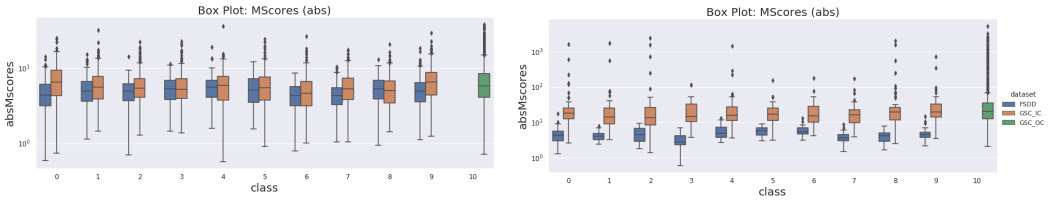

Figure 8: Distribution of TIME-LAPSE scores with trained Kymatio encoder. (left) Kymatio encoder trained with GSC 0-9 data, achieving 0.75 classification accuracy. The InD scores from the two datasets (blue and orange) are indistinguishable (right) Kymatio encoder trained with FSDD 0-9 data, achieving 0.97 classification accuracy. The InD scores between the two datasets are separable, showing that it has overfit to the dataset statistics.

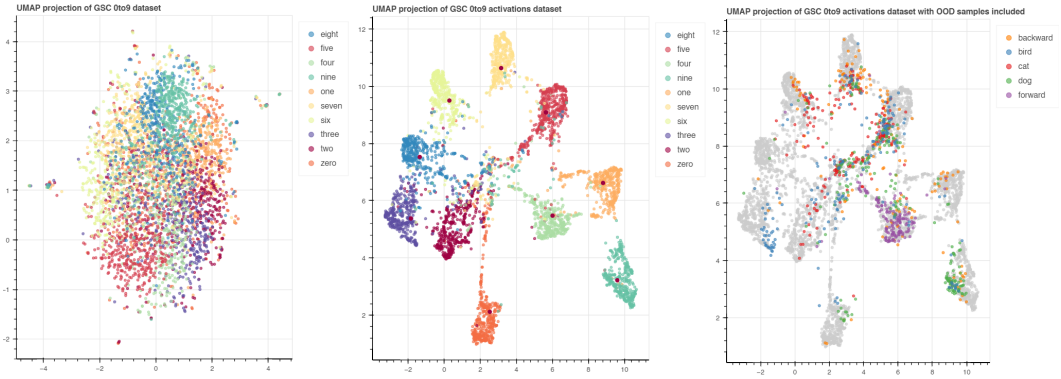

Figure 9: UMAP embeddings of GSC data (left) raw InD data (middle) representations from the trained encoder, i.e. a 2D projection of the latent space (c) OOD samples added to the mix (only few classes are shown for visualization purposes).

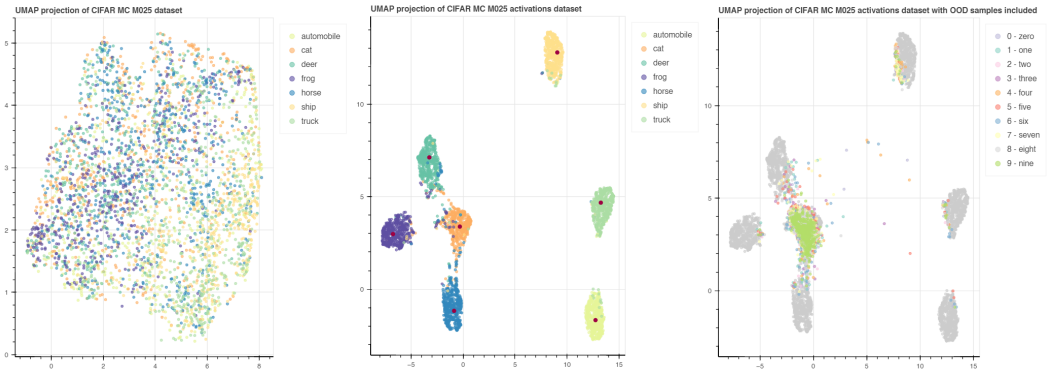

Figure 10: UMAP embeddings of CIFAR data (left) raw InD data (middle) representations from the trained encoder, i.e. a 2D projection of the latent space (c) OOD samples added to the mix (only few classes are shown for visualization purposes).

E Qualitative Analyses

We plot the UMAP (McInnes et al., 2018) embeddings of our raw InD data and their extracted latent space representations from the trained encoders along with those of OOD data. Fig. 9 shows the plots for our audio GSC dataset. We see that the raw data are too high-dimensional to be captured according to their class labels by UMAP. On the other hand, the extracted representations arrange themselves in well-separable clusters showing the representational capacity of the trained encoder. Most distance-based OOD detectors assume OOD representations will get mapped far away from the InD representations in the latent-space. This is a good assumption and holds in many cases. Indeed, the spatial scorer in TIME-LAPSE uses this hypothesis (along with others) to derive its predictive uncertainty score. However, OOD inputs can get mapped very close to InD inputs in the representational space under most distance metrics. Fig. 9 (right) shows the overlay of OOD inputs in the latent-space. Although it is just 2D and the extracted representations have higher dimensions, this figure still shows that OOD inputs can get mapped near to InD inputs. In that case, spatial scores lose their power. Since TIME-LAPSE has a temporal component to it, the temporal scorer acts as an additional constraint to ensure OOD inputs are separable. Thus, within a few samples, TIME-LAPSE can indicate when samples are going out-of-distribution with much better accuracy.

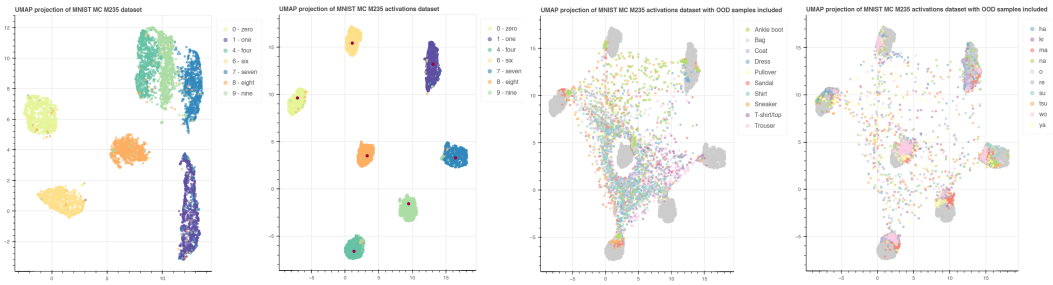

Figure 11: UMAP embeddings of MNIST data (left) raw InD data (middle) representations from the trained encoder, i.e. a 2D projection of the latent space (c) OOD samples added to the mix (only few classes are shown for visualization purposes).

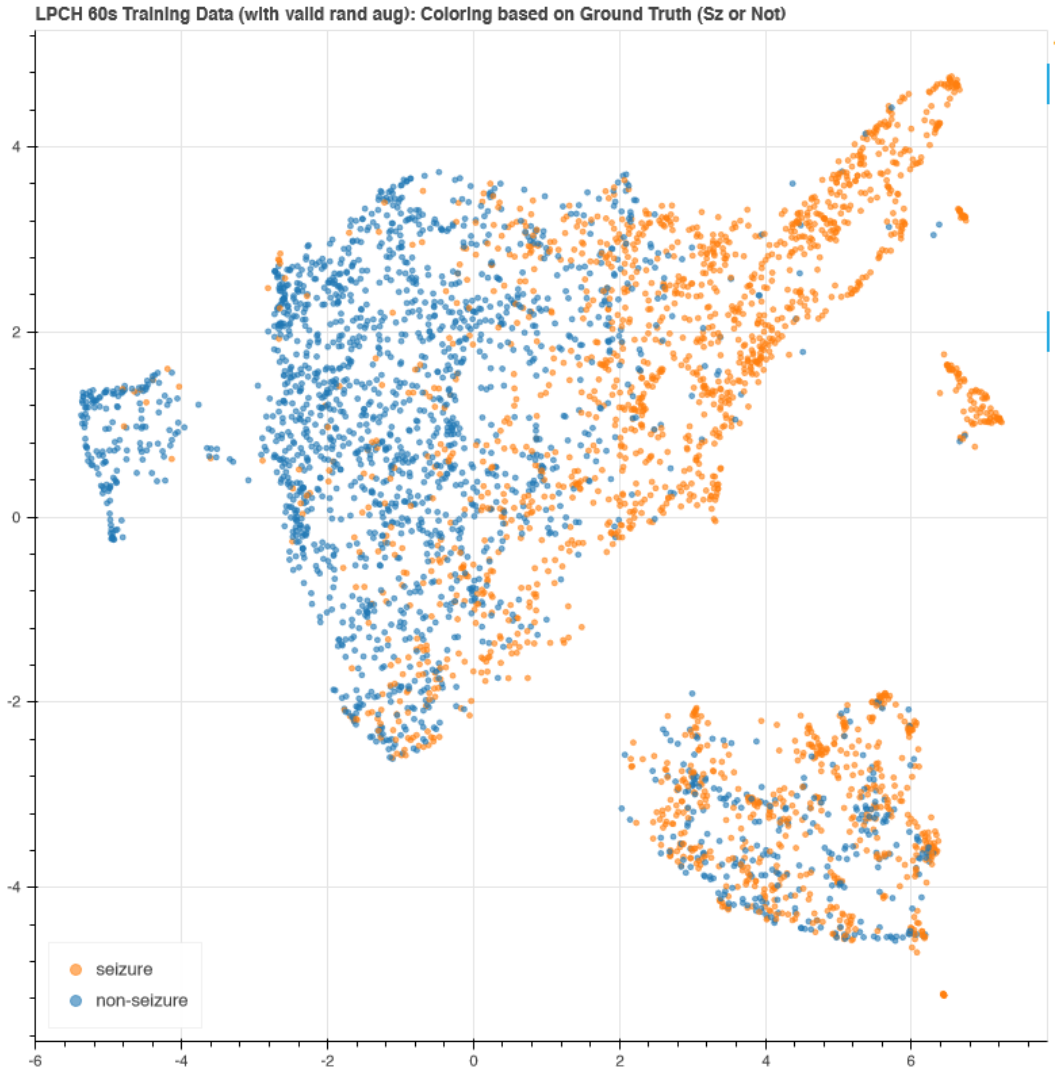

Figure 12: UMAP embeddings of EEG data (left) raw InD data (middle) representations from the trained encoder, i.e. a 2D projection of the latent space (c) OOD samples added to the mix (only few classes are shown for visualization purposes).

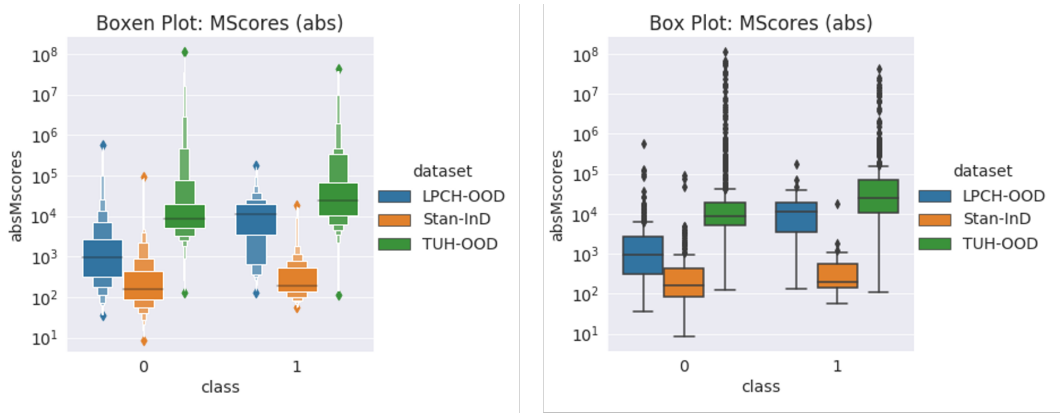

Figure 13: TIME-LAPSE score distribution on EEG data

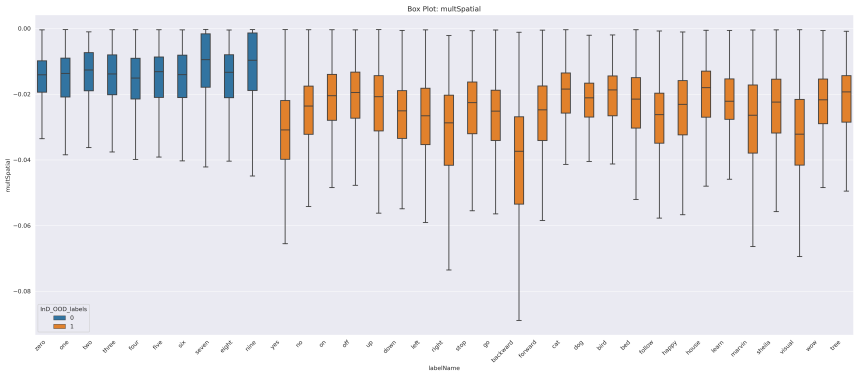

Figure 14: TIME-LAPSE score distribution on GSC

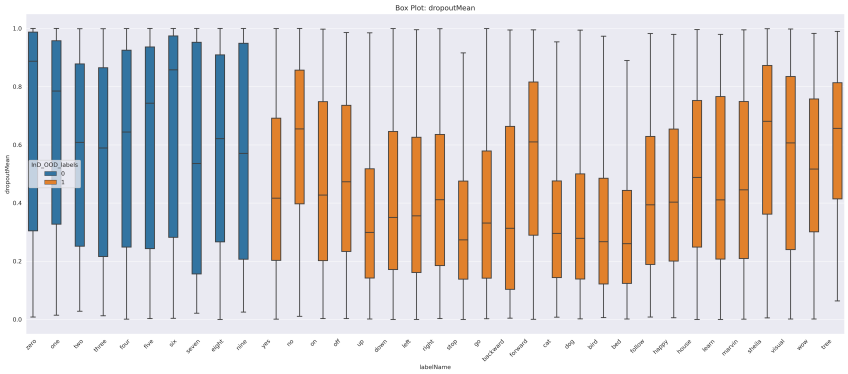

Figure 15: Test-time Dropout score distribution on GSC