

Latent Space Autoregression for Novelty Detection

Davide Abati Angelo Porrello Simone Calderara Rita Cucchiara

University of Modena and Reggio Emilia

{name.surname}@unimore.it

Abstract

Novelty detection is commonly referred to as the discrimination of observations that do not conform to a learned model of regularity. Despite its importance in different application settings, designing a novelty detector is utterly complex due to the unpredictable nature of novelties and its inaccessibility during the training procedure, factors which expose the unsupervised nature of the problem. In our proposal, we design a general framework where we equip a deep autoencoder with a parametric density estimator that learns the probability distribution underlying its latent representations through an autoregressive procedure. We show that a maximum likelihood objective, optimized in conjunction with the reconstruction of normal samples, effectively acts as a regularizer for the task at hand, by minimizing the differential entropy of the distribution spanned by latent vectors. In addition to providing a very general formulation, extensive experiments of our model on publicly available datasets deliver on-par or superior performances if compared to state-of-the-art methods in one-class and video anomaly detection settings. Differently from prior works, our proposal does not make any assumption about the nature of the novelties, making our work readily applicable to diverse contexts.

1. Introduction

Novelty detection is defined as the identification of samples which exhibit significantly different traits with respect to an underlying model of regularity, built from a collection of normal samples. The awareness of an autonomous system to recognize unknown events enables applications in several domains, ranging from video surveillance [7, 11], to defect detection [19] to medical imaging [35]. Moreover, the surprise induced by unseen events is emerging as a crucial aspect in reinforcement learning settings, as an enabling factor in curiosity-driven exploration [31].

However, in this setting, the definition and labeling of novel examples are not possible. Accordingly, the literature

agrees on approximating the ideal shape of the boundary separating normal and novel samples by modeling the intrinsic characteristics of the former. Therefore, prior works tackle such problem by following principles derived from the unsupervised learning paradigm [9, 34, 11, 23, 27]. Due to the lack of a supervision signal, the process of feature extraction and the rule for their normality assessment can only be guided by a proxy objective, assuming the latter will define an appropriate boundary for the application at hand.

According to cognitive psychology [4], novelty can be expressed either in terms of capabilities to *remember* an event or as a degree of *surprisal* [39] aroused by its observation. The latter is mathematically modeled in terms of low probability to occur under an expected model, or by lowering a variational free energy [15]. In this framework, prior models take advantage of either parametric [46] or non-parametric [13] density estimators. Differently, remembering an event implies the adoption of a memory represented either by a dictionary of normal prototypes - as in sparse coding approaches [9] - or by a low dimensional representation of the input space, as in the self-organizing maps [18] or, more recently, in deep autoencoders. Thus, in novelty detection, the remembering capability for a given sample is evaluated either by measuring reconstruction errors [11, 23] or by performing discriminative in-distribution tests [34].

Our proposal contributes to the field by merging remembering and surprisal aspects into a unique framework: we design a generative unsupervised model (i.e., an autoencoder, represented in Fig. 1i) that exploits end-to-end training in order to maximize remembering effectiveness for normal samples whilst minimizing the surprisal of their latent representation. This latter point is enabled by the maximization of the likelihood of latent representations through an autoregressive density estimator, which is performed in conjunction with the reconstruction error minimization. We show that, by optimizing both terms jointly, the model implicitly seeks for minimum entropy representations maintaining its remembering/reconstructive power. While entropy minimization approaches have been adopted in deep neural compression [3], to our knowledge this is the first proposal

tailored for novelty detection. In memory terms, our procedure resembles the concept of prototyping the normality using as few templates as possible. Moreover, evaluating the output of the estimator enables the assessment of the surprisal aroused by a given sample.

2. Related work

Reconstruction-based methods. On the one hand, many works lean toward learning a parametric projection and reconstruction of normal data, assuming outliers will yield higher residuals. Traditional sparse-coding algorithms [45, 9, 24] adhere to such framework, and represent normal patterns as a linear combination of a few basis components, under the hypotheses that novel examples would exhibit a non-sparse representation in the learned subspace. In recent works, the projection step is typically drawn from deep autoencoders [11]. In [27] the authors recover sparse coding principles by imposing a sparsity regularization over the learned representations, while a recurrent neural network enforces their smoothness along the time dimension. In [34], instead, the authors take advantage of an adversarial framework in which a discriminator network is employed as the actual novelty detector, spotting anomalies by performing a discrete in-distribution test. Oppositely, future frame prediction [23] maximizes the expectation of the next frame exploiting its knowledge of the past ones; at test time, observed deviations against the predicted content advise for abnormality. Differently from the above-mentioned works, our proposal relies on modeling the prior distribution of latent representations. This choice is coherent with recent works from the density estimation community [38, 6]. However, to the best of our knowledge, our work is the first advocating for the importance of such a design choice for novelty detection.

Probabilistic methods. A complementary line of research investigates different strategies to approximate the density function of normal appearance and motion features. The primary issue raising in this field concerns how to estimate such densities in a high-dimensional and complex feature space. In this respect, prior works involve hand-crafted features such as optical flow or trajectory analysis and, on top of that, employ both non-parametric [1] and parametric [5, 28, 22] estimators, as well as graphical modeling [16, 20]. Modern approaches rely on deep representations (e.g., captured by autoencoders), as in Gaussian classifiers [33] and Gaussian Mixtures [46]. In [13] the authors involve a Kernel Density Estimator (KDE) modeling activations from an auxiliary object detection network. A recent research trend considers training Generative Adversarial Networks (GANs) on normal samples. However, as such models approximate an implicit density function, they can be queried for new samples

but not for likelihood values. Therefore, GAN-based models employ different heuristics for the evaluation of novelty. For instance, in [35] a guided latent space search is exploited to infer it, whereas [32] directly queries the discriminator for a normality score.

3. Proposed model

Maximizing the probability of latent representations is analogous to lowering the surprisal of the model for a normal configuration, defined as the negative log-density of a latent variable instance [39]. Conversely, remembering capabilities can be evaluated by the reconstruction accuracy of a given sample under its latent representation.

We model the aforementioned aspects in a latent variable model setting, where the density function of training samples $p(\mathbf{x})$ is modeled through an auxiliary random variable \mathbf{z} , describing the set of causal factors underlying all observations. By factorizing

$$p(\mathbf{x}) = \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z}, \quad (1)$$

where $p(\mathbf{x}|\mathbf{z})$ is the conditional likelihood of the observation given a latent representation \mathbf{z} with prior distribution $p(\mathbf{z})$, we can explicit both the memory and surprisal contribution to novelty. We approximate the marginalization by means of an inference model responsible for the identification of latent space vector for which the contribution of $p(\mathbf{x}|\mathbf{z})$ is maximal. Formally, we employ a deep autoencoder, in which the reconstruction error plays the role of the negative logarithm of $p(\mathbf{x}|\mathbf{z})$, under the hypothesis that $p(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}|\tilde{\mathbf{x}}, I)$ where $\tilde{\mathbf{x}}$ denotes the output reconstruction. Additionally, surprisal is injected in the process by equipping the autoencoder with an auxiliary deep parametric estimator learning the prior distribution $p(\mathbf{z})$ of latent vectors, and training it by means of Maximum Likelihood Estimation (MLE). Our architecture is therefore composed of three building blocks (Fig. 1i): an encoder $f(\mathbf{x}; \theta_f)$, a decoder $g(\mathbf{z}; \theta_g)$ and a probabilistic model $h(\mathbf{z}; \theta_h)$:

$$\begin{aligned} f(\mathbf{x}; \theta_f) : \mathbb{R}^m &\rightarrow \mathbb{R}^d, & g(\mathbf{z}; \theta_g) : \mathbb{R}^d &\rightarrow \mathbb{R}^m, \\ h(\mathbf{z}; \theta_h) : \mathbb{R}^d &\rightarrow [0, 1]. \end{aligned} \quad (2)$$

The encoder processes input \mathbf{x} and maps it into a compressed representation $\mathbf{z} = f(\mathbf{x}; \theta_f)$, whereas the decoder provides a reconstructed version of the input $\tilde{\mathbf{x}} = g(\mathbf{z}; \theta_g)$. The probabilistic model $h(\mathbf{z}; \theta_h)$ estimates the density in \mathbf{z} via an autoregressive process, allowing to avoid the adoption of a specific family of distributions (i.e., Gaussian), potentially unrewarding for the task at hand. On this latter point, please refer to supplementary materials for comparison w.r.t. variational autoencoders [17].

With such modules, at test time, we can assess the two sources of novelty: elements whose observation is poorly

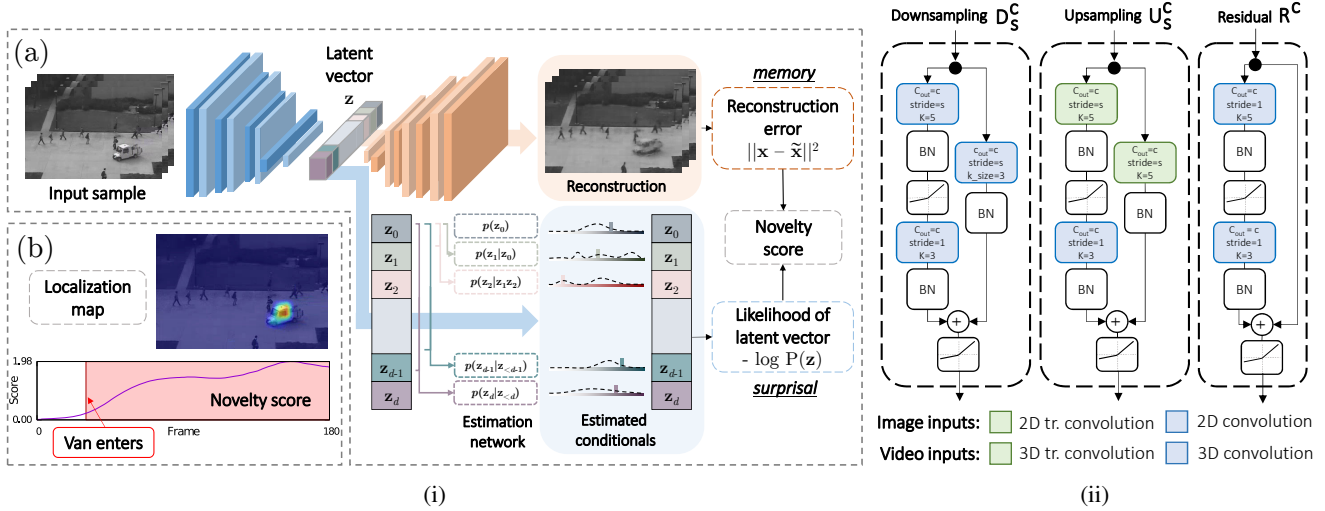


Figure 1: (i) The proposed novelty detection framework. The overall architecture, depicted in (a), consists of a deep autoencoder and an autoregressive estimation network operating on its latent space. The joint minimization of their respective objective leads to a measure of novelty - (b) - obtained by assessing the remembrance of the model when looking to a new sample, combined with its surprise aroused by causal factors. (ii) Building blocks employed in the autoencoder’s architecture.

explained by the causal factors inducted by normal samples (i.e., high reconstruction error); elements exhibiting good reconstructions whilst showing surprising underlying representations under the learned prior.

Autoregressive density estimation. Autoregressive models provide a general formulation for tasks involving sequential predictions, in which each output depends on previous observations [25, 29]. We adopt such a technique to factorize a joint distribution, thus avoiding to define its landscape a priori [21, 40]. Formally, $p(\mathbf{z})$ is factorized as

$$p(\mathbf{z}) = \prod_{i=1}^d p(z_i | \mathbf{z}_{<i}), \quad (3)$$

so that estimating $p(\mathbf{z})$ reduces to the estimation of each single Conditional Probability Density (CPD) expressed as $p(z_i | \mathbf{z}_{<i})$, where the symbol $<$ implies an order over random variables. Some prior models obey handcrafted orderings [43, 42], whereas others rely on order agnostic training [41, 10]. Nevertheless, it is still not clear how to estimate the proper order for a given set of variables. In our model, this issue is directly tackled by the optimization. Indeed, since we perform autoregression on learned latent representations, the MLE objective encourages the autoencoder to impose over them a pre-defined causal structure. Empirical evidence of this phenomenon is given in the supplementary material.

From a technical perspective, the estimator $h(\mathbf{z}; \theta_h)$ outputs parameters for d distributions $p(z_i | \mathbf{z}_{<i})$. In our implementation, each CPD is modeled as a multinomial over $B=100$ quantization bins. To ensure a conditional estimate of each

underlying density, we design proper layers guaranteeing that the CPD of each symbol z_i is computed from inputs $\{z_1, \dots, z_{i-1}\}$ only.

Objective and connection with differential entropy. The three components f , g and h are jointly trained to minimize $\mathcal{L} \equiv \mathcal{L}(\theta_f, \theta_g, \theta_h)$ as follows:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{\text{REC}}(\theta_f, \theta_g) + \lambda \mathcal{L}_{\text{LLK}}(\theta_f, \theta_h) \\ &= \mathbb{E}_{\mathbf{x}} \left[\underbrace{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}_{\text{reconstruction term}} - \lambda \underbrace{\log(h(\mathbf{z}; \theta_h))}_{\text{log-likelihood term}} \right], \end{aligned} \quad (4)$$

where λ is a hyper-parameter controlling the weight of the \mathcal{L}_{LLK} term. It is worth noting that it is possible to express the log-likelihood term as

$$\begin{aligned} &\mathbb{E}_{\mathbf{z} \sim p^*(\mathbf{z}; \theta_f)} [-\log h(\mathbf{z}; \theta_h)] \\ &= \mathbb{E}_{\mathbf{z} \sim p^*(\mathbf{z}; \theta_f)} [-\log h(\mathbf{z}; \theta_h) + \log p^*(\mathbf{z}; \theta_f) - \log p^*(\mathbf{z}; \theta_f)] \\ &= D_{\text{KL}}(p^*(\mathbf{z}; \theta_f) \parallel h(\mathbf{z}; \theta_h)) + \mathbb{H}[p^*(\mathbf{z}; \theta_f)], \end{aligned} \quad (5)$$

where $p^*(\mathbf{z}; \theta_f)$ denotes the true distribution of the codes produced by the encoder, and is therefore parametrized by θ_f . This reformulation of the MLE objective yields meaningful insights about the entities involved in the optimization. On the one hand, the Kullback-Leibler divergence ensures that the information gap between our parametric model h and the true distribution p^* is small. On the other hand, this framework leads to the minimization of the differential entropy of the distribution underlying the codes produced by the encoder f . Such constraint constitutes a crucial point when learning normality. Intuitively, if we think




Figure 2: Estimated differential entropies delivered on each MNIST class in the presence of different regularization strategies: our, divergence w.r.t a Gaussian prior (VAE) and input perturbation (DAE). For each class, the estimate is computed on the training samples’ hidden representations, whose distribution are fit utilizing a Gaussian KDE in a 3D-space. All models being equal, ours exhibits lower entropies on all classes.

about the encoder as a source emitting symbols (namely, the latent representations), its desired behavior, when modeling normal aspects in the data, should converge to a ‘boring’ process characterized by an intrinsic low entropy, since surprising and novel events are unlikely to arise during the training phase. Accordingly, among all the possible settings of the hidden representations, the objective begs the encoder to exhibit a low differential entropy, leading to the extraction of features that are easily predictable, therefore common and recurrent within the training set. This kind of features is indeed the most useful to distinguish novel samples from the normal ones, making our proposal a suitable regularizer in the anomaly detection setting.

We report empirical evidence of the decreasing differential entropy in Fig. 2, that compares the behavior of the same model under different regularization strategies.

3.1. Architectural Components

Autoencoder blocks. Encoder and decoder are respectively composed by downsampling and upsampling residual blocks depicted in Fig. 1ii. The encoder ends with fully connected (FC) layers. When dealing with video inputs, we employ *causal* 3D convolutions [2] within the encoder (i.e., only accessing information from previous time-steps). Moreover, at the end of the encoder, we employ a temporally-shared full connection (TFC, namely a linear projection sharing parameters across the time axis on the input feature maps) resulting in a temporal series of feature vectors. This way, the encoding procedure does not shuffle information across time-steps, ensuring temporal ordering.

Autoregressive layers. To guarantee the autoregressive nature of each output CPD, we need to ensure proper

connectivity patterns in each layer of the estimator h . Moreover, since latent representations exhibit different shapes depending on the input nature (image or video), we propose two different solutions.

When dealing with images, the encoder provides feature vectors with dimensionality d . The autoregressive estimator is composed by stacking multiple Masked Fully Connections (MFC, Fig. 3-(a)). Formally, it computes output feature map $\mathbf{o} \in \mathbb{R}^{d \times co}$ (where co is the number of output channels) given the input $\mathbf{h} \in \mathbb{R}^{d \times ci}$ (assuming $ci = 1$ at the input layer). The connection between the input element \mathbf{h}_i^k in position i , channel k and the output element \mathbf{o}_j^l is parametrized by

$$\begin{cases} w_{i,j}^{k,l} & \text{if } i < j \\ \begin{cases} w_{i,j}^{k,l} & \text{if type = B} \\ 0 & \text{if type = A} \end{cases} & \text{if } i = j \\ 0 & \text{if } i > j. \end{cases} \quad (6)$$

Type A forces a strict dependence on previous elements (and is employed only as the first estimator layer), whereas type B masks only succeeding elements. Assuming each CPD modeled as a multinomial, the output of the last autoregressive layer (in $\mathbb{R}^{d \times B}$) provides probability estimates for the B bins that compose the space quantization.

On the other hand, the compressed representation of video clips has dimensionality $t \times d$, being t the number of temporal time-steps and d the length of the code. Accordingly, the estimation network is designed to capture two-dimensional patterns within observed elements of the code. However, naively plugging 2D convolutional layers would assume translation invariance on both axes of the input map, whereas, due to the way the compressed representation is built, this assumption is only correct along the temporal axis. To cope with this, we apply d different convolutional kernels along the code axis, allowing the observation of the whole feature vector in the previous time-step as well as a portion of the current one. Every convolution is free to stride along the time axis and captures temporal patterns. In such operation, named Masked Stacked Convolution (MSC, Fig. 3-(b)), the i -th convolution is equipped with a kernel $\mathbf{w}^{(i)} \in \mathbb{R}^{3 \times d}$ kernel, that gets multiplied by the binary mask $\mathbf{M}^{(i)}$, defined as

$$m_{j,k}^{(i)} \in \mathbf{M}^{(i)} = \begin{cases} 1 & \text{if } j = 0 \\ 1 & \text{if } j = 1 \text{ and } k < i \text{ and type=A} \\ 1 & \text{if } j = 1 \text{ and } k \leq i \text{ and type=B} \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where j indexes the temporal axis and k the code axis.

Every single convolution yields a column vector, as a result of its stride along time. The set of column vectors resulting




Figure 3: Proposed autoregressive layers, namely the Masked Fully Connection (a, Eq. 6) and the Masked Stacked Convolution (b, Eq. 7). For both layers, we represent type A structure. Different kernel colors represent different parametrizations.

from the application of the d convolutions to the input tensor $\mathbf{h} \in \mathbb{R}^{t \times d \times ci}$ are horizontally stacked to build the output tensor $\mathbf{o} \in \mathbb{R}^{t \times d \times co}$, as follows:

$$\mathbf{o} = \left\| \prod_{i=1}^d [(\mathbf{M}^{(i)} \odot \mathbf{w}^{(i)}) * \mathbf{h}] \right\|, \quad (8)$$

where $\|$ represents the horizontal concatenation operation.

4. Experiments¹

We test our solution in three different settings: images, videos, and cognitive data. In all experiments the novelty assessment on the i -th example is carried out by summing the reconstruction term (REC_i) and the log-likelihood term (LLK_i) in Eq. 4 in a single novelty score NS_i :

$$NS_i = norm_S(REC_i) + norm_S(LLK_i). \quad (9)$$

Individual scores are normalized using a reference set of examples S (different for every experiment),

$$norm_S(L_i) = \frac{L_i - \max_{j \in S} L_j}{\max_{j \in S} L_j - \min_{j \in S} L_j}. \quad (10)$$

Further implementation details and architectural hyperparameters are in the supplementary material.

4.1. One-class novelty detection on images

To assess the model's performances in one class settings, we train it on each class of either MNIST or CIFAR-10 separately. In the test phase, we present the corresponding test set, which is composed of 10000 examples of all classes, and expect our model to assign a lower novelty score to images sharing the label with training samples. We use standard train/test splits, and isolate 10% of training samples for

validation purposes, and employ it as the normalization set (S in Eq. 9) for the computation of the novelty score. As for the baselines, we consider the following:

- standard methods such as OC-SVM [36] and Kernel Density Estimator (KDE), employed out of features extracted by PCA-whitening;
- a denoising autoencoder (DAE) sharing the same architecture as our proposal, but defective of the density estimation module. The reconstruction error is employed as a measure of normality vs. novelty;
- a variational autoencoder (VAE) [17], also sharing the same architecture as our model, in which the Evidence Lower Bound (ELBO) is employed as the score;
- Pix-CNN [42], modeling the density by applying autoregression directly in the image space;
- the GAN-based approach illustrated in [35].

We report the comparison in Tab. 1 in which performances are measured by the Area Under Receiver Operating Characteristic (AUROC), which is the standard metric for the task. As the table shows, our proposal outperforms all baselines in both settings.

Considering MNIST, most methods perform favorably. Notably, Pix-CNN fails in modeling distributions for all digits but one, possibly due to the complexity of modeling densities directly on pixel space and following a fixed autoregression order. Such poor test performances are registered despite good quality samples that we observed during training: indeed, the weak correlation between sample quality and test log-likelihood of the model has been motivated in [37]. Surprisingly, OC-SVM outperforms most deep learning based models in this setting.

On the contrary, CIFAR10 represents a much more significant challenge, as testified by the low performances of most models, possibly due to the poor image resolution and visual clutter between classes. Specifically, we observe

¹Code to reproduce results in this section is released at <https://github.com/aimagelab/novelty-detection>.

	MNIST							CIFAR10						
	OC SVM	KDE	DAE	VAE	Pix CNN	GAN	ours	OC SVM	KDE	DAE	VAE	Pix CNN	GAN	ours
0	0.988	0.885	0.991	0.998	0.531	0.926	0.993	0.630	0.658	0.718	0.688	0.788	0.708	0.735
1	0.999	0.996	0.999	0.999	0.995	0.995	0.999	0.440	0.520	0.401	0.403	0.428	0.458	0.580
2	0.902	0.710	0.891	0.962	0.476	0.805	0.959	0.649	0.657	0.685	0.679	0.617	0.664	0.690
3	0.950	0.693	0.935	0.947	0.517	0.818	0.966	0.487	0.497	0.556	0.528	0.574	0.510	0.542
4	0.955	0.844	0.921	0.965	0.739	0.823	0.956	0.735	0.727	0.740	0.748	0.511	0.722	0.761
5	0.968	0.776	0.937	0.963	0.542	0.803	0.964	0.500	0.496	0.547	0.519	0.571	0.505	0.546
6	0.978	0.861	0.981	0.995	0.592	0.890	0.994	0.725	0.758	0.642	0.695	0.422	0.707	0.751
7	0.965	0.884	0.964	0.974	0.789	0.898	0.980	0.533	0.564	0.497	0.500	0.454	0.471	0.535
8	0.853	0.669	0.841	0.905	0.340	0.817	0.953	0.649	0.680	0.724	0.700	0.715	0.713	0.717
9	0.955	0.825	0.960	0.978	0.662	0.887	0.981	0.508	0.540	0.389	0.398	0.426	0.458	0.548
avg	0.951	0.814	0.942	0.969	0.618	0.866	0.975	0.586	0.610	0.590	0.586	0.551	0.592	0.641

Table 1: AUROC results for novelty detection on MNIST and CIFAR10. Each row represents a different class on which baselines and our model are trained.

that our proposal is the only model outperforming a simple KDE baseline; however, this finding should be put into perspective by considering the nature of non-parametric estimators. Indeed, non-parametric models are allowed to access the whole training set for the evaluation of each sample. Consequently, despite they benefit large sample sets in terms of density modeling, they lead into an unfeasible inference as the dataset grows in size.

The possible reasons behind the difference in performance w.r.t. DAE are twofold. Firstly, DAE can recognize novel samples solely based on the reconstruction error, hence relying on its memorization capabilities, whereas our proposal also considers the likelihood of their representations under the learned prior, thus exploiting surprisal as well. Secondly, by minimizing the differential entropy of the latent distribution, our proposal increases the discriminative capability of the reconstruction. Intuitively, this last statement can be motivated observing that novelty samples are forced to reside in a high probability region of the latent space, the latter bounded to solely capture unsurprising factors of variation arising from the training set. On the other hand, the gap w.r.t. VAE suggests that, for the task at hand, a more flexible autoregressive prior should be pre-

ferred over the isotropic multivariate Gaussian. On this last point, VAE seeks representations whose average surprisal converges to a fixed and expected value (i.e., the differential entropy of its prior), whereas our solution minimizes such quantity within its MLE objective. This flexibility allows modulating the richness of the latent representation vs. the reconstructing capability of the model. On the contrary, in VAEs, the fixed prior acts as a blind regularizer, potentially leading to over-smooth representations; this aspect is also appreciable when sampling from the model as shown in the supplementary material.

Fig. 4 reports an ablation study questioning the loss functions aggregation presented in Eq. 9. The figure illustrates ROC curves under three different novelty scores: i) the log-likelihood term, ii) the reconstruction term, and iii) the proposed scheme that accounts for both. As highlighted in the picture, accounting for both memorization and surprisal aspects is advantageous in each dataset. Please refer to the supplementary material for additional evidence.

4.2. Video anomaly detection

In video surveillance contexts, novelty is often considered in terms of abnormal human behavior. Thus, we evaluate our proposal against state-of-the-art anomaly detection models. For this purpose, we considered two standard benchmarks in literature, namely UCSD Ped2 [8] and ShanghaiTech [27]. Despite the differences in the number of videos and their resolution, they both contain anomalies that typically arise in surveillance scenarios (e.g., vehicles in pedestrian walkways, pick-pocketing, brawling). For UCSD Ped, we preprocessed input clips of 16 frames to extract smaller patches (we refer to supplementary materials for details) and perturbed such inputs with random Gaussian noise with $\sigma = 0.025$. We compute the novelty score of each input clip as the mean novelty score among all patches. Concerning ShanghaiTech, we removed the dependency on




Figure 4: ROC curves delivered by different scoring strategies on MNIST and CIFAR-10 test sets. Each curve is an interpolation over the ten classes.

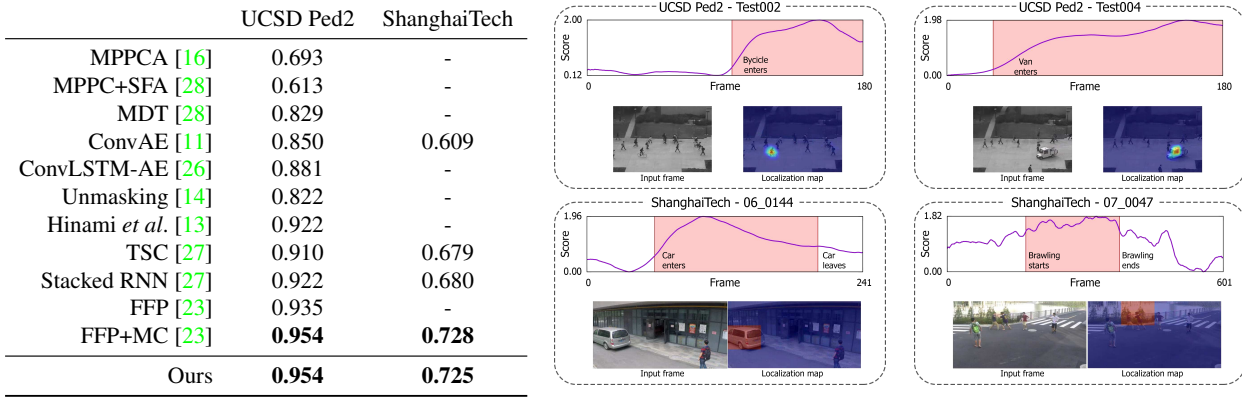


Figure 5: On the left, AUROC performances of our model w.r.t. state-of-the-art competitors. On the right, novelty scores and localizations maps for samples drawn from UCSD Ped2 and ShanghaiTech. For each example, we report the trend of the assessed score, highlighting with a different color the time range in which an anomalous subject comes into the scene.

the scenario by estimating the foreground for each frame of a clip with a standard MOG-based approach and removing the background. We fed the model with 16-frames clips, but ground-truth anomalies are labeled at frame level. In order to recover the novelty score of each frame, we compute the mean score of all clips in which it appears. We then merge the two terms of the loss function following the same strategy illustrated in Eq. 9, computing however normalization coefficients in a per-sequence basis, following the standard approach in the anomaly detection literature. The scores for each sequence are then concatenated to compute the overall AUROC of the model. Additionally, we envision localization strategies for both datasets. To this aim, for UCSD, we denote a patch exhibiting the highest novelty score in a frame as anomalous. Differently, in ShanghaiTech, we adopt a sliding-window approach [44]: as expected, when occluding the source of the anomaly with a rectangular patch, the novelty score drops significantly. Fig. 5 reports results in comparison with prior works, along with qualitative assessments regarding the novelty score and localization capabilities. Despite a more general formulation, our proposal scores on-par with the current state-of-the-art solutions specifically designed for video applications and taking advantage of optical flow estimation and motion constraints. Indeed, in the absence of such hypotheses (FFP entry in Fig. 5), our method outperforms future frame prediction on UCSD Ped2.

4.3. Model Analysis

CIFAR-10 with semantic features. We investigate the behavior of our model in the presence of different assumptions regarding the expected nature of novel samples. We expect that, as the correctness of such assumptions increases, novelty detection performances will scale accordingly. Such a trait is particularly desirable for applications in which prior beliefs about novel examples

can be envisioned. To this end, we leverage the CIFAR-10 benchmark described in Sec. 4.1 and change the type of information provided as input. Specifically, instead of raw images, we feed our model with semantic representations extracted by ResNet-50 [12], either pre-trained on Imagenet (i.e., assume semantic novelty) or CIFAR-10 itself (i.e., assume data-specific novelty). The two models achieved respectively 79.26 and 95.4 top-1 classification accuracies on the respective test sets. Even though this procedure is to be considered unfair in novelty detection, it serves as a sanity check delivering the upper-bound performances our model can achieve when applied to even better features. To deal with dense inputs, we employ a fully connected autoencoder and MFC layers within the estimation network. Fig. 6-(a) illustrates the resulting ROC curves, where semantic descriptors improve AUROC w.r.t. raw image inputs (entry “Unsupervised”). Such results suggest that our model profitably takes advantage of the separation between normal and abnormal input representations and scales accordingly, even up to optimal performances for the task under consideration. Nevertheless, it is interesting to note how different degrees of supervision deliver significantly different performances. As expected, dataset-specific supervision increases the AUROC from 0.64 up to 0.99 (a perfect score). Surprisingly, semantic feature vectors trained on Imagenet (which contains all CIFAR classes) provide a much lower boost, yielding an AUROC of 0.72. Such result suggests that, even in the rare cases where the semantic of novelty can be known in advance, its contribution has a limited impact in modeling the normality, mostly because novelty can depend on other cues (e.g., low-level statistics).

Autoregression via recurrent layers. To measure the contribution of the proposed MFC and MSC layers described in Sec. 3, we test on CIFAR-10 and UCSD

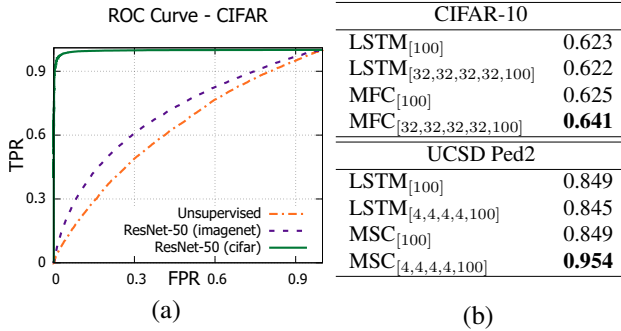


Figure 6: (a) CIFAR-10 ROC curves with semantic input vectors. Each curve is an interpolation among the ten classes. (b) Comparison of different architectures for the autoregressive density estimation in feature space. We indicate with LSTM_[F₁,F₂,...,F_N] - same goes for MFC and MSC - the output shape for each of the N layers composing the estimator. Results are reported in terms of test AUROC.

Ped2, alternative solutions for the autoregressive density estimator. Specifically, we investigate recurrent networks, as they represent the most natural alternative featuring autoregressive properties. We benchmark the proposed building blocks against an estimator composed of LSTM layers, which is designed to sequentially observe latent symbols $\mathbf{z}_{<i}$ and output the CPD of z_i as the hidden state of the last layer. We test MFC, MSC and LSTM in single-layer and multi-layer settings, and report all outcomes in Fig. 6-(b).

It emerges that, even though our solutions perform similarly to the recurrent baseline when employed in a shallow setting, they significantly take advantage of their depth when stacked in consecutive layers. MFC and MSC, indeed, employ disentangled parametrizations for each output CPD. This property is equivalent to the adoption of a specialized estimator network for each z_i , thus increasing the proficiency in modeling the density of its designated CPD. On the contrary, LSTM networks embed all the history (i.e., the observed symbols) in their memory cells, but manipulate each input of the sequence through the same weight matrices. In such a regime, the recurrent module needs to learn parameters shared among symbols, losing specialization and eroding its modeling capabilities.

4.4. Novelty in cognitive temporal processes

As a potential application of our proposal, we investigate its capability in modeling human attentional behavior. To this end, we employ the DR(eye)VE dataset [30], introduced for the prediction of focus of attention in driving contexts. It features 74 driving videos where frame-wise fixation maps are provided, highlighting the region of the scene attended by the driver. In order to capture the dynamics of attentional patterns, we purposely discard the visual content of




Figure 7: Left, the distribution of novelty scores assigned to normal patterns against attentional shifts labeled within the DR(eye)VE dataset. Right, DR(eye)VE clips yielding the highest novelty score (i.e., clips in which the attentional pattern shifts from the expected behavior). Interestingly, they depict some peculiar situations such as waiting for the traffic light or approaching a roundabout.

the scene and optimize our model on clips of fixation maps, randomly extracted from the training set. After training, we rely on the novelty score of each clip as a proxy for the uncommonness of an attentional pattern. Moreover, since the dataset features annotations of peculiar and unfrequent patterns (such as distractions, recording errors), we can measure the correlation of the captured novelty w.r.t. those. In terms of AUROC, our model scores 0.926, highlighting that novelty can arise from unexpected behaviors of the driver, such as distractions or other shifts in attention. Fig. 7 reports the different distribution of novelty scores for ordinary and peculiar events.

5. Conclusions

We propose a comprehensive framework for novelty detection. We formalize our model to capture the twofold nature of novelties, which concerns the incapability to remember unseen data and the surprisal aroused by the observation of their latent representations. From a technical perspective, both terms are modeled by a deep generative autoencoder, paired with an additional autoregressive density estimator learning the distribution of latent vectors by maximum likelihood principles. To this aim, we introduce two different masked layers suitable for image and video data. We show that the introduction of such an auxiliary module, operating in latent space, leads to the minimization of the encoder’s differential entropy, which proves to be a suitable regularizer for the task at hand. Experimental results show state-of-the-art performances in one-class and anomaly detection settings, fostering the flexibility of our framework for different tasks without making any data-related assumption.

Acknowledgements. We gratefully acknowledge Facebook Artificial Intelligence Research and Panasonic Silicon Valley Lab for the donation of GPUs used for this research.

References

- [1] Amit Adam, Ehud Rivlin, Ilan Shimshoni, and Daviv Reinitz. Robust real-time unusual event detection using multiple fixed-location monitors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):555–560, 2008. 2
- [2] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv:1803.01271*, 2018. 4
- [3] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *International Conference on Learning Representations*, 2017. 1
- [4] Andrew Barto, Marco Mirolli, and Gianluca Baldassarre. Novelty or surprise? *Frontiers in psychology*, 4:907, 2013. 1
- [5] Arslan Basharat, Alexei Gritai, and Mubarak Shah. Learning object motion patterns for anomaly detection and improved object detection. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2
- [6] Matthias Bauer and Andriy Mnih. Resampled priors for variational autoencoders. *International Conference on Artificial Intelligence and Statistics*, 2019. 2
- [7] Simone Calderara, Uri Heinemann, Andrea Prati, Rita Cucchiara, and Naftali Tishby. Detecting anomalies in peoples trajectories using spectral graph analysis. *Computer Vision and Image Understanding*, 115(8):1099–1111, 2011. 1
- [8] Antoni Chan and Nuno Vasconcelos. Ucsd pedestrian database. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2008. 6
- [9] Yang Cong, Junsong Yuan, and Ji Liu. Sparse reconstruction cost for abnormal event detection. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 3449–3456. IEEE, 2011. 1, 2
- [10] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pages 881–889, 2015. 3
- [11] Mahmudul Hasan, Jonghyun Choi, Jan Neumann, Amit K Roy-Chowdhury, and Larry S Davis. Learning temporal regularity in video sequences. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 733–742. IEEE, 2016. 1, 2, 7
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 7
- [13] Ryota Hinami, Tao Mei, and Shin’ichi Satoh. Joint detection and recounting of abnormal events by learning deep generic knowledge. In *IEEE International Conference on Computer Vision*, pages 3639–3647, 2017. 1, 2, 7
- [14] Radu Tudor Ionescu, Sorina Smeureanu, Bogdan Alexe, and Marius Popescu. Unmasking the abnormal events in video. *IEEE International Conference on Computer Vision*, 2017. 7
- [15] Laurent Itti and Pierre Baldi. Bayesian surprise attracts human attention. *Vision research*, 49(10):1295–1306, 2009. 1
- [16] Jaechul Kim and Kristen Grauman. Observe locally, infer globally: a space-time mrf for detecting abnormal activities with incremental updates. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 2921–2928. IEEE, 2009. 2, 7
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2014. 2, 5
- [18] Teuvo Kohonen. *Self-organization and associative memory*, volume 8. Springer Science & Business Media, 2012. 1
- [19] Ajay Kumar. Computer-vision-based fabric defect detection: A survey. *IEEE Transactions on Industrial Electronics*, 55(1):348–363, 2008. 1
- [20] Junseok Kwon and Kyoung Mu Lee. A unified framework for event summarization and rare event detection from multiple views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1737–1750, 2015. 2
- [21] Hugo Larochelle and Iain Murray. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 29–37, 2011. 3
- [22] Weixin Li, Vijay Mahadevan, and Nuno Vasconcelos. Anomaly detection and localization in crowded scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(1):18–32, 2014. 2
- [23] Wen Liu, Weixin Luo, Dongze Lian, and Shenghua Gao. Future frame prediction for anomaly detection – a new baseline. *IEEE International Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2, 7
- [24] Cewu Lu, Jianping Shi, and Jiaya Jia. Abnormal event detection at 150 fps in matlab. In *IEEE International Conference on Computer Vision*, pages 2720–2727. IEEE, 2013. 2
- [25] Pauline Luc, Natalia Neverova, Camille Couprie, Jakob Verbeek, and Yann LeCun. Predicting deeper into the future of semantic segmentation. In *IEEE International Conference on Computer Vision*, pages 648–657, 2017. 3
- [26] Weixin Luo, Wen Liu, and Shenghua Gao. Remembering history with convolutional lstm for anomaly detection. In *IEEE International Conference on Multimedia and Expo*, pages 439–444. IEEE, 2017. 7
- [27] Weixin Luo, Wen Liu, and Shenghua Gao. A revisit of sparse coding based anomaly detection in stacked rnn framework. *IEEE International Conference on Computer Vision*, 2017. 1, 2, 6, 7
- [28] Vijay Mahadevan, Weixin Li, Viral Bhalodia, and Nuno Vasconcelos. Anomaly detection in crowded scenes. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 1975–1981. IEEE, 2010. 2, 7
- [29] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016. 3
- [30] Andrea Palazzi, Davide Abati, Simone Calderara, Francesco Solera, and Rita Cucchiara. Predicting the driver’s focus of attention: the dr(eye)ve project. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 8
- [31] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised pre-

- diction. In *International Conference on Machine Learning*, volume 2017, 2017. 1
- [32] Mahdyar Ravanbakhsh, Enver Sangineto, Moin Nabi, and Nicu Sebe. Training adversarial discriminators for cross-channel abnormal event detection in crowds. *arXiv preprint arXiv:1706.07680*, 2017. 2
- [33] Mohammad Sabokrou, Mohsen Fayyaz, Mahmood Fathy, Zahra Moayed, and Reinhard Klette. Deep-anomaly: Fully convolutional neural network for fast anomaly detection in crowded scenes. *Computer Vision and Image Understanding*, 2018. 2
- [34] Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 3379–3388, 2018. 1, 2
- [35] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, pages 146–157. Springer, 2017. 1, 2, 5
- [36] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support vector method for novelty detection. In *Neural Information Processing Systems*, 2000. 5
- [37] Lucas Theis, Aaron van den Oord, and Matthias Bethge. A note on the evaluation of generative models. *International Conference on Learning Representations*, 2016. 5
- [38] Jakub M Tomczak and Max Welling. Vae with a vamp prior. *International Conference on Artificial Intelligence and Statistics*, 2018. 2
- [39] Myron Tribus. *Thermostatistics and thermodynamics: an introduction to energy, information and states of matter, with engineering applications*. van Nostrand, CS7, 1961. 1, 2
- [40] Benigno Uria, Iain Murray, and Hugo Larochelle. Rnade: The real-valued neural autoregressive density-estimator. In *Advances in Neural Information Processing Systems*, pages 2175–2183, 2013. 3
- [41] Benigno Uria, Iain Murray, and Hugo Larochelle. A deep and tractable density estimator. In *International Conference on Machine Learning*, pages 467–475, 2014. 3
- [42] A. van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, and A. Graves. Conditional image generation with pixelcnn decoders. In *Neural Information Processing Systems*, 2016. 3, 5
- [43] A. van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *International Conference on Machine Learning*, 2016. 3
- [44] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014. 7
- [45] Bin Zhao, Li Fei-Fei, and Eric P Xing. Online detection of unusual events in videos via dynamic sparse coding. In *IEEE International Conference on Computer Vision and Pattern Recognition*, pages 3313–3320. IEEE, 2011. 2
- [46] 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. In *International Conference on Learning Representations*, 2018. 1, 2