
Uncertainty-Aware Explanations Through Probabilistic Self-Explainable Neural Networks

Jon Vadillo¹, Roberto Santana¹, Jose A. Lozano^{1,2}, Marta Kwiatkowska³

¹Department of Computer Science and Artificial Intelligence,
University of the Basque Country UPV/EHU

²Basque Center for Applied Mathematics (BCAM)

³Department of Computer Science, University of Oxford
{jon.vadillo, roberto.santana, ja.lozano}@ehu.eus,
marta.kwiatkowska@cs.ox.ac.uk

Abstract

The lack of transparency of Deep Neural Networks continues to be a limitation that severely undermines their reliability and usage in high-stakes applications. Promising approaches to overcome such limitations are Prototype-Based Self-Explainable Neural Networks (PSENNs), whose predictions rely on the similarity between the input at hand and a set of prototypical representations of the output classes, offering therefore a deep, yet transparent-by-design, architecture. In this paper, we introduce a probabilistic reformulation of PSENNs, called Prob-PSENN, which replaces point estimates for the prototypes with probability distributions over their values. This provides not only a more flexible framework for an end-to-end learning of prototypes, but can also capture the explanatory uncertainty of the model, which is a missing feature in previous approaches. In addition, since the prototypes determine both the explanation and the prediction, Prob-PSENNs allow us to detect when the model is making uninformed or uncertain predictions, and to obtain valid explanations for them. Our experiments demonstrate that Prob-PSENNs provide more meaningful and robust explanations than their non-probabilistic counterparts, while remaining competitive in terms of predictive performance, thus enhancing the explainability and reliability of the models.

1 Introduction

The complexity in the architectures of state-of-the-art Deep Neural Networks (DNNs) largely accounts for their “black box” nature [34], which is in conflict with one of the basic requirements for trustworthiness [4]: to be able to explain and understand the decisions of the model. Although several strategies have been proposed in order to explain black box models [46], they often provide only partial information about their inner workings, such as which parts of the input at hand most condition the output [5, 35, 37–40]. However, these approaches do not provide information about how the model processes that information or why they imply the output [1, 17, 28, 34]. In order to achieve a more transparent classification process, recent works advocate the use of Prototype-Based Self-Explainable Neural Networks (PSENNs, for simplicity) [3, 7, 13, 17, 27], which are trained to jointly maximize their prediction performance and their explainability. These approaches rely primarily on prototype-based models, in which the prototypes are learned during the training phase, aiming to capture discriminative and also semantically-meaningful features. In this way, the output classification is based on the similarity between the input and the learned prototypes, providing an intuitive and human-understandable classification process.

On the other hand, current DNNs, including the PSENNs discussed above, are generally optimized with the aim of finding the point estimate values of the parameters that minimize a given loss term, leading to models with a deterministic inference process. An important drawback of these approaches is that, generally, it is not possible to measure to what extent the model is confident in its own prediction, known in the literature as the epistemic uncertainty of the model [12, 19, 24]. Indeed, it has been shown that DNNs tend to be overconfident in their predictions [16], classifying inputs with high probability even when random or out-of-distribution inputs are provided. Being unable to detect when the model is making such unreliable or uninformed decisions dramatically reduces the trustworthiness and reliability of the model, and is, therefore, of great concern in high-stakes applications, as recognized in the EU AI Act [11]. To address this issue, recent works proposed replacing point estimates for the parameters of the model with a probability distribution over their values [15]. This leads to a stochastic model, for which it is possible to measure the aforementioned types of uncertainty, and, thus, identify when the model is making uncertain predictions that should not be trusted.

In this paper, we aim to address the aforementioned drawbacks and challenges by leveraging recent tools from uncertainty estimation in DNNs, and combining them with PSENN architectures in order to develop *Probabilistic PSENNs (Prob-PSENNs)*. The key feature of this model is that the prototypes are defined as random variables, for which suitable probability distributions are learned during the training phase of the model – jointly with the rest of the parameters of the network. This probabilistic paradigm allows us not only to improve the explainability of the classification process, but also to capture different sources of uncertainty regarding both the output and the explanation. Furthermore, since the output of the PSENNs directly depends on the explanatory component [7, 13, 27], we can establish connections between their corresponding uncertainties, enabling a more thorough analysis of the prediction process, and allowing the models to self-explain their own uncertainties.

The main contributions of this work are summarized below:

1. We introduce Prob-PSENN, a probabilistic reformulation of PSENNs that enhances trustworthiness by addressing two critical aspects: (i) interpretability and transparency, and (ii) uncertainty quantification.
2. We demonstrate how the probabilistic reinterpretation of the prototypes substantially enhances the explanatory capabilities of the model, enabling more diverse, meaningful and robust explanations than those achievable with the non-probabilistic counterparts.
3. Prob-PSENNs allow us, for the first time, to model the uncertainty in both the explanations and predictions, which is a missing feature in conventional PSENNs. We formalize different explanatory uncertainty notions, and show how Prob-PSENNs can (i) quantify those uncertainties, (ii) explain them, and (iii) leverage them to improve their predictive performance.
4. The effectiveness of Prob-PSENNs is evaluated in several classification tasks, demonstrating their reliability and trustworthiness.

2 Preliminaries: Prototype-Based Self-Explainable Neural Networks (PSENNs)

Let us consider a classification problem in which the goal is to classify d -dimensional inputs $x \in \mathbb{R}^d$ into one of the c possible classes in $Y = \{y_1, y_2, \dots, y_c\}$. Building on the frameworks introduced in [13, 27], the core components of the PSENNs are described below. We refer the reader to Appendix A for a justification of the choice of these frameworks, as well as a discussion of related works.

The first component of these architectures is an encoder module $e : \mathbb{R}^d \rightarrow \mathbb{R}^l$, where l denotes the dimensionality of the *latent* space to which the input is encoded. The latent representation computed by the encoder $e(\cdot)$ for the input x will be denoted as $e_x = e(x)$. Conversely, let $g : \mathbb{R}^l \rightarrow \mathbb{R}^d$ denote a decoding module, capable of reconstructing the original input x from the corresponding encoded representation e_x . The key component of PSENNs involves the generation of m prototypes in the latent space: $R = (r_1, r_2, \dots, r_m)$, with $r_i \in \mathbb{R}^l$, $1 \leq i \leq m$. The aim of generating these prototypes is to capture diverse, discriminative and semantically meaningful features capable of representing the classes of the problem. In this way, the similarity between an encoded input e_x and each of the m prototypes can be employed to determine the output class for x in an easily interpretable manner,

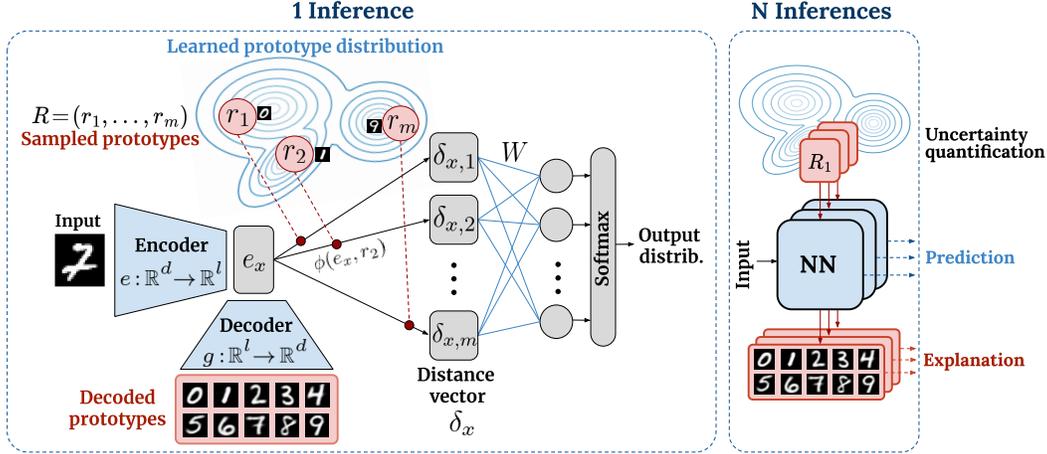


Figure 1: Architecture of the proposed Probabilistic PSENN (Prob-PSENN). The model replaces deterministic prototypes with probabilistic prototype distributions, enabling the sampling of diverse yet representative prototypes from the latent space. To account for prototype variability, N samples of the prototypes can be considered during the inference process, producing diverse predictions that enable capturing both explanatory and predictive uncertainties.

based on the assumption that an input x belonging to class y_i will be closer to those prototypes representing the class y_i than to the prototypes corresponding to the remaining classes. To formalize the aforementioned similarities, let the column-vector $\delta_x = (\delta_{x,1}, \delta_{x,2}, \dots, \delta_{x,m})^\top$ represent the distance between an encoded input e_x and the m prototypes in R . More specifically, the distance between e_x and the i -th prototype r_i will be computed as $\delta_{x,i} = \phi(e_x, r_i)$, $1 \leq i \leq m$, based on a distance metric $\phi: \mathbb{R}^l \times \mathbb{R}^l \rightarrow \mathbb{R}$. The distance vector δ_x will condition both the output classification and the corresponding explanation, as detailed below.

Classification Taking δ_x as the input, the output provided by the model will be determined by a prototype classifier $h: \mathbb{R}^m \rightarrow [0, 1]^c$. As in [27], a classification network $h(\delta_x) = s(W\delta_x)$ will be assumed, where W is a $c \times m$ weight matrix and $s(\cdot)$ the *softmax* function. For simplicity, the end-to-end classification process will be denoted as $f(x) = h(\delta_x)$. Notice that $f(x)$ outputs a vector representing the probability with which the input belongs to each of the possible classes of the problem, according to the model. The probability assigned to each class y_i will be denoted as $f_i(x)$, $1 \leq i \leq c$, and thus we can define the output as: $f(x) = (f_1(x), f_2(x), \dots, f_c(x))$.

Explanation Since the output class is based on the distance-vector δ_x , the explanation ξ_x for the prediction $f(x)$ can be determined also by the similarity between the prototypes and the input at hand: $\xi_x = \{\{g(r_1), \delta_{x,1}\}, \{g(r_2), \delta_{x,2}\}, \dots, \{g(r_m), \delta_{x,m}\}\}$, where $g(\cdot)$ is the decoder defined above. The interpretability of the model relies, therefore, on to what extent the decoded prototypes are capable of capturing relevant and semantically-meaningful features representing each class of the problem, which is enforced during the training phase by means of including interpretability regularizers in the loss function [13, 27].

3 Probabilistic Prototype-Based Self-Explainable Neural Networks (Prob-PSENNs)

As mentioned above, PSENNs treat the prototypes as network parameters, for which point estimate values are computed and hence are fixed once the training is completed. In this section, we introduce Prob-PSENNs, which generalize the prototypes by means of defining and learning a probability distribution over their values, as we illustrate in Figure 1. The main motivations for this probabilistic formulation are discussed below.

First, for each class of the problem, there might exist several sets of prototypes R capable of representing the data with comparable effectiveness. For illustration, let us consider a handwritten

digit classification scenario, and the fact that there are several ways of representing each digit. This implies that there are several representative prototypes for each class. Consequently, instead of relying on a single set of prototypes, it might be more reasonable to consider all those possible variations in order to classify an input. Furthermore, in previous approaches, the number of prototypes m needs to be finite and defined beforehand, even if, in practice, it might be cumbersome or infeasible to determine how many prototypes will be required to capture sufficiently diverse characterizations of the classes. In contrast, our probabilistic reformulation enables a more flexible way to decide the variety of the prototypes required to represent the classes, which will be determined, assuming one prototype distribution per class, by the variance of the distributions.

Finally, the choice of the prototypes in R conditions not only the output but also the explanation. Therefore, placing a probability distribution over the prototypes allows us to enable a novel probabilistic framework for the generation of the explanation as well, which can be employed to capture different sources of uncertainty in the explanations of the model, and thereby, produce more informative and useful explanations.

3.1 Predictive distribution

Instead of considering a point estimate for R , we will use a random variable, denoted \mathbf{R} . In this way, we can consider $m = c$ random class-prototypes, such that \mathbf{r}_i represents the class y_i , $1 \leq i \leq c$. This eliminates the need to estimate m and avoids having to restrict the model to a fixed number of prototypes per class, as each \mathbf{r}_i will model the required variability within the corresponding class. Due to the randomness in the network’s prototypes, the inference of the model will be stochastic, and will be denoted as $f^{\mathbf{R}}(x)$. In this way, a predictive distribution for the output can be defined:

$$p(\mathbf{y}|x) = \mathbb{E}_{p(\mathbf{R})}[p(\mathbf{y}|x, \mathbf{R})] = \int p(\mathbf{y}|x, \mathbf{R}) \cdot p(\mathbf{R}) d\mathbf{R} = \int f^{\mathbf{R}}(x) \cdot p(\mathbf{R}) d\mathbf{R}. \quad (1)$$

In practice, $p(\mathbf{y}|x)$ can be approximated by considering a finite number of N samples from $p(\mathbf{R})$:

$$\tilde{p}(\mathbf{y}|x) = \tilde{f}(x) = \frac{1}{N} \sum_{n=1}^N f^{R_n}(x), \quad \text{with } R_n \sim p(\mathbf{R}), 1 \leq n \leq N. \quad (2)$$

3.2 Capturing the predictive uncertainty

The predictive distribution defined in Equation (1) allows us to better model the *predictive uncertainty* of a Prob-PSENN compared to its deterministic counterpart. First, the total predictive uncertainty can be measured by means of the entropy [36] of the prediction: $H[p(\mathbf{y}|x)] = -\sum_{i=1}^c p(\mathbf{y} = y_i|x) \cdot \log p(\mathbf{y} = y_i|x)$. However, for a model with fixed (pointwise) values for the parameters, the source of this uncertainty will be only *aleatoric* uncertainty, which refers to the possibly unavoidable noise, randomness or ambiguity inherent in the data. Considering a probability distribution for the parameter values instead of point estimates allows us to also capture the epistemic uncertainty [19, 21, 24], which accounts for the model uncertainty about the true parameters modeling the data (e.g., as a consequence of insufficient training data). This enables a safer use of the model in critical tasks, since it allows us to identify when the model is making unreliable predictions. An illustrative comparison between aleatoric and epistemic uncertainty is provided in Figure 2 (left).

Thus, for the predictive distribution defined in Equation (1), $H[p(\mathbf{y}|x)]$ encapsulates both aleatoric and epistemic uncertainty. To measure epistemic uncertainty only, the mutual information between the output \mathbf{y} and the parameters \mathbf{R} can be employed [10, 12, 18, 19], which, in our case, can be formalized as: $I(\mathbf{y}, \mathbf{R}|x) = H[p(\mathbf{y}|x)] - \mathbb{E}_{p(\mathbf{R})}[H[p(\mathbf{y}|x, \mathbf{R})]]$, while $\mathbb{E}_{p(\mathbf{R})}[H[p(\mathbf{y}|x, \mathbf{R})]]$ represents a measure of aleatoric uncertainty. Notice that the values of all these metrics can be normalized to the range $[0, 1]$, where 1 represents the maximum uncertainty, by dividing the resulting value by the entropy of the uniform categorical distribution of c categories. In practice, $H[p(\mathbf{y}|x)]$ can be approximated by $H[\tilde{p}(\mathbf{y}|x)]$ (see Equation 2).

3.3 Capturing the explanation uncertainty

With a PSENN with fixed prototypes, only a point estimate δ_x can be provided for the input at hand, that is, a point estimate on how close the input is to the particular prototypes selected to

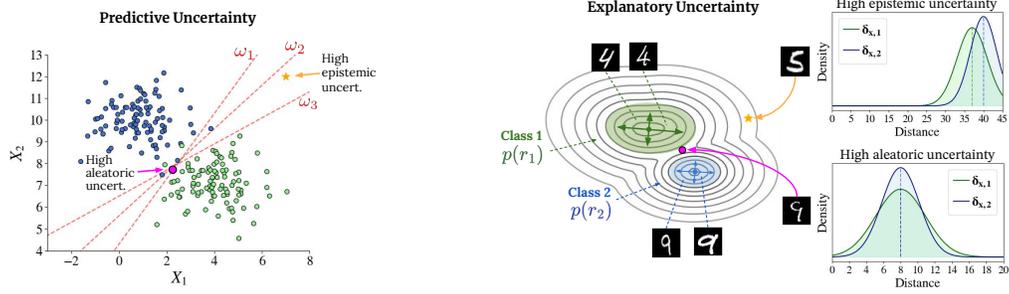


Figure 2: Comparison between the predictive and explanatory uncertainty, illustrated for a binary classification problem with a two-dimensional input space. In the left subfigure, dashed lines represent different decision boundaries. In the right subfigure, colored regions represent high density regions for the prototype distribution $p(r_1)$ and $p(r_2)$.

represent each class, as illustrated in Figure 3-(a). In contrast, the probability distribution over the prototypes considered in Prob-PSENNs induces a probability distribution over each value of δ_x , $p(\delta_{x,i}) = p(\phi(e_x, r_i))$, as exemplified in Figure 3-(b) and Figure 3-(c). This allows us to compute not only more representative information for each value of δ_x , but also to compute the uncertainty in the explanation, as will be detailed below. Furthermore, this enables a direct and explicit connection between the explanation uncertainty and the output uncertainty, which is a key advantage of Prob-PSENNs.

3.3.1 Types of explanation uncertainty

Similarly to the predictive uncertainty, different types of uncertainty can be attributed to the explanations of Prob-PSENNs. In the context of explanations, we will refer to *aleatoric* uncertainty as the uncertainty in determining which class-prototypes best represent the input at hand, that is, in determining which prototypes are closest to the encoded input, possibly due to ambiguity in the representation of the input. Therefore, high aleatoric uncertainty will imply highly overlapping distance-distributions $p(\delta_{x,i})$. In contrast, *epistemic* explanatory uncertainty will be high for those inputs that lie far from the training distribution, and, consequently, also far from the distribution over the prototypes, resulting in uninformative and non-representative explanations for that input. Therefore, high epistemic uncertainty will imply large distances with respect to all the prototypes. An illustrative representation of these types of explanatory uncertainty is provided in Figure 2 (right).

3.3.2 Quantifying explanation uncertainty

Given that, in the general case, computing $p(\delta_{x,i})$ will be intractable both analytically and numerically, we will rely on sample-based estimators in order to compute meaningful information about those distributions. To begin with, assuming $\delta_x^{(1)}, \dots, \delta_x^{(N)}$ are the distance-vectors obtained for N inferences of the model (i.e., for N random sets of prototypes), the sample mean $\bar{\delta}_{x,i} = \frac{1}{N} \sum_{n=1}^N \delta_{x,i}^{(n)}$ can inform us about the average similarity between the input and the prototypes of the class y_i , while the sample variance $\text{Var}(p(\delta_{x,i})) \approx \frac{1}{N-1} \sum_{n=1}^N (\delta_{x,i}^{(n)} - \bar{\delta}_{x,i})^2$ can be considered a measure of uncertainty about those similarities.

While the aforementioned two metrics describe each distribution $p(\delta_{x,i})$ separately, considering all these distributions jointly allows us to estimate aleatoric and epistemic uncertainty in the explanations. As mentioned earlier, aleatoric uncertainty can be determined based on the overlap between the probability density functions describing $p(\delta_{x,i})$, which will be denoted as $f_{\delta_{x,i}}(\cdot)$. For a binary case, the overlap can be measured by means of (normalized) overlapping indices [32]: $\int_{\mathbb{R}} \min[f_{\delta_{x,1}}(z), f_{\delta_{x,2}}(z)] dz$. In order to address multi-class problems, we will generalize the metric as the maximum overlap between the density function $f_{\delta_{x,i^*}}(\cdot)$ corresponding to the most likely class y_{i^*} and the functions corresponding to the remaining classes:

$$\mathcal{U}_A(x) = \max_{\substack{j=1,\dots,c \\ j \neq i^*}} \left[\int_{\mathbb{R}} \min[f_{\delta_{x,i^*}}(z), f_{\delta_{x,j}}(z)] dz \right]. \quad (3)$$

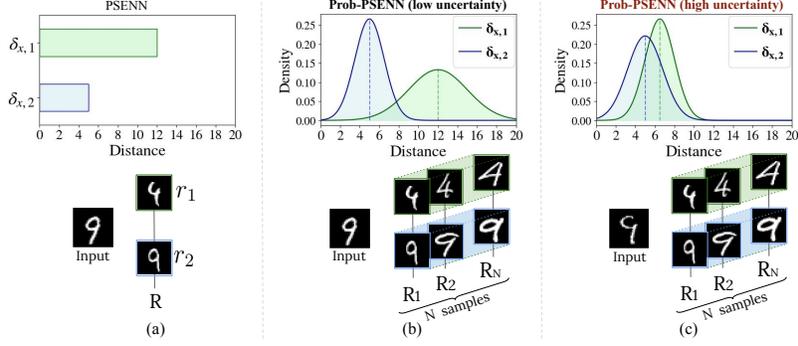


Figure 3: Illustrative comparison between a conventional PSENN (left column) and the proposed Probabilistic-PSENN (middle and right columns) regarding the computation of the distance-vector δ_x . Whereas a fixed set of prototypes R is used in (a), leading to a point estimate δ_x , placing a probability distribution over R , as in (b) and (c), induces a probability distribution over the distance $\delta_{x,i}$ between the encoded input and the random prototypes r_i representing the class y_i , $1 \leq i \leq c$.

In practice, the density functions can be estimated by means of kernel-density estimation methods, based on the distances obtained for a finite number of inferences N .

Regarding epistemic explanatory uncertainty, as explained above, high uncertainty will be achieved for those inputs for which the distances to all the prototypes are large. On that account, the minimum sample mean: $\bar{\delta}_{x,i*} = \min_{1 \leq j \leq c} \bar{\delta}_{x,j}$ represents a suitable base metric for evaluating this uncertainty. Furthermore, in order to determine the severity of the uncertainty in a normalized and scale-independent manner, a quantile-based metric will be defined, close to the approach proposed in [33]. This metric will be based on comparing the value at hand with the values obtained for the data distribution \mathcal{D} , which can be estimated, in practice, using a training set D :

$$\mathcal{U}_E(x) = P_{\substack{(x',y) \sim \mathcal{D} \\ y=y_{i*}}}(\bar{\delta}_{x',i*} \leq \bar{\delta}_{x,i*}). \quad (4)$$

3.4 Creating more meaningful explanations

In practice, we expect to sample N different explanations, and, thereby, N different yet representative prototypes for each class $y_i \in Y$ (assuming $m = c$). Although presenting all these prototypes as part of the explanation might be impractical, different strategies can be considered to summarize or distill the results, for the sake of more meaningful and useful explanations. First, among the prototypes corresponding to the most likely class, the ones closest to the input at hand can be considered the *best candidates* to justify the prediction of the model. Furthermore, the closest prototypes corresponding to the remaining classes can be taken as *counterarguments* against the most likely class, which can be used, for instance, to identify and explain why an input can be considered ambiguous or challenging to classify. On the other hand, the farther prototypes corresponding to the most likely class can be used to exemplify alternative yet still valid representations for that class, which can provide useful complementary information, and is thus valuable for *knowledge-acquisition* purposes.

3.5 Training procedure

Let $p_\lambda(\mathbf{R})$ represent a parametric probability distribution over the prototypes. The goal of the training procedure will be to optimize the parameters λ , jointly with the rest of parameters of the model, in order to maximize both the predictive performance and interpretability of the model. For simplicity, the prototypes $\mathbf{r}_1, \dots, \mathbf{r}_c$ will be assumed to be independent random variables, with each \mathbf{r}_i distributed according to a distribution $p_{\lambda_i}(\mathbf{r}_i)$, $1 \leq i \leq c$, and, hence, $p_\lambda(\mathbf{R}) = p_\lambda(\mathbf{r}_1, \dots, \mathbf{r}_c) = \prod_{i=1}^c p_{\lambda_i}(\mathbf{r}_i)$. The following generalized loss function will be employed to train the entire architecture:

$$\mathcal{L} = \tau_1 \cdot \mathcal{L}_{NLL} + \tau_2 \cdot \mathcal{L}_{REC} + \tau_3 \cdot \mathcal{L}_{INT}$$

where,

$$\mathcal{L}_{NLL} = -\frac{1}{N} \sum_{n=1}^N \sum_{(x,y) \in D} \sum_{i=1}^c \mathbb{1}(y=y_i) \cdot \log(f_i^{R_n}(x)), \quad R_n = [r_i]_{i=1}^c \text{ s.t. } r_i \sim p_{\lambda_i}(\mathbf{r}_i), \quad 1 \leq n \leq N,$$

$$\mathcal{L}_{\text{REC}} = \sum_{(x, \cdot) \in D} \|x - g(e(x))\|_2^2, \quad \mathcal{L}_{\text{INT}} = - \sum_{(x, y) \in D} \sum_{i=1}^c \mathbb{1}(y = y_i) \cdot \log p_{\lambda_i}(e(x)),$$

and $\tau_1, \tau_2, \tau_3 \in \mathbb{R}^+$ weight the contribution of each term. Notice that \mathcal{L}_{NLL} represents the classification (cross-entropy) loss, averaged for N samples of prototypes, \mathcal{L}_{REC} the reconstruction loss, and \mathcal{L}_{INT} an interpretability loss, as it will encourage the distribution over the prototypes to be similar to the distribution of the encoded data, which is a crucial requirement for interpretability [27]. Based on this methodology, the entire architecture can be jointly optimized by means of conventional gradient descent approaches.

4 Experiments

In this section, we describe the experimental setup designed to evaluate the effectiveness and capabilities of the proposed model, followed by a presentation and discussion of the corresponding results. The primary goal of these experiments is to demonstrate how Prob-PSENNs can: (i) learn highly diverse and prototypical representations of the classes, (ii) provide reliable uncertainty quantification for both the explanation and the classification, and (iii) leverage the uncertainty quantification mechanism to improve the accuracy, reliability and trustworthiness of the model.

4.1 Setup

The effectiveness and capabilities of the proposed model will be assessed on the MNIST [26], Fashion-MNIST [44], K-MNIST [8] and SVHN [31] datasets. This choice of datasets is based on the fact that they provide a clear and coherent notion of prototypes to facilitate the assessment of our results, particularly regarding explanation quality (e.g., prototype diversity) and uncertainty quantification (e.g., ambiguity).

The autoencoder module of Prob-PSENN will be implemented by a CNN-based Variational AutoEncoder (VAE), which adds an additional regularization term to the overall loss function, namely, the Kullback–Leibler (KL) divergence between the learned latent distribution and a prior (Gaussian) distribution [23]. The full implementation details are reported in Appendix B. For each prototype distribution $p_{\lambda_i}(\mathbf{r}_i)$, $1 \leq i \leq c$, a full covariance Gaussian distribution will be assumed, and, as in [27], the squared Euclidean distance will be used as the distance metric ϕ to measure the distance between the encoded input and the sampled prototypes (in the latent space). The final classification layer W will be set as $W = -I$, where I represents the identity matrix of dimension c , thus ensuring that the association between the prototype \mathbf{r}_i and the class y_i , $1 \leq i \leq c$, is explicit and unequivocal, further promoting transparency. The Adam optimization method [22] has been used to train the networks, using a learning rate of 0.001. Our code is publicly available at <https://github.com/vadel/Prob-PSENN>.

4.2 Learned latent spaces and distributions over the prototypes

First, for illustration, the results obtained for the case in which $l = 2$ (i.e., a two-dimensional latent space is considered) are shown in Figure 4, for the MNIST dataset. The left subfigure shows the learned latent space, visualizing the latent representations assigned to a random set of test data (dots), as well as the learned distribution over the prototypes $p_{\lambda_i}(\mathbf{r}_i)$, $1 \leq i \leq c$ (contour lines). The right subfigure represents 15 random sets of prototypes R sampled from the learnt distributions, and decoded by the decoder g . As can be seen, even for a very constrained (low-dimensional) latent space, the model is capable of capturing and decoding prototypical and varied representations for each class.

Figure 5 shows randomly sampled sets of prototypes for the four datasets considered, for Prob-PSENNs trained with a latent dimensionality of $l = 40$. As can be observed, the increase in the dimensionality of the latent space makes it possible to capture prototype distributions with a higher degree of complexity in their structure, and to obtain more detailed decodings. Furthermore, even for complex datasets such as K-MNIST (see Appendix E), the learned prototype distributions are able to capture representative, high-quality and diverse characterizations of the classes. More results are shown in Appendices C, D and E.

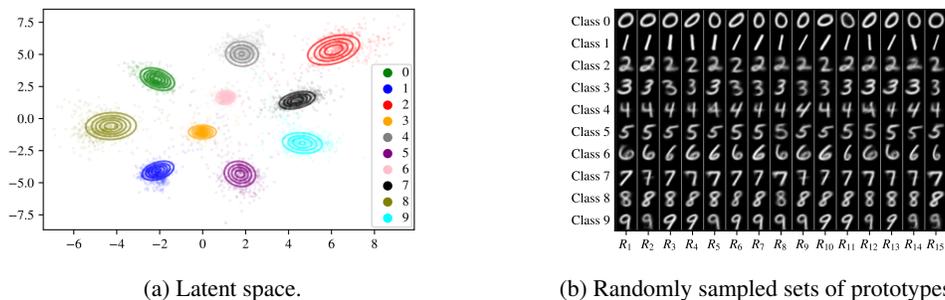


Figure 4: Illustrative results obtained with a Prob-PSENN with $l = 2$ for the MNIST dataset. (a) Learned latent space, where the contours represent the prototype distributions $p_{\lambda_i}(r_i)$ and the dots represent encoded input samples. (b) 15 random sets of prototypes $R_n = (r_1, \dots, r_c)$ s.t. $r_i \sim p_{\lambda_j}(r_j)$, $1 \leq i \leq c$, $1 \leq n \leq 15$.

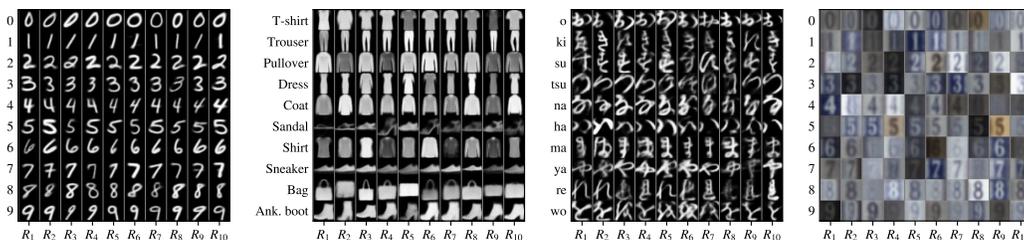


Figure 5: Visualization of prototypes sampled from the distributions learned with Prob-PSENNs ($l = 40$), for the MNIST, F-MNIST, K-MNIST and SVHN datasets.

4.3 Evaluating and illustrating the explanatory uncertainties

Figure 6 compares, for a Prob-PSENN with $l = 2$ and the MNIST dataset, the areas of the latent space with high explanatory uncertainty. The aleatoric and epistemic explanatory uncertainties (second and third columns) have been evaluated using the metrics $\mathcal{U}_A(x)$ and $\mathcal{U}_E(x)$ defined in Equations (3) and (4), respectively. The fourth column shows the areas solely with high epistemic uncertainty (i.e., with no aleatoric uncertainty), which has been computed as $\mathcal{U}_E(x) - \mathcal{U}_A(x)$. As can be seen from the results, regions with high aleatoric uncertainty correspond to areas in the latent space that lie between distinct class clusters. On the other hand, the areas with high epistemic uncertainty are those far from the regions to which the prototype distributions assign high density, and, therefore, the explanations are no longer representative of those regions.

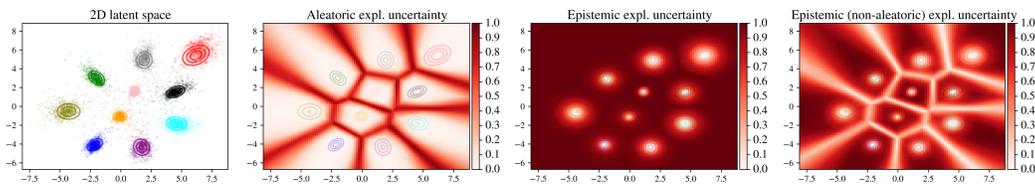


Figure 6: Evaluating the explanatory uncertainty in different regions of the latent space learned by Prob-PSENN.

Epistemic uncertainty Figure 7 compares the epistemic uncertainty achieved by a Prob-PSENN, trained for the MNIST dataset, when classifying 1000 test samples of the MNIST, F-MNIST and K-MNIST datasets. As can be seen, when inputs from the F-MNIST and K-MNIST dataset are presented to the model, a very high explanatory epistemic is captured for almost all the inputs, while for the MNIST test data the uncertainty is uniformly distributed. Furthermore, Figure 8 shows examples of MNIST inputs that achieve the lowest and highest uncertainties. As can be seen, even for

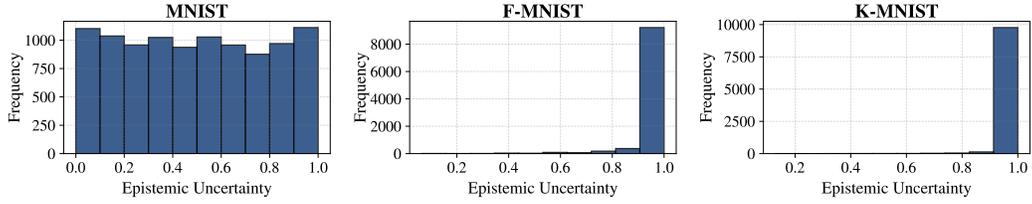


Figure 7: Histograms showing the explanatory epistemic uncertainty obtained by a Prob-PSENN (trained on the MNIST dataset) for 10000 test samples from MNIST, F-MNIST and K-MNIST. Notice that the scales of the Y-axis differ in the three plots.

the MNIST dataset, highly unusual inputs that could be considered out-of-distribution instances can be found, which are effectively detected by the uncertainty estimation capabilities of Prob-PSENN.

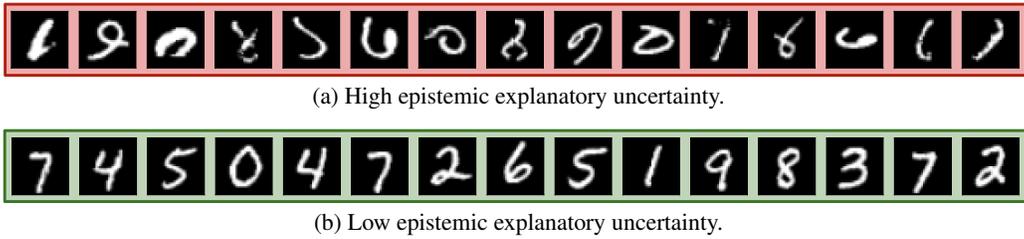


Figure 8: Examples of inputs receiving (a) high or (b) low explanatory epistemic uncertainty from a Prob-PSENN trained on the MNIST dataset.

Aleatoric uncertainty Finally, the prediction and explanation for a particular input with high aleatoric uncertainty is shown in Figure 9. The top row includes the density function over the distances corresponding to the two most likely classes, $\delta_{x,4}$ and $\delta_{x,9}$, estimated by Gaussian-kernel density estimation, as well as a box plot showing the variability in the output probabilities assigned to each class, considering all the inferences. Note that the distributions of both the distances and the output probabilities are highly overlapping, evidencing a high explanatory and predictive aleatoric uncertainty. The second row of the figure includes the input and a subset of the sampled prototypes, sorted by distance to the input in increasing order. As can be seen, both distributions over the prototypes capture and sample prototypical representations that could be taken as equally descriptive of the input at hand, thus explaining its ambiguity.

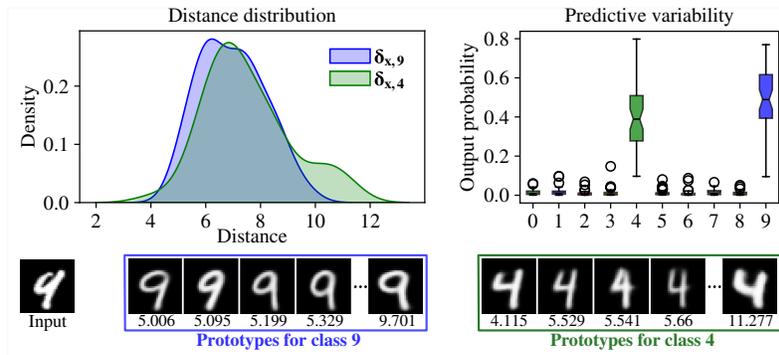


Figure 9: Prediction and explanation results obtained for a test input with high aleatoric uncertainty. The prototypes are sorted by distance to the input, in increasing order.

Table 1: Predictive accuracies (mean and standard deviation for 5 runs) obtained for the MNIST, Fashion-MNIST and K-MNIST datasets. The accuracies are shown in percentages. Best results in the default setting are in bold, while the best results in the uncertainty-aware setting (*Unc. Aware*; see Section 4.4.1) are in bold and italics.

| Model | Configuration | MNIST | F-MNIST | K-MNIST |
|---------------------|---------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| Li et al. [27] | $(m = 10, l = 40)$ | 99.0 ± 0.16 | 90.2 ± 0.33 | 91.1 ± 0.47 |
| Li et al. [27] | $(m = 15, l = 40)$ | 90.0 ± 0.11 | 90.1 ± 0.61 | 90.9 ± 0.41 |
| Li et al. [27] | $(m = 30, l = 40)$ | 99.1 ± 0.15 | 89.9 ± 0.33 | 91.0 ± 0.96 |
| Li et al. [27] | $(m = 50, l = 40)$ | 99.0 ± 0.11 | 89.8 ± 0.44 | 91.2 ± 0.74 |
| Gautam et al. [13] | $(m = 50, l = 256)$ | 99.4 ± 0.10 | 91.9 ± 0.20 | 96.1 ± 0.16 |
| Prob-PSENN | $(m = 10, l = 20)$ | 99.4 ± 0.05 | 91.9 ± 0.14 | 96.7 ± 0.21 |
| Prob-PSENN | $(m = 10, l = 40)$ | 99.4 ± 0.02 | 92.0 ± 0.24 | 96.6 ± 0.17 |
| + <i>Unc. Aware</i> | $(\alpha = 0.99)$ | 99.9 ± 0.02 | 95.6 ± 0.35 | 99.5 ± 0.03 |
| + <i>Unc. Aware</i> | $(\alpha = 0.95)$ | 99.9 ± 0.01 | 97.3 ± 0.21 | 99.7 ± 0.02 |

4.4 Comparison between conventional PSENN and Prob-PSENN

To analyze the predictive performance of Prob-PSENNs, we begin by comparing the accuracies obtained for the MNIST, F-MNIST and K-MNIST test sets, as shown in Table 1. For reference, we include the results of the deterministic baseline models proposed in Li et al. [27] and Gautam et al. [13], reproduced using the implementation released by the authors (see Appendix A for a further discussion on these works). The results show that Prob-PSENNs achieve either superior or competitive performance in all the datasets, while having as benefits greater robustness, trustworthiness, and enhanced explanatory capabilities. These capabilities are showcased in Figure 10, which compares, for the MNIST task, the explanation computed by a conventional PSENN ($m = 15$) with that computed by Prob-PSENN. More comparisons are provided in Appendices D and E. As can be seen, the prototypes sampled by Prob-PSENN exhibit considerably more diversity than the fixed prototypes learned by a conventional PSENN, which enables one to find more tailored prototypes for the input at hand, as well as to provide alternative forms that the class might take. In addition, while PSENNs only return point distances to the prototypes, Prob-PSENNs provide a distribution over those distances, enabling assessment of the uncertainty of the prediction, and hence, accounting for a more trustworthy approach.

4.4.1 Harnessing uncertainty to improve predictive accuracy and reliability

We experimentally validated that taking model uncertainty into consideration can provide substantial improvements in accuracy. Specifically, by discarding inputs with high explanatory epistemic

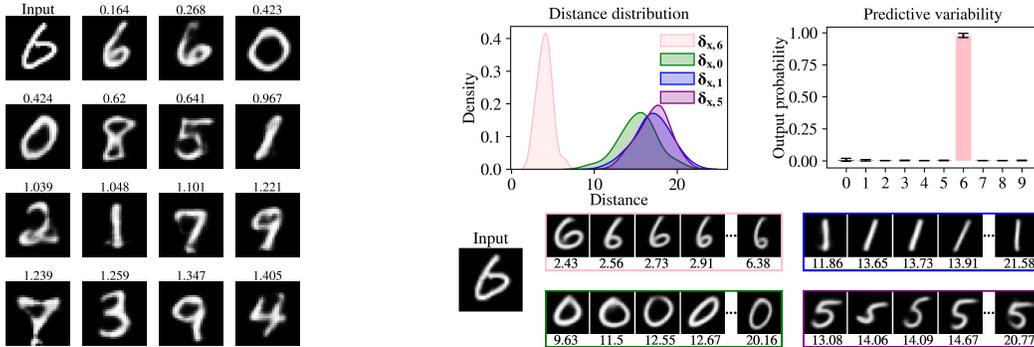


Figure 10: Comparison between the explanation provided by a PSENN [27] with fixed prototypes (left) and a Prob-PSENN (right). For simplicity, the results corresponding to the four most likely classes are visualized for Prob-PSENN. For each of these classes, we display the four prototypes closest to the input, as well as the most distant prototype.

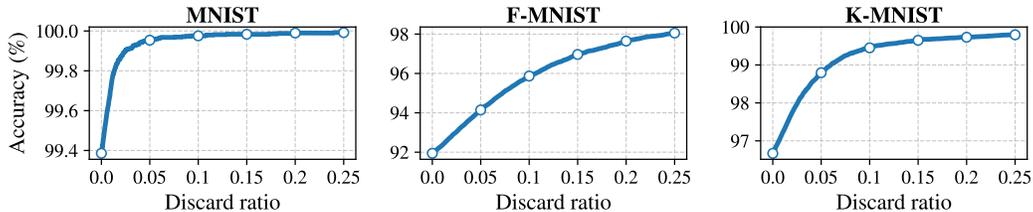


Figure 11: Trade-off between the accuracy gains and the discarded ratio, when the inputs with highest epistemic uncertainty are rejected. Results averaged for 5 runs. Notice that the scales of the Y-axis differ in the three plots.

uncertainty¹ (e.g., greater than $\alpha = 95\%$ or $\alpha = 99\%$), the accuracy of Prob-PSENN increases in all the datasets and configurations, as we show in Table 1 (last two rows). To complement these results, Table 2 shows the percentage of inputs that are discarded in each case. The trade-off between the accuracy and the ratio of discarded samples (i.e., those with highest epistemic uncertainty) is also shown in Figure 11. At the same time, a large number of the discarded inputs are anomalous or challenging to classify, which would justify discarding them for the sake of higher reliability. As evidence, we show in Table 2 that the error rate of the deterministic baseline model is dramatically higher in the inputs which receive high uncertainty than in the rest. For instance, the average error rate percentage of the model in the discarded inputs ($\alpha = 0.99$) is roughly 23%, 40% and 51% for MNIST, F-MNIST and K-MINST, respectively. Thus, this evaluation shows how capturing model uncertainty provides greater accuracy, reliability and safety.

4.4.2 Ablation studies

To isolate the impact of the probabilistic reformulation of the prototypes, Appendix F compares the performance of Prob-PSENN with its deterministic counterpart, using the same architecture and training pipeline. The point-wise prototypes have been optimized using the loss function proposed in [27]. The results obtained with deterministic PSENNs are summarized in Table F.1, computed considering $l \in \{20, 40\}$ latent dimensions and $m \in \{10, 30\}$ prototypes. The results for Prob-PSENN are included in the last rows of the table, as reference. As can be observed, both approaches achieve nearly identical performance, with average accuracy differences below 0.2%, further demonstrating that the advantages of our probabilistic model do not imply a reduction in its predictive performance.

In addition, Appendix F further examines the benefits of using a VAE over a vanilla AE as the backbone architecture. Table F.2 compares the predictive performance of Prob-PSENNs when replacing the VAE with an AE while keeping all other architectural and training details unchanged.

Table 2: Percentage of inputs discarded by Prob-PSENNs at different uncertainty thresholds (α) and the error rates (%) of the deterministic baseline model [27] for high-uncertainty (discarded) and low-uncertainty (retained) inputs across MNIST, F-MNIST, and K-MNIST.

| Threshold | Metric | MNIST | F-MNIST | K-MNIST |
|------------------|------------------------|------------------|------------------|------------------|
| $\alpha = 0.999$ | Discarded Inputs | 1.38 ± 0.10 | 2.23 ± 0.49 | 5.84 ± 0.21 |
| | Error Rate (High-Unc.) | 31.30 ± 3.22 | 45.49 ± 2.33 | 64.20 ± 2.61 |
| | Error Rate (Low-Unc.) | 0.79 ± 0.20 | 9.05 ± 0.57 | 6.28 ± 0.98 |
| $\alpha = 0.99$ | Discarded Inputs | 3.02 ± 0.03 | 9.19 ± 0.73 | 11.22 ± 0.23 |
| | Error Rate (High-Unc.) | 23.72 ± 2.86 | 40.01 ± 1.73 | 50.95 ± 2.73 |
| | Error Rate (Low-Unc.) | 0.51 ± 0.18 | 6.80 ± 0.46 | 4.44 ± 0.85 |
| $\alpha = 0.95$ | Discarded Inputs | 6.78 ± 0.08 | 17.52 ± 0.70 | 19.84 ± 0.28 |
| | Error Rate (High-Unc.) | 13.15 ± 1.68 | 33.28 ± 1.56 | 36.59 ± 2.75 |
| | Error Rate (Low-Unc.) | 0.35 ± 0.16 | 4.87 ± 0.34 | 3.00 ± 0.63 |

¹We justify the choice of the mentioned uncertainty metric based on the fact that it is reasonable to consider a prediction unreliable if all the prototypes have low similarity to the input, given the distance-based classification pipeline of the model.

The results show that AEs lead to a performance drop ranging from approximately 0.1% to 1.2%, depending on the dataset. In addition, Figure F.21 visualizes randomly sampled prototypes from Prob-PSENNs with different backbone architectures, showing that while vanilla AE backbones provide comparable results in lower latent dimensions (e.g., $l = 10$), they struggle to maintain high-quality prototype distributions as the dimensionality increases. In contrast, VAEs enable effective scaling to higher-dimensional spaces. This demonstrates that, while Prob-PSENNs are not inherently restricted to VAE backbones, more regularized latent spaces enhance both predictive performance and prototype quality in high-dimensional settings.

4.5 Scalability

In order to show that our approach can scale to larger numbers of classes, Appendix G includes results on the E-MNIST dataset [9], an extension of MNIST which contains handwritten representations of numbers and letters and is composed of 47 classes. As can be seen, Prob-PSENN achieves a highly diverse and representative sets of prototypes for all the classes, which would require an excessively large number of pointwise prototypes to achieve using deterministic approaches. Furthermore, to show that Prob-PSENNs can scale to colour datasets, we tested our model on the SVHN dataset. The accuracy of the model in this dataset is reported in Table 3. The results show that the predictive performance of Prob-PSENNs is on par with previous deterministic approaches, when not superior.

Table 3: Results on the SVHN dataset.

| Configuration | Accuracy (%) |
|--------------------------------|----------------|
| Gautam et al. [13] (ProtoPNet) | 88.6 ± 0.3 |
| Gautam et al. [13] (ProtoVAE) | 92.2 ± 0.3 |
| Gautam et al. [13] (SENN) | 91.5 ± 0.4 |
| Prob-PSENN | 92.1 ± 0.2 |

5 Conclusions

In this paper we introduced Prob-PSENN, a novel Prototype-Based Self-Explainable Neural Network paradigm relying on a probabilistic generalization of the prototypes, which enhances the capabilities of the model in two key aspects. First, unlike previous approaches, our model is not restricted to a fixed number or set of prototypes. This enables Prob-PSENN to learn much more diverse prototypes, and to consider all of them when classifying an input, which represents a more flexible and robust approach. Second, our probabilistic approach equips Prob-PSENNs with tools to capture the uncertainty in both the prediction and the explanation of the network. This novel feature makes it possible to detect when the model is providing uncertain outcomes, substantially increasing the reliability of the model. We experimentally demonstrated how Prob-PSENNs provide more informative and trustworthy explanations than their non-probabilistic counterparts, while achieving superior or competitive predictive accuracies. Moreover, we show that the accuracy of Prob-PSENNs can be improved by leveraging their novel uncertainty-quantification capabilities, further enhancing their reliability.

6 Limitations and future work

As future work, inspired by recent advances in Bayesian Neural Networks [12, 15, 30], we plan to develop a Bayesian formulation for the distribution over the prototypes, which might enable a more theoretically-grounded approach for uncertainty quantification. Furthermore, we also believe that a fully-probabilistic model (i.e., placing a probability distribution on the rest of the network parameters as well) could enhance the quantification of the uncertainty in Prob-PSENNs, or even induce new notions of uncertainty, such as the *encoding* or *class-representation uncertainty*. Apart from that, while Prob-PSENNs have demonstrated high predictive and explanatory performance, we found it challenging to model probability distributions over the prototypes in very high-dimensional spaces, and, therefore, resort to relatively low dimensional latent spaces in our experiments. We plan to address this limitation in the future, in order to scale the proposed approach to higher-dimensional datasets.

Acknowledgements

MK received partial support from ELSA: European Lighthouse on Secure and Safe AI project (Grant No. 101070617 under UK guarantee) and the ERC under the European Union’s Horizon 2020 research and innovation program (FUN2MODEL, Grant No. 834115). JV, RS and JAL acknowledge support from the Spanish Ministry of Science, Innovation and Universities (Project PID2022-137442NB-I00 and BCAM Severo Ochoa Accreditation CEX2021-001142-S), and the Basque Government (Grant Nos. KK-2024/00030 and IT1504-22).

References

- [1] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, and B. Kim. Sanity Checks for Saliency Maps. In *Advances in Neural Information Processing Systems*, volume 31, pages 9505–9515, 2018.
- [2] U. Aivodji, H. Arai, O. Fortineau, S. Gambs, S. Hara, and A. Tapp. Fairwashing: The Risk of Rationalization. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 161–170, 2019.
- [3] D. Alvarez-Melis and T. Jaakkola. Towards Robust Interpretability with Self-Explaining Neural Networks. In *Advances in Neural Information Processing Systems*, volume 31, pages 7775–7784, 2018.
- [4] M. Ashoori and J. D. Weisz. In AI We Trust? Factors That Influence Trustworthiness of AI-infused Decision-Making Processes. *arXiv preprint arXiv:1912.02675*, 2019. doi: 10.48550/arXiv.1912.02675.
- [5] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. *PLOS ONE*, 10(7):e0130140, 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0130140.
- [6] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli. Evasion Attacks against Machine Learning at Test Time. In *Machine Learning and Knowledge Discovery in Databases*, Lecture Notes in Computer Science, pages 387–402, 2013. ISBN 978-3-642-40994-3. doi: 10.1007/978-3-642-40994-3_25.
- [7] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin, and J. K. Su. This Looks Like That: Deep Learning for Interpretable Image Recognition. In *Advances in Neural Information Processing Systems*, volume 32, pages 8930–8941, 2019.
- [8] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep Learning for Classical Japanese Literature. In *NeurIPS 2018 Workshop on Machine Learning for Creativity and Design*, 2018.
- [9] G. Cohen, S. Afshar, J. Tapson, and A. van Schaik. EMNIST: Extending MNIST to handwritten letters. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 2921–2926, 2017. doi: 10.1109/IJCNN.2017.7966217.
- [10] S. Depeweg, J.-M. Hernandez-Lobato, F. Doshi-Velez, and S. Udfluft. Decomposition of Uncertainty in Bayesian Deep Learning for Efficient and Risk-sensitive Learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 1184–1193, 2018.
- [11] European Union (EU). Artificial Intelligence Act, 2024. URL <https://artificialintelligenceact.eu/the-act/>.
- [12] Y. Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.
- [13] S. Gautam, A. Boubekki, S. Hansen, S. Salahuddin, R. Jenssen, M. Höhne, and M. Kampffmeyer. ProtoVAE: A Trustworthy Self-Explainable Prototypical Variational Model. In *Advances in Neural Information Processing Systems*, volume 35, pages 17940–17952, 2022.

- [14] A. Ghorbani, A. Abid, and J. Zou. Interpretation of Neural Networks Is Fragile. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3681–3688, 2019. doi: 10.1609/aaai.v33i01.33013681.
- [15] E. Goan and C. Fookes. Bayesian Neural Networks: An Introduction and Survey. *Case Studies in Applied Bayesian Data Science*, 2259:45–87, 2020. doi: 10.1007/978-3-030-42553-1_3.
- [16] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On Calibration of Modern Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70, pages 1321–1330, 2017.
- [17] P. Hase, C. Chen, O. Li, and C. Rudin. Interpretable Image Recognition with Hierarchical Prototypes. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 32–40, 2019. ISBN 978-1-57735-820-6.
- [18] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel. Bayesian Active Learning for Classification and Preference Learning. *arXiv preprint arXiv:1112.5745*, 2011. doi: 10.48550/arXiv.1112.5745.
- [19] E. Hüllermeier and W. Waegeman. Aleatoric and Epistemic Uncertainty in Machine Learning: An Introduction to Concepts and Methods. *Machine Learning*, 110(3):457–506, 2021. ISSN 1573-0565. doi: 10.1007/s10994-021-05946-3.
- [20] H. Joo, J. Kim, H. Han, and J. Lee. Distributional Prototypical Methods for Reliable Explanation Space Construction. *IEEE Access*, 11:34821–34834, 2023. ISSN 2169-3536. doi: 10.1109/ACCESS.2023.3264794.
- [21] A. Kendall and Y. Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *Advances in Neural Information Processing Systems*, volume 30, pages 5574–5584, 2017.
- [22] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2015. doi: 10.48550/arXiv.1412.6980.
- [23] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*, 2014.
- [24] Y. Kwon, J.-H. Won, B. J. Kim, and M. C. Paik. Uncertainty Quantification Using Bayesian Neural Networks in Classification: Application to Ischemic Stroke Lesion Segmentation. In *Medical Imaging with Deep Learning (MIDL)*, 2022.
- [25] H. Lakkaraju and O. Bastani. "How do I fool you?": Manipulating User Trust via Misleading Black Box Explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, pages 79–85, 2020. ISBN 978-1-4503-7110-0. doi: 10.1145/3375627.3375833.
- [26] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. ISSN 1558-2256. doi: 10.1109/5.726791.
- [27] O. Li, H. Liu, C. Chen, and C. Rudin. Deep Learning for Case-Based Reasoning Through Prototypes: A Neural Network That Explains Its Predictions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, pages 3530–3537, 2018.
- [28] Z. C. Lipton. The Mythos of Model Interpretability: In Machine Learning, the Concept of Interpretability Is Both Important and Slippery. *Queue*, 16(3):31–57, 2018. ISSN 1542-7730. doi: 10.1145/3236386.3241340.
- [29] M. Nauta, R. Van Bree, and C. Seifert. Neural Prototype Trees for Interpretable Fine-grained Image Recognition. In *Proceedings of the 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14933–14943, 2021. ISBN 978-1-66544-509-2. doi: 10.1109/CVPR46437.2021.01469.

- [30] R. M. Neal. *Bayesian Learning for Neural Networks*, volume 118. Springer Science & Business Media, New York, NY, 2012.
- [31] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng. Reading Digits in Natural Images with Unsupervised Feature Learning. In *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [32] M. Pastore and A. Calcagni. Measuring Distribution Similarities Between Samples: A Distribution-Free Overlapping Index. *Frontiers in Psychology*, 10(1089), 2019. ISSN 1664-1078. doi: 10.3389/fpsyg.2019.01089.
- [33] A. Petrov and M. Kwiatkowska. Robustness of Unsupervised Representation Learning without Labels. *arXiv preprint arXiv:2210.04076*, 2022. doi: 10.48550/arXiv.2210.04076.
- [34] C. Rudin. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence*, 1(5):206–215, 2019. ISSN 2522-5839. doi: 10.1038/s42256-019-0048-x.
- [35] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017. doi: 10.1109/ICCV.2017.74.
- [36] C. E. Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27(3):379–423, 1948. ISSN 0005-8580. doi: 10.1002/j.1538-7305.1948.tb01338.x.
- [37] A. Shrikumar, P. Greenside, and A. Kundaje. Learning Important Features Through Propagating Activation Differences. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70, pages 3145–3153, 2017.
- [38] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. In *Workshop of the 2014 International Conference on Learning Representations (ICLR)*, 2014.
- [39] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for Simplicity: The All Convolutional Net. In *Workshop of the 2015 International Conference on Learning Representations (ICLR)*, 2015. doi: 10.48550/arXiv.1412.6806.
- [40] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic Attribution for Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70, pages 3319–3328, 2017.
- [41] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [42] J. Vadillo, R. Santana, and J. A. Lozano. Adversarial Attacks in Explainable Machine Learning: A Survey of Threats Against Models and Humans. *WIREs Data Mining and Knowledge Discovery*, 15(1):e1567, 2025. ISSN 1942-4795. doi: 10.1002/widm.1567.
- [43] C. Wang, Y. Chen, F. Liu, D. J. McCarthy, H. Frazer, and G. Carneiro. Mixture of Gaussian-distributed Prototypes with Generative Modelling for Interpretable Image Classification. *arXiv preprint arXiv:2312.00092*, 2023. doi: 10.48550/arXiv.2312.00092.
- [44] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint arXiv:2107.01943*, 2017. doi: 10.48550/arXiv.1708.07747.
- [45] X. Yuan, P. He, Q. Zhu, and X. Li. Adversarial Examples: Attacks and Defenses for Deep Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 30(9):2805–2824, 2019. ISSN 2162-2388. doi: 10.1109/TNNLS.2018.2886017.
- [46] Y. Zhang, P. Tiño, A. Leonardis, and K. Tang. A Survey on Neural Network Interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 5(5):726–742, 2021. ISSN 2471-285X. doi: 10.1109/TETCI.2021.3100641.

Appendix A Discussion on previous PSENN approaches

As introduced in Section 1, the development of PSENNs arises as an alternative to black-box DNNs and posthoc explanations, with the goal of ensuring a human-understandable decision process by construction, while maintaining competitive predictive performance. This is achieved by integrating a transparent case-based reasoning as the core element, so that the prediction of the model is determined based on the similarity between the input and a set of prototypical representations of the output classes (see Section 2 for further details).

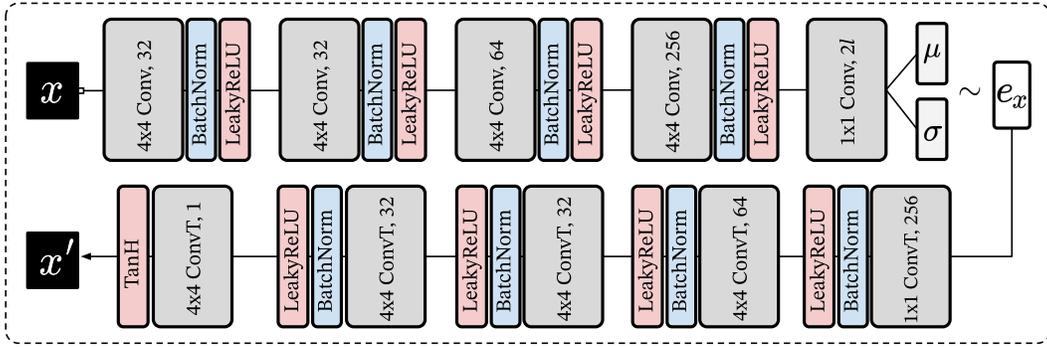
The way in which such prototypes are defined has led to different types of PSENNs in the literature. In the seminal work of Li et al. [27], mirrored by recent follow-up works [13, 20], the prototypes represent complete characterizations of each class or concept (e.g., full representations of the digits in handwritten digit recognition). In contrast, other follow-up works aim to capture more abstract [3] or specific features (e.g., prototypical parts or patches) [7, 17, 29] as prototypes. However, an important drawback of these approaches is that the prototypes are restricted to being samples or sample parts of (encoded) training inputs [3, 7, 17, 29]. This implies a flexibility loss regarding the end-to-end optimization of the prototypes, as well as a transparency loss when *surrogate prototypes* (e.g., nearest training images) are selected to visualize the explanations [3, 13]. For these reasons, in order to guarantee a fully transparent and end-to-end trainable architecture, we followed the same spirit as in [13, 27], and used the corresponding approaches as references to build our probabilistic framework.

Closer to our approach, concurrent works [20, 43] leverage the use of distributions over the prototypes to improve the consistency between image similarity and latent space location [20] and to increase the representation power of the model [43]. Contrary to our work, both approaches employ only the means of the learned distributions as prototypes, whereas our sampling-based inference process allows us to sample multiple sets of prototypes, avoiding the constraint of committing to a fixed set, and harnessing all of them for both the explanation and classification.

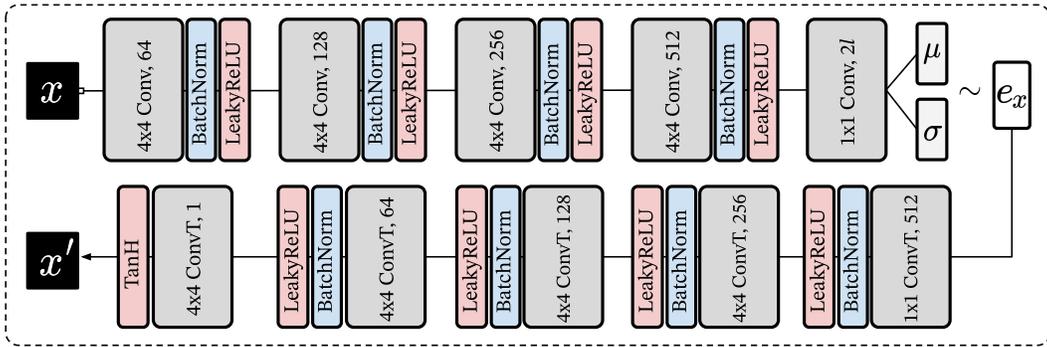
Furthermore, the approach in [43] follows the strategy of [3, 7, 17] and relies on surrogate prototypes, replacing the distribution means with nearest training instances, with the corresponding flexibility and transparency loss discussed above. Another key difference between our approach and that in [20] is that, in their work, the similarities between the input’s distributional embedding and the prototype distributions, which fed the final classification layer, are measured by considering the entire distributions, using distributional similarity metrics. However, as stated in [20], such similarities neither form nor rely on observable prototypes, which is contrary to the purpose of explaining and justifying the classification. Instead, we advocate for a more transparent approach, in which the model predictions are solely derived from the distances to the sampled sets of prototypes, ensuring a complete and transparent justification of model outputs.

Appendix B Implementation details

The Prob-PSENN autoencoder module has been implemented as a CNN-based Variational Autoencoder, consisting of four convolutional layers with 4×4 kernels, batch normalization, and LeakyReLU activations, followed by a 1×1 convolutional layer that controls the dimensionality of the latent space. The decoder mirrors this architecture by means of deconvolutional layers. Further details on the architecture and the corresponding hyperparameters are provided in Figure B.12. The inputs values have been normalized to the range $[-1, 1]$, and, for grayscale datasets, the input shapes have been upscaled from 28×28 to 32×32 using bilinear interpolation. For the SVHN dataset, random brightness adjustments and zoom transformations have been incorporated during training to improve generalization. The experiments have been run for $\tau_1 = 1.0, \tau_2 = 10.0, \tau_3 = 0.05, \tau_4 = 0.05$ and a batch size of 128. The models have been trained for 100 epochs on grayscale datasets and 30 epochs on the SVHN dataset. The Adam optimization method [22] has been used to train the networks, using a learning rate of 0.001. The final classification layer W has been set as $W = -I$, where I represents the identity matrix of dimension c , thus ensuring that the association between the prototype r_i and the class $y_i, 1 \leq i \leq c$, is explicit and unequivocal. In order to ensure higher stability during the training process, the value of \mathcal{L}_{INT} has been cut to the range $[-1, 10^3]$. For all datasets, standard training and test partitions have been used, with 10% of the training data reserved for validation. The results of [13, 27] have been reproduced using the implementations released by the authors. Our code is publicly available at <https://github.com/vadel/Prob-PSENN>.



(a) MNIST, F-MNIST, K-MNIST



(b) SVHN

Figure B.12: VAE architectures used in Prob-PSENNs for (a) MNIST, F-MNIST, K-MNIST, and (b) SVHN.

Appendix C Complementary results on the MNIST dataset

In order to complement the results reported in Section 4, Figure C.13 shows, for the MNIST dataset, 15 randomly sampled sets of prototypes from a Prob-PSENN with $l = 5$ (left) and $l = 10$ (right). As can be observed, the increase in the dimensionality of the latent space makes it possible to capture prototype distributions with a higher degree of complexity in their structure, as well as to obtain more detailed decodings.

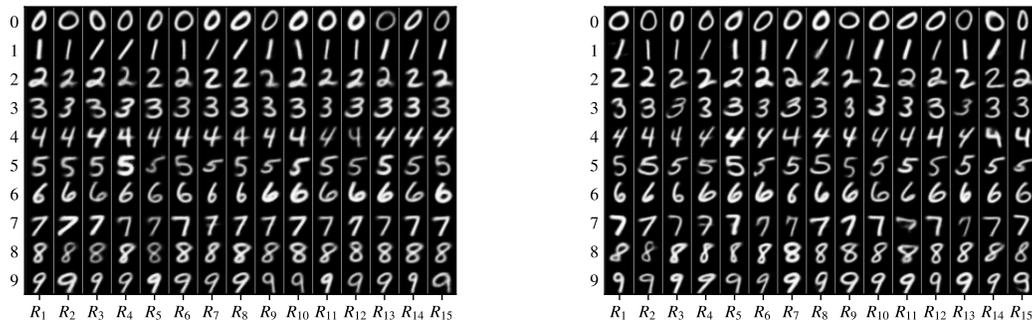


Figure C.13: 15 random sets of prototypes, obtained with a Prob-PSENN with $l = 5$ (left) and $l = 40$ (right).

Appendix D Results on the Fashion-MNIST dataset

This section includes complementary results on the Fashion-MNIST dataset. First, Figure D.14 shows randomly sampled prototypes for a Prob-PSENN with $l = 5$ (left) and $l = 40$ (right). These figures demonstrate that the probabilistic redefinition of the prototypes introduced in Prob-PSENNs enables the model to learn, even for relatively low-dimensional latent spaces, diverse and prototypical representations of each class.

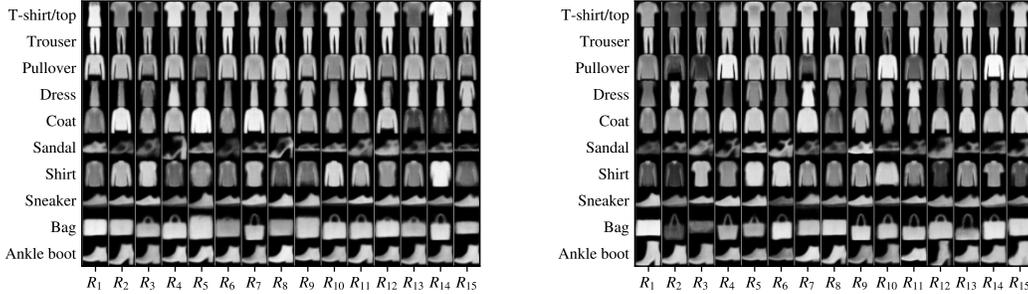


Figure D.14: Sampled prototypes for the Fashion-MNIST dataset, using a Prob-PSENN with $l = 5$ (left) and $l = 40$ (right).

Furthermore, a visual comparison between the explanations provided by a non-probabilistic PSENN and a Prob-PSENN is provided in Figures D.15 and D.16. For the sake of completeness, the former figure includes the explanations provided by a PSENN with $m = 15$ (left) and $m = 50$ (right) prototypes. In both cases, the original input is shown in the top-left corner, and the prototypes are colored according to the classes returned by the model when the prototypes are classified, in order to make it clearer which class each prototype is most associated with. Notice that this connection is explicit and unequivocal by construction in our Prob-PSENNs, being the class y_i represented by the prototype r_i , $1 \leq i \leq c$, resulting in more transparent decision processes and easier-to-interpret explanations. As can be observed in Figures D.15 and D.16, the prototypes learned by the PSENNs are, even for $m = 50$ prototypes, considerably less meaningful and varied than the ones captured by Prob-PSENN. Furthermore, the uncertainty metrics introduced in Section 3.3 enable us to quantify a low aleatoric and epistemic explanatory uncertainty for this input (0.004 and 0.16, respectively), which increases the trustworthiness and reliability of the outcomes of Prob-PSENN.

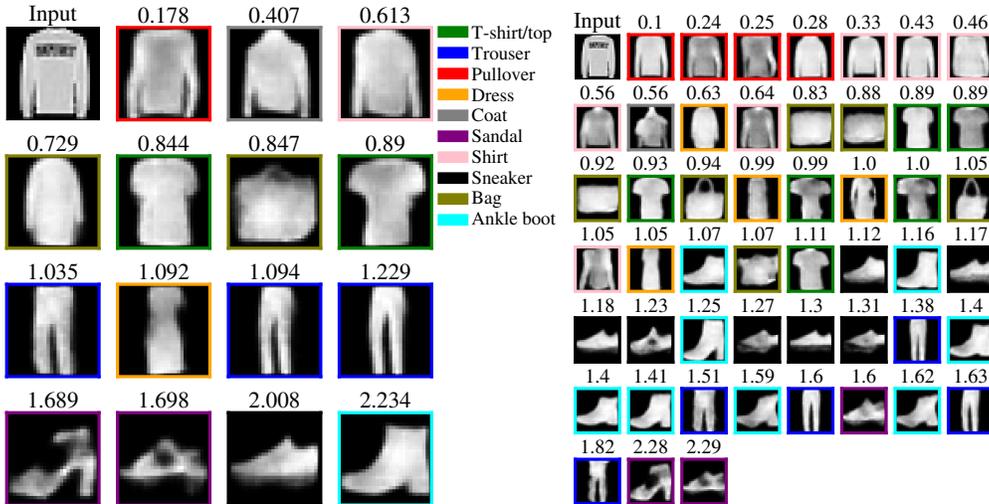


Figure D.15: Explanations provided by a PSENN [27] with $m = 15$ (left) and $m = 50$ (right) for the Fashion-MNIST dataset. Both models classify the input (shown in the top-left corner of each grid) as the class *pullover* with a probability of 0.99. Note that the prototypes are sorted based on their distances to the input (displayed on top of each figure), in increasing order.

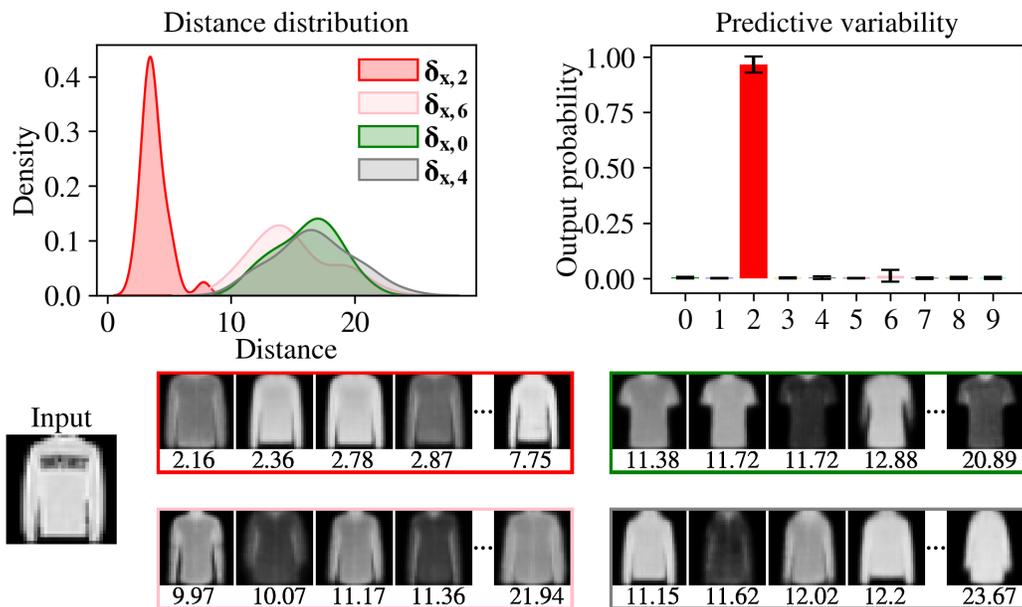


Figure D.16: Explanation provided by a Prob-PSENN, for the Fashion-MNIST dataset. The top row shows the distance distribution corresponding to the four most likely classes (left), and the variability in the output probabilities (right). The bottom part displays the prototypes for the four most likely classes, along with their corresponding distances to the input. Note that the prototypes are ordered by their distance, in ascending order.

Appendix E Results on the K-MNIST dataset

This section contains additional results on the 10-class K-MNIST dataset [8], which includes handwritten representations of Japanese letters. Random images from this dataset are shown in Figure E.17, from which it can be noticed a high variance in the representation of each letter. Randomly sampled prototypes from Prob-PSENNs with $l = 5$ and $l = 10$ are displayed in Figure E.18. As can be assessed in the figure, despite the complexity and high variability of this dataset, Prob-PSENNs are capable of capturing the most frequent representations of each class, ensuring a sufficiently diverse and meaningful characterization of the classes.

As for the previous datasets, a comparison between the explanations of PSENNs and Prob-PSENN for the K-MNIST dataset is provided in Figures E.19 and E.20. Notice how, also due to the higher variability in the representations of each class, the prototypes learned by a PSENN with $m = 15$ prototypes (Figure E.19-left) are not sufficiently representative of the classes, and, consequently, we need to resort to a large number of prototypes ($m = 50$, Figure E.19-right). However, even for $m = 50$, the prototypes learned by PSENN are considerably less varied and informative than the ones sampled by Prob-PSENN, which evidences the higher flexibility of our approach, as well as its increased transparency and explanatory capabilities.

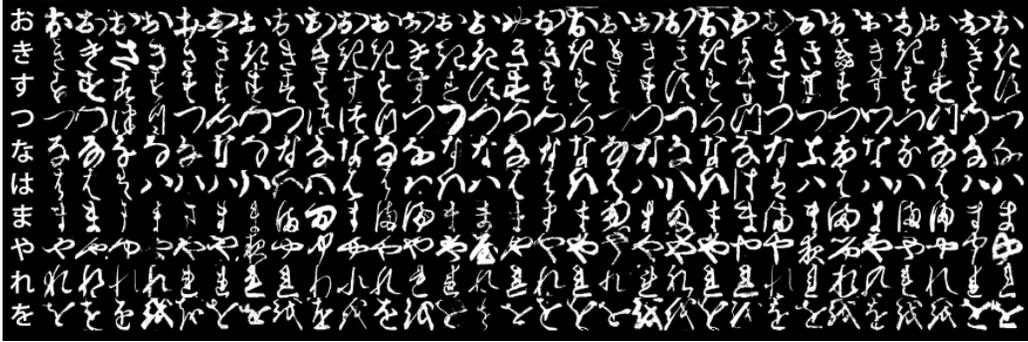


Figure E.17: Examples of the K-MNIST dataset [8], as reported in the original repository <https://github.com/rois-codh/kmnist> (CC BY-SA 4.0 license). Each row corresponds to one class. The first column of the figure shows the modern representation of each letter, as a reference.

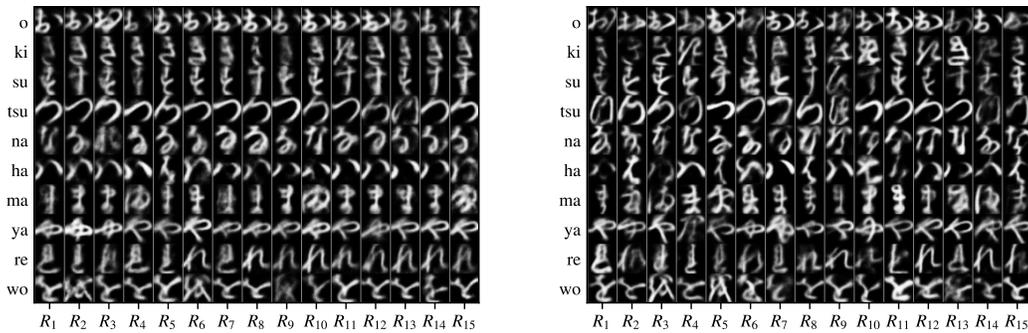


Figure E.18: Sampled prototypes for the K-MNIST dataset, using a Prob-PSENN with $l = 5$ (left) and $l = 40$ (right).

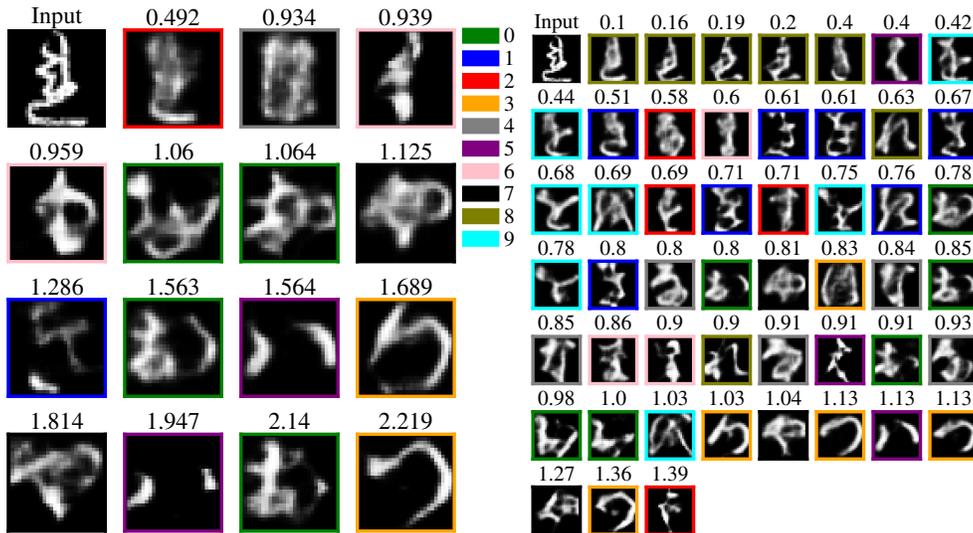


Figure E.19: Explanations provided by a PSENN [27] with $m = 15$ (left) and $m = 50$ (right) for the K-MNIST dataset. Both models classify the input (shown in the top-left corner of each grid) as the class 8 with a probability of 0.99. Note that the prototypes are sorted based on their distances to the input (displayed on top of each figure), in increasing order.

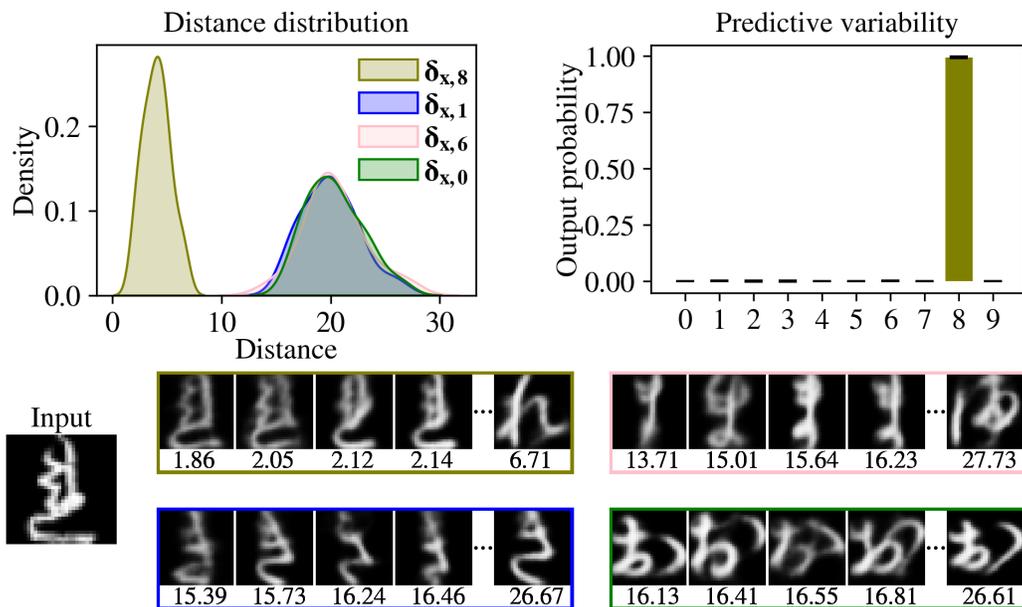


Figure E.20: Explanation provided by a Prob-PSENN, for the K-MNIST dataset. The top row shows the distance distribution corresponding to the four most likely classes (left), and the variability in the output probabilities (right). The bottom part displays the prototypes for the four most likely classes, along with their corresponding distances to the input. Note that the prototypes are ordered by their distance, in ascending order.

Appendix F Ablation studies: complementary results

This Appendix complements the results from Section 4.4 by comparing the performance of Prob-PSENN with its deterministic counterpart, using the same architecture and training pipeline. The point-wise prototypes have been optimized using the loss function proposed in [27]. The results obtained with the deterministic PSENN are summarized in Table F.1, computed considering $l = 20$ and $l = 40$ dimensions for the latent space and various numbers of prototypes m . The results for Prob-PSENN are included in the last rows of the table (for $l = 20$ and $l = 40$), as reference. As can be observed, the performance of both approaches is nearly identical, with differences in average accuracy below 0.2%. These results further demonstrate that the advantages of our probabilistic model do not imply a reduction in its predictive performance, while bringing as benefits enhanced trustworthiness, uncertainty quantification and explanatory capabilities.

Table F.1: Predictive accuracy percentages (mean and standard deviation for 5 runs) for deterministic PSENNs using the same architecture and training pipeline as Prob-PSENNs. The **best results for each approach** (PSENN and Prob-PSENN) are highlighted in bold, in order to facilitate comparison and emphasize their similar performance.

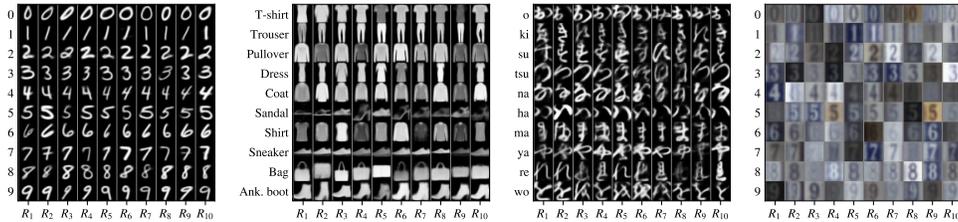
| Model | Configuration | MNIST | F-MNIST | K-MNIST |
|------------|----------------|------------------------------------|------------------------------------|------------------------------------|
| PSENN | $(m=10, l=20)$ | 99.39 \pm 0.04 | 92.11 \pm 0.10 | 96.71 \pm 0.12 |
| PSENN | $(m=10, l=40)$ | 99.42 \pm 0.02 | 91.94 \pm 0.25 | 96.79 \pm 0.08 |
| PSENN | $(m=30, l=20)$ | 99.33 \pm 0.10 | 90.92 \pm 0.14 | 96.42 \pm 0.18 |
| PSENN | $(m=30, l=40)$ | 99.29 \pm 0.08 | 90.30 \pm 0.51 | 96.72 \pm 0.11 |
| Prob-PSENN | $(m=10, l=20)$ | 99.42 \pm 0.05 | 91.91 \pm 0.14 | 96.72 \pm 0.21 |
| Prob-PSENN | $(m=10, l=40)$ | 99.38 \pm 0.02 | 91.98 \pm 0.24 | 96.61 \pm 0.17 |

Next, in Table F.2 we report the results obtained with Prob-PSENNs for which VAEs have been replaced with vanilla AutoEncoders, matching all the remaining architectural details and the training pipeline. The results show that vanilla AutoEncoders lead to a performance drop ranging from

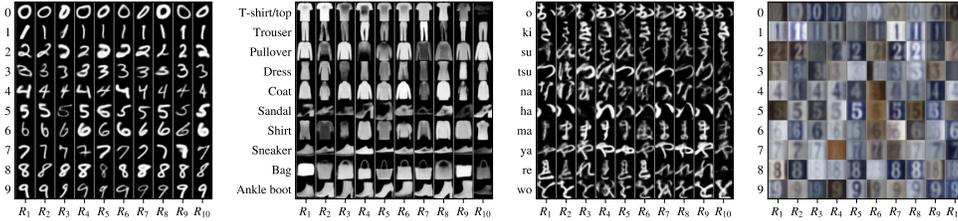
approximately 0.1% to 1.2%, depending on the dataset. In addition, Figure F.21 visualizes different sets of randomly sampled prototypes from Prob-PSENNs with different backbone architectures, highlighting the advantage of regularized latent spaces for improving prototype quality in high-dimensional settings.

Table F.2: Comparison between the accuracy of Prob-PSENNs using vanilla AutoEncoder (AE) and Variational AutoEncoder (VAE) backbones.

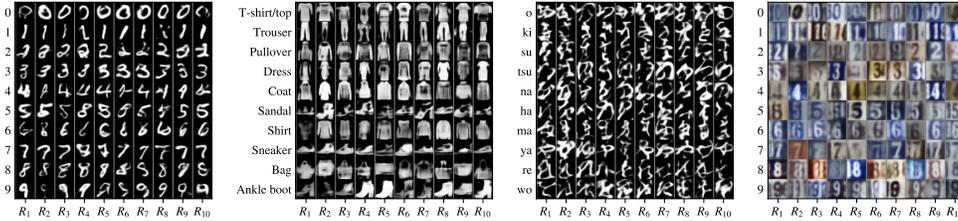
| Model | Configuration | MNIST | F-MNIST | K-MNIST | SVHN |
|------------|------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| Prob-PSENN | (AE, $l = 10$) | 99.26 \pm 0.06 | 91.21 \pm 0.48 | 95.58 \pm 0.36 | 90.91 \pm 0.26 |
| Prob-PSENN | (AE, $l = 20$) | 99.36 \pm 0.05 | 91.76 \pm 0.21 | 96.01 \pm 0.12 | 90.72 \pm 0.41 |
| Prob-PSENN | (AE, $l = 40$) | 99.35 \pm 0.09 | 91.65 \pm 0.30 | 96.00 \pm 0.15 | 90.37 \pm 0.22 |
| Prob-PSENN | (VAE, $l = 20$) | 99.42 \pm 0.05 | 91.91 \pm 0.14 | 96.72 \pm 0.21 | 91.62 \pm 0.56 |
| Prob-PSENN | (VAE, $l = 40$) | 99.38 \pm 0.02 | 91.98 \pm 0.24 | 96.61 \pm 0.17 | 92.14 \pm 0.24 |



(a) Prob-PSENN with a VAE backbone, $l = 40$.



(b) Prob-PSENN with an AE backbone, $l = 10$.



(c) Prob-PSENN with an AE backbone, $l = 40$.

Figure F.21: Visualization of randomly sampled prototype sets from Prob-PSENNs with the following backbone configurations: (a) VAE and $l = 40$, (b) AE and $l = 10$, and (c) AE and $l = 40$. Prob-PSENNs with AE backbones provide a good performance in lower dimensions (e.g., $l = 10$), but struggle to learn high-quality prototype distributions as dimensionality increases, whereas VAEs enable effective scaling to higher-dimensional spaces.

Appendix G Results on the E-MNIST dataset

In order to analyze the scalability of Prob-PSENN in tasks consisting of larger number of classes, we evaluated its performance in the E-MNIST dataset [9] (47 classes), an extension of MNIST which includes both numeric and alphabetic characters. As it can be seen in Figure G.22, Prob-PSENN is still capable of learning suitable distribution over the prototypes for the 47 classes, yielding realistic and diverse prototypes, while achieving a test accuracy of 0.89.

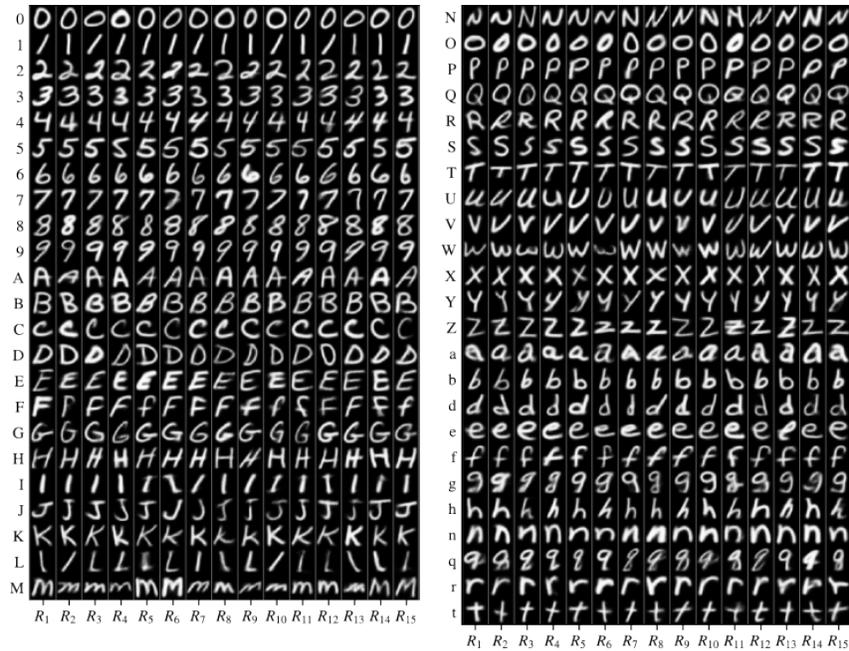


Figure G.22: Prototypes sampled for a Prob-PSENN trained on the E-MNIST [9] dataset (47 classes), consisting of handwritten representations of numbers and letters.

Appendix H Broader Impact

As specified above, the goal of Prob-PSENNs is to promote a transparent-by-design DNN architecture while maintaining competitive predictive performance. In addition to transparency, Prob-PSENNs provide tools to quantify the uncertainty in the model’s predictions and explanations, which further enhances its safety and reliability. Therefore, our work is aligned with three important needs of deep learning models: interpretability, transparency, and reliability. However, a current concern in deep learning models is their lack of robustness to adversarial attacks [6, 41, 45], a vulnerability that also threatens explanation methods and self-explainable models [2, 14, 25, 42]. Therefore, in order to ensure a safe and responsible deployment of Prob-PSENNs in practice, it remains crucial to carefully assess the implications of adversarial attacks, and, consequently, to evaluate the robustness of the model against them.