

# 000 RECONSTRUCTIONNET: A NEURAL NETWORK AR- 001 CHITECTURE FOR UNCERTAINTY-AWARE PREDI- 002 CITIONS WITH EXPLAINABILITY

003  
004  
005  
006 **Anonymous authors**  
007 Paper under double-blind review  
008  
009  
010  
011  
012

## ABSTRACT

013 Uncertainty estimation quantifies a model’s confidence in its predictions, foster-  
014 ing calibrated trust among users. Existing approaches face two key limitations:  
015 (1) most capture only a single type of uncertainty, and (2) they incur additional  
016 training or inference overhead. We propose ReconstructionNet, a neural network  
017 that addresses these limitations by modeling the joint input–output distribution  
018 with class-specific autoencoders. This enables simultaneous prediction and esti-  
019 mation of both aleatoric and distributional uncertainty in a single pass. Across  
020 five real-world datasets, ReconstructionNet matches or surpasses baseline classi-  
021 fiers while producing uncertainty estimates with greater reliability, selectivity, ro-  
022 bustness to false negatives, and strong out-of-distribution detection. Furthermore,  
023 ReconstructionNet’s architecture naturally supports uncertainty explanations, re-  
024 vealing how individual features contribute to prediction uncertainty without extra  
025 computation. Experiments demonstrate that these explanations highlight misclas-  
026 sified regions consistent with human intuition. Together, these contributions es-  
027 tablish ReconstructionNet as a unified framework for trustworthy and interpretable  
028 artificial intelligence.

## 029 1 INTRODUCTION

030 Uncertainty estimation refers to the task of quantifying how uncertain a machine learning model is  
031 about its prediction for each instance. Reliable uncertainty estimates foster calibrated trust by alert-  
032 ing users to cases where the model is likely to be uncertain and erroneous (Toh et al., 2025). Com-  
033 mon methods for uncertainty estimation include Bayesian Neural Networks (BNNs) (Jospin et al.,  
034 2022), Monte Carlo Dropout (Gal & Ghahramani, 2016) and Deep Ensemble (Lakshminarayanan  
035 et al., 2017). These methods quantify one or more of three main types of uncertainty (Malinin &  
036 Gales, 2018): 1) Aleatoric (data), 2) Epistemic (model), and 3) Distributional uncertainty. While  
037 existing work in uncertainty estimation shows promise, it often faces several limitations. Most un-  
038 certainty estimates quantify only a single type of uncertainty and are unable to differentiate between  
039 various sources of uncertainty. Furthermore, many uncertainty estimation methods incur increased  
040 training and inference time.

041 We introduce **ReconstructionNet**, a neural network architecture designed to address the aforemen-  
042 tioned limitations. ReconstructionNet quantifies aleatoric and distributional uncertainty by modeling  
043 the joint input–output distribution with class-specific autoencoders, effectively measuring the dis-  
044 tance of an instance from the training data of each class. This design reduces epistemic uncertainty  
045 by constraining the state space and allows ReconstructionNet to distinguish between aleatoric and  
046 distributional uncertainty.

047 Beyond identifying when a model is uncertain, understanding which features contribute to prediction  
048 uncertainty is equally valuable, giving rise to the emerging field of uncertainty explanation (Wang  
049 et al., 2025; Antorán et al., 2021; Fan, 2025). Uncertainty explanations are typically represented as  
050 vectors of real values, where each value quantifies the significance and impact of an input feature on  
051 the model’s uncertainty. State-of-the-art methods include applying out-of-the-box eXplainable Ar-  
052 tificial Intelligence (XAI) techniques, such as Integrated Gradients (IG) (Sundararajan et al., 2017),  
053 to explain existing uncertainty estimates (Mougan & Nielsen, 2023; Iversen et al., 2024). Other

054 approaches involve observing how input perturbations affect prediction uncertainty (Antorán et al.,  
 055 2021; Wang et al., 2025) and learning which features significantly reduce uncertainty. However,  
 056 these methods require extra modules on top of uncertainty estimation, increasing inference time.  
 057

058 ReconstructionNet’s design enables inbuilt uncertainty explanation: each class-specific autoencoder  
 059 produces feature-wise reconstruction errors scaled by learned error weights for classification. Those  
 060 weighed reconstruction errors quantify each feature’s contribution to the uncertainty, providing un-  
 061 certainty explanations without additional modules or computation.  
 062

063 The contributions of this research are as follows:  
 064

- 064 1. Propose ReconstructionNet, a neural architecture that minimises epistemic uncertainty  
 065 while quantifying and explaining both aleatoric and distributional uncertainty.
- 066 2. Provide a theoretical evaluation of ReconstructionNet’s uncertainty explanations.
- 067 3. Demonstrate the efficacy of ReconstructionNet for prediction, uncertainty estimation, and  
 068 explanation on real-world applications in healthcare and finance.  
 069

## 070 2 RELATED WORK

### 071 2.1 UNCERTAINTY ESTIMATION

072 **Definition 1 (Uncertainty Estimation)** *For instance  $\mathbf{x} \in \mathbb{R}^d$  and a model  $f$ , an uncertainty esti-  
 073 mator  $\sigma(\mathbf{x}; f) : \mathbb{R}^d \rightarrow \mathbb{R}$  assigns a real-valued measure of the prediction uncertainty for  $\mathbf{x}$ .*

074 Uncertainty estimates generally aim to quantify three types of uncertainty:  
 075

- 076 1. **Aleatoric:** Uncertainty arising from inherent noise in the training data. In classification  
 077 problems, this manifests as overlapping classes.
- 078 2. **Epistemic:** Uncertainty from inadequate model parameter fit, reducible by expanding the  
 079 dataset or narrowing the hypothesis space.
- 080 3. **Distributional:** Uncertainty caused by data shifts between the training and prediction set.  
 081

082 **Bayesian methods** capture epistemic uncertainty by modeling distributions over network weights.  
 083 For instance, Bayesian Neural Networks (BNNs) (Jospin et al., 2022) sample from these distribu-  
 084 tions to produce multiple predictions, with the variability reflecting uncertainty. Deep Ensembles  
 085 (Lakshminarayanan et al., 2017) and Monte Carlo Dropout (MCD) (Gal & Ghahramani, 2016) ap-  
 086 proximate this sampling via multiple models or stochastic weight activations within a single net-  
 087 work. More recently, efficient Bayesian last-layer approaches such as Variational Bayesian Last  
 088 Layers (Harrison et al., 2024) and Bayesian Non-negative Decision Layers (Hu et al., 2025) reduce  
 089 computational cost by restricting stochasticity to the final layer. While Bayesian methods remain  
 090 powerful, they typically require additional forward passes compared to deterministic models.  
 091

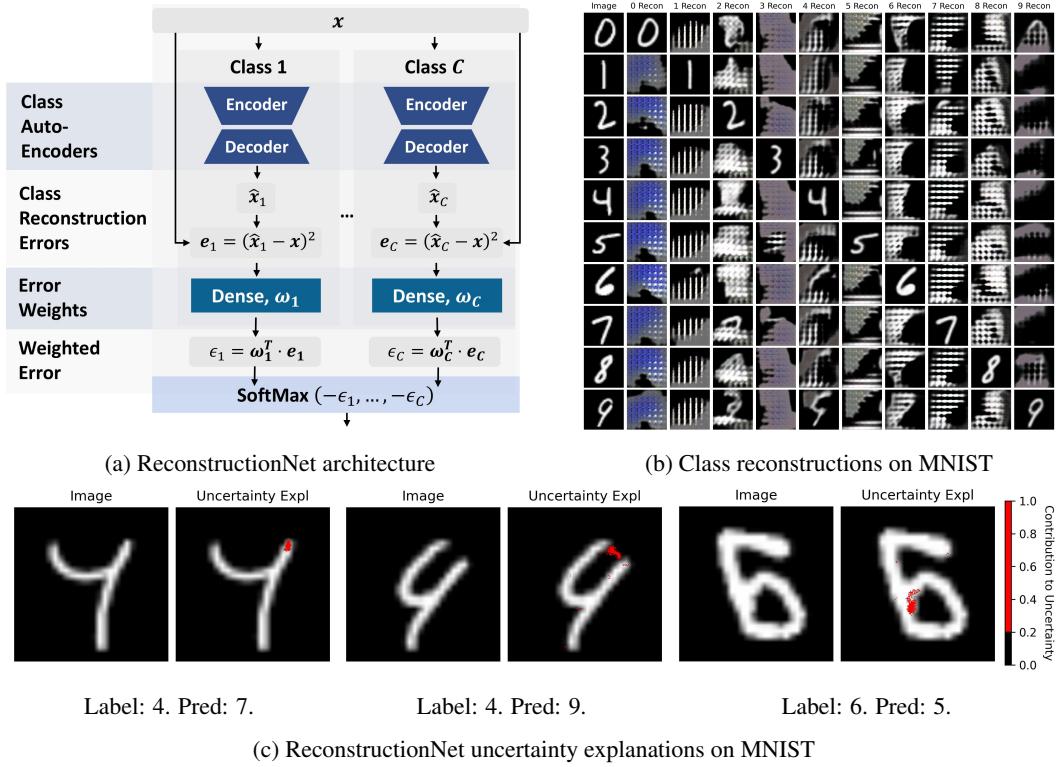
092 **Evidential Deep Learning (EDL)** methods (Ulmer et al., 2023; Amini et al., 2020) are more com-  
 093 putationally efficient than Bayesian approaches, requiring only a single forward pass and one trained  
 094 model. They assume the output follows a well-characterized distribution, predicting the parameters  
 095 of the output distribution from which uncertainty can be derived. For classification, this corre-  
 096 sponds to Dirichlet parameters (Sensoy et al., 2018), with Posterior Networks (PN) (Charpentier  
 097 et al., 2020) extending this via normalizing flows. Despite capturing aleatoric and distributional  
 098 uncertainty, these methods rely on restrictive assumptions about the output distribution.  
 099

100 **Deterministic Uncertainty Methods (DUMs)** (Postels et al., 2022; Charpentier et al., 2023; Ze-  
 101 lenka et al., 2023a) make minimal assumptions about the output distribution, estimating distribu-  
 102 tional uncertainty as the distance of an instance from the training set. A recent DUM, the Re-  
 103 construction Uncertainty Estimate (RUE) (Wang et al., 2024; Korte et al., 2024), uses a decoder  
 104 to reconstruct inputs from the latent representation, with reconstruction error serving as the un-  
 105 certainty measure. ReconstructionNet extends RUE by also capturing aleatoric uncertainty alongside  
 106 distributional uncertainty.  
 107

108 2.2 UNCERTAINTY EXPLANATIONS  
109110 Knowing when a model is unreliable is valuable, but uncertainty explanations provide deeper insight  
111 by quantifying the contribution of each feature to the model’s overall uncertainty.  
112113 **Definition 2 (Uncertainty Explanation)** *Given an instance  $\mathbf{x} \in \mathbb{R}^d$ , a model  $f$ , and an uncertainty  
114 estimator  $\sigma(\mathbf{x}; f) : \mathbb{R}^d \rightarrow \mathbb{R}$ , an uncertainty explanation method  $\zeta(\mathbf{x}; f, \sigma) : \mathbb{R}^d \rightarrow \mathbb{R}^d$  assigns each  
115 input feature a real value reflecting its contribution to  $f(\mathbf{x})$ ’s uncertainty.*  
116117 Uncertainty explanation is a nascent field with two primary approaches:  
118119 **Gradient-Based Methods** (Mougan & Nielsen, 2023; Iversen et al., 2024) applied standard eX-  
120 plainable Artificial Intelligence (XAI) methods, such as Integrated Gradients (IG) (Sundararajan  
121 et al., 2017), to explain uncertainty estimates. While easy to implement, gradient-based expla-  
122 nations can be sparse due to vanishing gradients, making them difficult to interpret.  
123124 **Perturbation-Based Methods** (Antorán et al., 2021; Wang et al., 2025) assess each feature’s contri-  
125 bution to uncertainty by perturbing inputs and measuring the impact on the uncertainty score. Their  
126 accuracy depends on the number of perturbations, making them computationally expensive.  
127128 Both methods require an additional module for uncertainty explanation, increasing inference time.  
129 In contrast, (Zelenka et al., 2023b) computes predictions, uncertainty, and explanations in a sin-  
130 gle forward pass using one model. Based on prototype networks (Snell et al., 2017), it classifies  
131 instances by similarity to class prototypes, with uncertainty explanations derived from the inner  
132 product between the predicted prototype and the instance’s feature map. Similarly, Reconstruc-  
133 tionNet leverages its architecture to provide ante-hoc uncertainty explanations efficiently during  
134 inference.  
135136 2.3 RECONSTRUCTION-BASED METHODS  
137138 **Anomaly Detection.** Reconstruction error has been widely used in autoencoder-based anomaly  
139 detection Chen et al. (2018) to identify deviations from the training distribution. While Reconstruc-  
140 tionNet also leverages reconstruction errors, it differs substantially in several aspects:  
141142 *Reconstruction for Classification and Uncertainty Estimation.* Traditional anomaly-detection au-  
143 toencoders detect out-of-distribution (OOD) inputs only. In contrast, ReconstructionNet uses class-  
144 specific autoencoders to model the joint input-output distribution and classifies by selecting the  
145 class with the lowest reconstruction error. This reconstruction error also serves as a distributional  
146 uncertainty estimate. Reconstruction thus supports both classification and uncertainty estimation in  
147 ReconstructionNet.  
148149 *Reconstruction Errors as Explanations.* Reconstruction errors in ReconstructionNet serve as  
150 feature-level uncertainty explanations, highlighting which input features drive uncertainty, unlike  
151 conventional anomaly detectors that produce only scalar anomaly scores.  
152153 *Training Objective and Architecture.* Unlike unsupervised anomaly detection, Reconstruc-  
154 tionNet uses a supervised training objective and class-specific autoencoders, where each class has its  
155 own encoder-decoder pair. This design enables joint modeling of input-output distributions.  
156157 **Reconstruction Error as Regularizers** Some works use reconstruction loss as a regularizer for  
158 classification Le et al. (2018); Ghifary et al. (2016). ReconstructionNet differs fundamentally in  
159 several aspects:  
160161 *Class-specific autoencoders and joint distribution modeling.* Prior works use a single autoencoder  
162 with conditional probabilities for classification. ReconstructionNet employs one autoencoder per  
163 class to model the joint input-output distribution, classifying via minimum reconstruction error.  
164165 *Uncertainty estimation and explanations.* Unlike prior works, ReconstructionNet leverages recon-  
166 struction errors to quantify uncertainty and provide feature-level explanations.  
167168 *Distinct architecture and prediction mechanism.* Instead of connecting encoders directly to the  
169 prediction head, ReconstructionNet computes weighted reconstruction errors across class-specific  
170 autoencoders to make predictions.  
171

162 

### 3 METHODOLOGY

163  
164 We present ReconstructionNet, a neural network architecture which offers the following features:  
165166 1. Uses the joint input-output probability for classification.  
167 2. Quantifies distributional and aleatoric uncertainty.  
168 3. Generates uncertainty explanations.  
169195 Figure 1: ReconstructionNet overview and examples. (a) ReconstructionNet architecture: Each  
196 class-specific autoencoder is trained to reconstruct only its own class, producing reconstruction  
197 errors inversely related to the joint input–output probability. After applying error weights and softmax  
198 normalization, the resulting probabilities are used for classification (Equation 2). (b) MNIST class  
199 reconstructions: Only the autoencoder for the true class yields a faithful reconstruction, leading to  
200 lower reconstruction error and a higher prediction probability. Other autoencoders generate artifacts  
201 (e.g., checkerboard patterns (Odena et al., 2016)) due to lack of training on mismatched classes. (c)  
202 MNIST uncertainty explanations: Weighted class reconstruction errors highlight uncertain regions.  
203 In the first example, the extra right vertical line increases uncertainty of it being a 7; in the middle,  
204 the missing connector raises uncertainty of it being a 9; in the last, the extra bottom-left vertical line  
205 increases uncertainty of it being a 5. See Appendix A.4 for implementation details.207 

#### 3.1 CLASSIFICATION

208 Consider a classification dataset with  $N$  instances and  $C$  classes, where each instance  $i$  has an input  
209 vector  $\mathbf{x}_i \in \mathbb{R}^d$  and label  $y_i$ . Let  $\mathcal{X}_j$  denote the set of training instances with target label  $j$ , containing  
210  $N_j$  instances. For each instance  $i$  and class  $j$ , the true and predicted class probabilities are  $p_{ij}$  and  
211  $\hat{p}_{ij}$ , respectively.213 Traditional neural networks predict the conditional probability of class  $j$  given instance  $\mathbf{x}_i$ . The  
214 class with the highest conditional probability is then selected as the final predicted class  $\hat{y}_i$ :

215 
$$\hat{y}_i = \arg \max_{j \in \{1, \dots, C\}} \Pr(\hat{y}_i = j \mid \mathbf{x}_i). \quad (1)$$

ReconstructionNet differs from traditional feedforward neural network classifiers in its inference process. ReconstructionNet instead predicts the joint probability of the target  $\hat{y}_i$  and input  $\mathbf{x}_i$  and the target class with the highest joint probability is the final predicted class (Equation 2):

$$\hat{y}_i = \arg \max_{j \in \{1, \dots, C\}} \Pr(\hat{y}_i = j, \mathbf{x}_i). \quad (2)$$

Equation 2 is a valid classification formulation, equivalent to Equation 1 via Bayes' theorem:

$$\arg \max_j \Pr(\hat{y}_i = j, \mathbf{x}_i) = \arg \max_j \Pr(\hat{y}_i = j \mid \mathbf{x}_i) \Pr(\mathbf{x}_i) = \arg \max_j \Pr(\hat{y}_i = j \mid \mathbf{x}_i). \quad (3)$$

To model joint input-output probability, we adopt the model architecture in Figure 1a. Given a classification problem with  $C$  classes, we construct  $C$  autoencoders  $g_1, \dots, g_C$  and train them simultaneously, such that each autoencoder models the joint probability  $\Pr(\hat{y}_i = j, \mathbf{x}_i)$ . During inference, the model computes feature-wise reconstruction errors  $\mathbf{e}_{ij}$  for each autoencoder (Equation 4).  $\hat{\mathbf{x}}_{ij} = g_j(\mathbf{x}_i)$  is the reconstructed input of instance  $i$  by class- $j$  autoencoder  $g_j$ ;  $\hat{\mathbf{x}}_{ij}^k$  is the reconstruction for feature  $k$ .

$$\mathbf{e}_{ij} = \begin{bmatrix} (\hat{\mathbf{x}}_{ij}^1 - \mathbf{x}_i^1)^2 & \dots & (\hat{\mathbf{x}}_{ij}^d - \mathbf{x}_i^d)^2 \end{bmatrix} \quad (4)$$

Next, it calculates a weighted reconstruction error  $\epsilon_{ij}$  for each class  $j$ , where the weights  $\omega_j$  are trainable parameters (Equation 5).

$$\epsilon_{ij} = \omega_j^\top \cdot \mathbf{e}_{ij}. \quad (5)$$

Finally, prediction probability  $\hat{p}_{ij}$  is obtained by applying softmax to the negative weighted errors  $\epsilon_i$ . Negation is used since the higher the reconstruction error, the lower the probability (Figure 1b).

$$\hat{p}_i = \text{softmax}(-\epsilon_i). \quad (6)$$

To train the ensemble of class autoencoders concurrently, we designed a loss function (Equation 7) consisting of two components, with  $\beta$  as a hyperparameter to balance both training objectives.

$$\mathcal{L}_{total} = \mathcal{L}_{CE} + \beta \cdot \mathcal{L}_{Class\_MSE}, \quad \mathcal{L}_{CE} = - \sum_{i=1}^N \sum_{j=1}^C p_{ij} \log(\hat{p}_{ij}) \quad (7)$$

$\mathcal{L}_{CE}$  is the cross-entropy loss; it optimizes the classification performance of the ReconstructionNet. To ensure that the class autoencoders also learn to model the joint probability  $\Pr(\hat{y}_i = j, \mathbf{x}_i)$  while also maintaining predictive accuracy, we introduce the class-dependent mean squared error (MSE):

$$\mathcal{L}_{Class\_MSE}^j = \frac{1}{N_j} \sum_{\mathbf{x}_i \in \mathcal{X}_j} \|\hat{\mathbf{x}}_{ij} - \mathbf{x}_i\|^2. \quad (8)$$

Where each class autoencoder  $g_j$  is trained to minimize reconstruction error for its corresponding instances in  $\mathcal{X}_j$ . The overall reconstruction loss,  $\mathcal{L}_{Class\_MSE}$ , is then the average across all classes:

$$\mathcal{L}_{Class\_MSE} = \frac{1}{N} \sum_{j=1}^C N_j \cdot \mathcal{L}_{Class\_MSE}^j. \quad (9)$$

The formulation of  $\mathcal{L}_{Class\_MSE}$  ensures each class autoencoder is trained exclusively to reconstruct samples from its ground truth class, thereby modeling the joint distribution of the input and target. Assume that the latent vector  $z_j$  from the encoder of class  $j$  follows an isotropic Gaussian distribution,  $\Pr(\mathbf{x}_i, y_i = j \mid z_{ij}) = \mathcal{N}(\hat{\mathbf{x}}, \sigma^2 I)$  (Doersch, 2021; Odaibo, 2019). The probability density function of the distribution of all  $\mathbf{x}$  from class  $j$  can be expressed as:

$$\Pr(\mathbf{x}_i, y_i = j \mid z_{ij}) = \frac{1}{(2\pi\sigma^2)^{d/2}} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \hat{\mathbf{x}}_{ij}\|^2\right). \quad (10)$$

To model this distribution, we optimize our model to maximize the log-likelihood over the training dataset with target class  $j$ :

$$\ell(\hat{\mathbf{x}}, \sigma) = \sum_{i \in \mathcal{X}_j} \left[ -\frac{d}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \|\mathbf{x}_i - \hat{\mathbf{x}}_{ij}\|^2 \right] \quad (11)$$

$$= \sum_{i \in \mathcal{X}_j} \left[ -\frac{d}{2} \log(2\pi\sigma^2) \right] - \frac{N_j}{2\sigma^2} \cdot \mathcal{L}_{Class\_MSE}^j. \quad (12)$$

We observe that maximizing the log-likelihood is equivalent to minimizing  $\mathcal{L}_{Class\_MSE}^j$ . This demonstrates that minimizing  $\mathcal{L}_{Class\_MSE}$  leads to the training of class autoencoders that model the joint distribution of the input and target, with reconstruction error inversely related to the joint probability. Additionally, the architecture and loss function of ReconstructionNet limit its hypothesis space, making it more resistant to epistemic uncertainty (Hüllermeier & Waegeman, 2021).

### 3.2 UNCERTAINTY ESTIMATION

By modelling the joint input-output probability, we can quantify:

1. **Aleatoric Uncertainty:** When instances lie in overlapping class regions, several classes have similarly *high* probabilities above threshold  $\theta_1$ , signalling high aleatoric uncertainty.

$$\exists \mathbf{C}_i \subseteq \{1, \dots, C\}, |\mathbf{C}_i| \geq 2 \text{ s.t.}$$

$$\Pr(\hat{y}_i = c_1, \mathbf{x}_i) \approx \Pr(\hat{y}_i = c_2, \mathbf{x}_i) \geq \theta_1, \forall c_1, c_2 \in \mathbf{C}_i. \quad (13)$$

2. **Distributional Uncertainty:** When instances lie beyond the training distribution, the joint probabilities of all classes are similarly *low*, below some threshold  $\theta_2$ , signaling high distributional uncertainty.

$$\forall c \in \{1, \dots, C\}, \Pr(\hat{y}_i = c, \mathbf{x}_i) \leq \theta_2. \quad (14)$$

Modelling joint probabilities allows differentiation between uncertainty types: instances show high aleatoric uncertainty when their most probable classes have probabilities above  $\theta_1$ , and high distributional uncertainty when all class probabilities are below  $\theta_2$ .

The notion of aleatoric uncertainty, as illustrated in Equation 13, is nicely captured by Shannon entropy, reflecting evenly spread high prediction probabilities across overlapping classes:

**Definition 3 (Aleatoric Uncertainty)** *For an instance  $\mathbf{x}_i$ , aleatoric uncertainty is quantified using the Shannon entropy (Shannon, 1948) of the prediction probabilities  $\hat{p}_{ij}$ :*

$$\sigma_{aleatoric}(\mathbf{x}_i) = - \sum_{j=1}^C \hat{p}_{ij} \log \hat{p}_{ij}.$$

Distributional uncertainty in Equation 14 is captured by the predicted class's reconstruction error. By design, the predicted class has the lowest error and highest probability among all classes; therefore, if its error is high (probability low), all other classes also have low probabilities.

**Definition 4 (Distributional Uncertainty)** *For an instance  $\mathbf{x}_i$ , distributional uncertainty is the reconstruction error of the predicted class  $\hat{y}_i$ 's autoencoder:*

$$\sigma_{dist}(\mathbf{x}_i) = \|\mathbf{e}_{i\hat{y}_i}\|_1.$$

### 3.3 UNCERTAINTY EXPLANATION

The weighted reconstruction errors  $\zeta$  of the predicted class  $\hat{y}$  serve as uncertainty explanations, as they represent feature uncertainties scaled by their importance to the prediction.

**Definition 5 (ReconstructionNet Explanation)** *The ReconstructionNet Explanation for instance  $\mathbf{x}_i$  and its predicted class  $\hat{y}_i$  is the weighted reconstruction errors of predicted class  $\hat{y}_i$ :*

$$\zeta(\mathbf{x}_i) = \omega_{\hat{y}_i} \odot \mathbf{e}_{i\hat{y}_i}.$$

ReconstructionNet uncertainty explanations exhibit the following three properties: (1) Implementation Invariance (Sundararajan et al., 2017), (2) Sensitivity (Sundararajan et al., 2017), and (3) Consistency (Lundberg & Lee, 2017). An explanation that satisfies all three properties is (1) consistent across different implementations, (2) does not attribute irrelevant features incorrectly, and (3) preserves the relative importance of features across models.

An uncertainty explanation is *implementation-invariant* if, for a pair of functionally equivalent prediction models, the same explanation is generated for any instance. A pair of functionally equivalent models are models that yield the same output for a given set of inputs.

324 **Property 1 (Implementation Invariance)** *Given two functionally equivalent prediction models,  $f$   
 325 and  $f'$ , an explanation function  $\chi$  is implementation-invariant if and only if, for any instance  $\mathbf{x}$ , the  
 326 explanations derived from both models using  $\chi$  are equivalent:  $\chi(\mathbf{x}; f) = \chi(\mathbf{x}; f')$ .*  
 327

328 ReconstructionNet’s uncertainty explanations  $\zeta$ , defined as the weighted reconstruction errors of the  
 329 predicted class (Definition 5), are implementation-invariant: for two functionally equivalent models,  
 330 both the error weights  $\omega$  and the reconstruction errors remain identical.

331 A *sensitive* explanation function allocates zero attribution to irrelevant features for prediction. A  
 332 feature is considered irrelevant if a change in its value does not impact the model’s prediction.  
 333

334 **Property 2 (Sensitivity)** *An explanation function  $\chi$  is sensitive if it assigns a zero feature attribu-  
 335 tion value,  $\chi(\mathbf{x}; f)_i = 0$ , to features  $i$  that are irrelevant to the prediction.*  
 336

337 ReconstructionNet’s uncertainty explanations  $\zeta$  are sensitive. Since the error weight  $\omega$  encodes each  
 338 feature’s contribution to the final prediction, any change in feature uncertainty that does not affect  
 339 the prediction must have a weight of zero, yielding a weighted reconstruction error  $\zeta$  of zero.

340 An explanation is *consistent* if a feature’s uncertainty attribution does not decrease when the model  
 341 is altered to increase that feature’s contribution.

342 **Property 3 (Consistency)** *Let  $f'$  be a modification of  $f$  where feature  $i$ ’s contribution is increased.  
 343 For an instance  $\mathbf{x}$  and  $\mathbf{x}^{\setminus i}$  with  $x_i = 0$ , an explanation  $\chi$  is consistent if:*

$$344 \quad f'(\mathbf{x}) - f'(\mathbf{x}^{\setminus i}) \geq f(\mathbf{x}) - f(\mathbf{x}^{\setminus i}) \quad \text{then} \quad \chi(\mathbf{x}; f')_i \geq \chi(\mathbf{x}; f)_i.$$

348 ReconstructionNet’s uncertainty explanations,  $\zeta$ , are consistent. For any instance, reconstruction  
 349 errors stay the same between  $f$  and  $f'$ , since the joint input-output probability modeled by the  
 350 autoencoders is independent of feature contributions. Thus, only the error weights  $\omega$  can change; if  
 351 a feature’s weight increases in  $f'$ , its uncertainty attribution also increases, satisfying consistency.

## 353 4 EXPERIMENTS

### 355 4.1 DATASETS

357 We use tabular datasets (Covid, Diabetes, Fund) to evaluate uncertainty reliability, selectivity, and ro-  
 358 bustness, medical image datasets (ISIC, OCTMNIST) for OOD detection, and MNIST, ISIC, Covid  
 359 and Diabetes to assess uncertainty explanation correctness. The datasets are summarized as follows:  
 360

- 361 1. **Covid** (Hinns et al., 2021) is a tabular dataset of United Kingdom’s COVID-19 policies  
 362 and regional case counts, labelled by whether  $R_t$  (effective reproduction number)  $> 1$ .  
 363
- 364 2. **Diabetes** (Mustafa, 2023) is a binary tabular dataset of demographics, pre-existing condi-  
 365 tions, and vital signs, with labels indicating diabetes status.  
 366
- 367 3. **Fund** (Kovvuri et al., 2023) is a binary tabular dataset from 4,330 funds, using macro-  
 368 economic indicators, fund allocations, HHI, and past performance to predict if a fund’s net  
 369 asset value (NAV) exceeds the previous quarter’s.  
 370
- 371 4. **ISIC** (Codella et al., 2019; Tschandl et al., 2018) is a dermoscopic image dataset of seven  
 372 skin conditions. For OOD detection, we created three datasets with decreasing similarity  
 373 to ISIC: BCN-IN (images from seen classes of BCN20000 (Combalia et al., 2019)), BCN-  
 374 OUT (images from the unseen Scar class), and ChestMNIST (Wang et al., 2017).  
 375
- 376 5. **OCTMNIST** (Kermany et al., 2018) is a retinal OCT dataset with four classes. For OOD  
 377 detection, we used three datasets of decreasing similarity to OCTMNIST: OCTDL-IN (im-  
 378 ages from seen classes of OCTDL (Kulyabin et al., 2024)), OCTDL-OUT (images from  
 379 unseen classes), and ChestMNIST (Wang et al., 2017).  
 380
- 381 6. **MNIST** (Deng, 2012) is an image dataset of ten handwritten digits.  
 382

378 4.2 BASELINES  
379380 We compared ReconstructionNet against six recent state-of-the-art uncertainty estimation methods:  
381 (1) Entropy, (2) MCD, (3) DE, (4) PN, (5) BNN, and (6) EDL. Implementation and tuning details of  
382 the baselines are provided in the appendix.383 4.3 EVALUATION METRICS  
384385 We evaluate prediction performance with Area Under receiver operating characteristic Curve (AUC)  
386 and accuracy. To evaluate uncertainty estimation performance, we used the following metrics:  
387388 1. **Correlation** (Mi et al., 2022; Upadhyay et al., 2022) measures reliability as the Pearson  
389 correlation between uncertainty and error, with higher values indicating better reliability.  
390 2. **AURC** (Ding et al., 2020) quantifies selectivity by measuring the area under the risk-  
391 coverage curve (AURC), with lower values indicating better selectivity.  
392 3.  $\sigma$ -**Risk Score** measures robustness to false negatives as errors for instances with normalized  
393 uncertainty below  $\sigma = \{0.1, 0.2, 0.3, 0.4\}$ , with lower values indicating greater resilience.  
394 4. **OOD Detection** (Lakshminarayanan et al., 2017; Postels et al., 2020; Malinin & Gales,  
395 2018) measures how well uncertainty distinguishes in-distribution from OOD instances  
396 using AUROC, with higher values indicating better detection.  
397398 4.4 RESULTS  
399400 4.4.1 PREDICTION PERFORMANCE  
401402 Table 1: Model prediction performance. NN refers to an MLP for tabular and a ResNet for image  
403 datasets. The best-performing model for each metric is **bolded**, while the second-best is underlined.  
404

	Covid (Tabular)		Diabetes (Tabular)		Fund (Tabular)		ISIC (Image)		OCTMNIST (Image)	
	AUC (↑)	Acc (↑)								
RN (Ours)	<u>0.95 ± 0.00</u>	<u>0.88 ± 0.00</u>	<u>0.97 ± 0.00</u>	<u>0.92 ± 0.01</u>	<b>0.75 ± 0.01</b>	<b>0.71 ± 0.00</b>	<b>0.91 ± 0.01</b>	<b>0.76 ± 0.01</b>	<b>0.99 ± 0.00</b>	<b>0.91 ± 0.02</b>
NN	0.95 ± 0.00	<b>0.89 ± 0.00</b>	<b>0.98 ± 0.00</b>	0.93 ± 0.01	0.60 ± 0.02	<b>0.71 ± 0.00</b>	0.87 ± 0.01	0.70 ± 0.01	<b>0.99 ± 0.01</b>	0.88 ± 0.03
MCD	0.95 ± 0.00	<b>0.89 ± 0.00</b>	<b>0.98 ± 0.00</b>	<u>0.93 ± 0.01</u>	0.59 ± 0.02	<b>0.71 ± 0.00</b>	0.87 ± 0.01	0.70 ± 0.01	<b>0.99 ± 0.01</b>	0.88 ± 0.03
DE	0.95 ± 0.00	<b>0.89 ± 0.00</b>	<b>0.98 ± 0.00</b>	0.93 ± 0.00	0.59 ± 0.01	<b>0.71 ± 0.00</b>	0.89 ± 0.00	0.72 ± 0.01	<b>0.99 ± 0.00</b>	0.89 ± 0.02
PN	<b>0.96 ± 0.01</b>	0.88 ± 0.01	0.97 ± 0.00	0.92 ± 0.01	0.53 ± 0.00	<b>0.71 ± 0.00</b>	0.70 ± 0.01	0.64 ± 0.01	0.94 ± 0.02	0.80 ± 0.05
BNN	0.91 ± 0.02	0.85 ± 0.02	0.97 ± 0.01	0.91 ± 0.01	0.54 ± 0.01	<b>0.71 ± 0.00</b>	0.60 ± 0.02	0.52 ± 0.02	<b>0.99 ± 0.00</b>	0.87 ± 0.02
EDL	0.92 ± 0.02	0.79 ± 0.18	0.87 ± 0.07	<b>0.95 ± 0.02</b>	0.58 ± 0.05	<u>0.62 ± 0.17</u>	0.53 ± 0.01	0.51 ± 0.01	<u>0.96 ± 0.00</u>	0.83 ± 0.02

414 Table 1 compares model performance. ReconstructionNet (RN) achieved strong results across  
415 datasets, with second-highest AUC and accuracy on COVID and Diabetes, and the highest AUC  
416 and accuracy on Fund, ISIC, and OCTMNIST. Notably, on Fund, ReconstructionNet outperformed  
417 others in AUC despite similar accuracy, suggesting it learns discriminative features robust to data  
418 shifts (caused by the COVID-19 pandemic) and is resilient to epistemic uncertainty.  
419420 4.4.2 UNCERTAINTY ESTIMATION PERFORMANCE  
421422 **Aleatoric Uncertainty:** Table 2 summarizes uncertainty estimation performance. Reconstruction-  
423 Net outperformed all models on COVID except  $\sigma$ -risk at  $\sigma = 0.1$ , where it matched MLP Entropy  
424 and MCD. On Diabetes and Fund, it consistently ranked among the top across metrics.  
425426 **Distributional Uncertainty:** Table 3 presents the OOD detection performance of all uncertainty  
427 estimates. ReconstructionNet achieves the highest AUROC across all OOD datasets and improves  
428 as datasets deviate further from in-distribution data, showing it effectively ranks dataset dissimilarity.  
429 In contrast, other estimates like Entropy, MCD, PN, and BNN drop on highly dissimilar datasets (see  
430 ISIC) or struggle with unseen classes, as seen on OCTMNIST.  
431432 **Distinguishing Between Aleatoric and Distributional:** To illustrate ReconstructionNet’s ability  
433 to separate uncertainty types, we visualize aleatoric and distributional uncertainty over time on the  
434 Fund dataset (Figure 2). Aleatoric uncertainty is highest before January 2020, while during the 2020

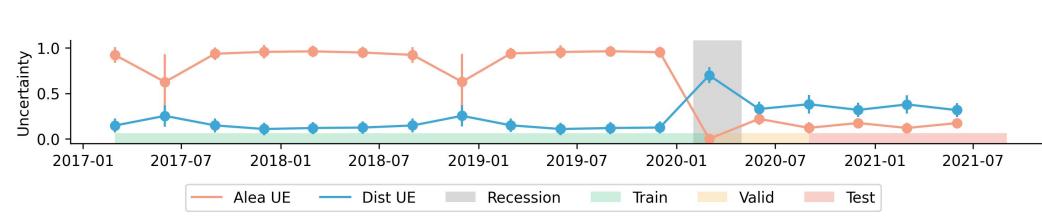
432 Table 2: Uncertainty estimation performance. The best-performing model for each metric is **bolded**,  
 433 while the second-best model is underlined.

Data	UE	Correlation ( $\uparrow$ )	AURC ( $\downarrow$ )	$\sigma = 0.1$ ( $\downarrow$ )	$\sigma = 0.2$ ( $\downarrow$ )	$\sigma = 0.3$ ( $\downarrow$ )	$\sigma = 0.4$ ( $\downarrow$ )
Covid	RN (Ours)	<b>0.823 <math>\pm</math> 0.016</b>	<b>0.028 <math>\pm</math> 0.002</b>	0.005 $\pm$ 0.001	<b>0.006 <math>\pm</math> 0.002</b>	<b>0.007 <math>\pm</math> 0.002</b>	<b>0.012 <math>\pm</math> 0.005</b>
	MLP Entropy	0.664 $\pm$ 0.012	0.030 $\pm$ 0.002	<b>0.003 <math>\pm</math> 0.004</b>	0.016 $\pm$ 0.006	0.022 $\pm$ 0.006	0.032 $\pm$ 0.005
	MLP MCD	0.639 $\pm$ 0.012	0.032 $\pm$ 0.003	0.004 $\pm$ 0.005	0.018 $\pm$ 0.009	0.029 $\pm$ 0.006	0.042 $\pm$ 0.005
	MLP DE	0.603 $\pm$ 0.031	0.030 $\pm$ 0.003	0.011 $\pm$ 0.003	0.020 $\pm$ 0.007	0.032 $\pm$ 0.007	0.040 $\pm$ 0.005
	PN Epis	0.313 $\pm$ 0.082	0.038 $\pm$ 0.011	0.021 $\pm$ 0.010	0.034 $\pm$ 0.015	0.041 $\pm$ 0.017	0.051 $\pm$ 0.016
	PN Alea	<u>0.690 <math>\pm</math> 0.034</u>	<u>0.029 <math>\pm</math> 0.007</u>	0.016 $\pm$ 0.005	0.022 $\pm$ 0.007	0.032 $\pm$ 0.006	0.039 $\pm$ 0.005
	BNN	0.030 $\pm$ 0.027	0.132 $\pm$ 0.018	0.104 $\pm$ 0.070	0.159 $\pm$ 0.023	0.161 $\pm$ 0.023	0.132 $\pm$ 0.015
Diabetes	EDL	0.674 $\pm$ 0.068	0.115 $\pm$ 0.094	0.238 $\pm$ 0.105	0.073 $\pm$ 0.009	0.058 $\pm$ 0.004	0.045 $\pm$ 0.005
	RN (Ours)	<u>0.902 <math>\pm</math> 0.011</u>	0.008 $\pm$ 0.001	<b>0.000 <math>\pm</math> 0.000</b>	<u>0.001 <math>\pm</math> 0.000</u>	<b>0.001 <math>\pm</math> 0.000</b>	<b>0.001 <math>\pm</math> 0.000</b>
	MLP Entropy	0.864 $\pm$ 0.009	<b>0.007 <math>\pm</math> 0.002</b>	<b>0.000 <math>\pm</math> 0.000</b>	<b>0.000 <math>\pm</math> 0.000</b>	<b>0.001 <math>\pm</math> 0.000</b>	<b>0.001 <math>\pm</math> 0.000</b>
	MLP MCD	0.845 $\pm$ 0.022	0.008 $\pm$ 0.002	<b>0.000 <math>\pm</math> 0.000</b>	<b>0.000 <math>\pm</math> 0.000</b>	<b>0.001 <math>\pm</math> 0.000</b>	0.002 $\pm$ 0.001
	MLP DE	0.733 $\pm$ 0.031	0.011 $\pm$ 0.001	<b>0.000 <math>\pm</math> 0.000</b>	0.001 $\pm$ 0.000	0.005 $\pm$ 0.001	0.011 $\pm$ 0.003
	PN Epis	0.525 $\pm$ 0.297	0.020 $\pm$ 0.021	0.006 $\pm$ 0.009	0.016 $\pm$ 0.023	0.023 $\pm$ 0.030	0.028 $\pm$ 0.033
	PN Alea	0.812 $\pm$ 0.022	0.009 $\pm$ 0.001	0.001 $\pm$ 0.000	0.001 $\pm$ 0.000	0.002 $\pm$ 0.001	0.003 $\pm$ 0.001
Fund	BNN	-0.619 $\pm$ 0.012	0.214 $\pm$ 0.016	0.262 $\pm$ 0.024	0.222 $\pm$ 0.016	0.185 $\pm$ 0.016	0.136 $\pm$ 0.013
	EDL	<b>0.942 <math>\pm</math> 0.017</b>	0.046 $\pm$ 0.034	<b>0.000 <math>\pm</math> 0.000</b>	<u>0.001 <math>\pm</math> 0.002</u>	0.006 $\pm$ 0.006	0.011 $\pm$ 0.010
	RN (Ours)	<b>0.385 <math>\pm</math> 0.008</b>	<b>0.141 <math>\pm</math> 0.003</b>	<u>0.039 <math>\pm</math> 0.004</u>	<b>0.048 <math>\pm</math> 0.005</b>	<b>0.077 <math>\pm</math> 0.015</b>	<b>0.126 <math>\pm</math> 0.023</b>
	MLP Entropy	0.131 $\pm$ 0.034	0.212 $\pm$ 0.023	0.064 $\pm$ 0.043	0.132 $\pm$ 0.063	0.187 $\pm$ 0.036	0.241 $\pm$ 0.012
	MLP MCD	0.128 $\pm$ 0.038	0.236 $\pm$ 0.018	0.073 $\pm$ 0.067	0.124 $\pm$ 0.062	0.165 $\pm$ 0.046	0.210 $\pm$ 0.028
	MLP DE	0.022 $\pm$ 0.045	0.287 $\pm$ 0.032	0.364 $\pm$ 0.331	0.330 $\pm$ 0.192	0.283 $\pm$ 0.031	0.286 $\pm$ 0.045
	PN Epis	0.016 $\pm$ 0.001	0.291 $\pm$ 0.017	0.279 $\pm$ 0.028	0.278 $\pm$ 0.025	0.278 $\pm$ 0.022	0.279 $\pm$ 0.021
Fund	PN Alea	<u>0.166 <math>\pm</math> 0.026</u>	0.271 $\pm$ 0.001	0.221 $\pm$ 0.007	0.225 $\pm$ 0.009	0.230 $\pm$ 0.013	0.235 $\pm$ 0.009
	BNN	-0.198 $\pm$ 0.028	0.364 $\pm$ 0.010	0.665 $\pm$ 0.189	0.543 $\pm$ 0.125	0.442 $\pm$ 0.065	0.405 $\pm$ 0.026
	EDL	0.162 $\pm$ 0.118	0.301 $\pm$ 0.202	<b>0.016 <math>\pm</math> 0.015</b>	<u>0.055 <math>\pm</math> 0.062</u>	0.101 $\pm$ 0.078	0.174 $\pm$ 0.092

456 Table 3: Out-of-distribution (OOD) detection performance, measured using AUROC. OOD datasets  
 457 are presented in order of increasing deviation from the in-distribution. The best-performing model  
 458 for each metric is **bolded**, while the second-best model is underlined.

	ISIC			OCTMNIST		
	BCN-IN ( $\uparrow$ )	BCN-OUT ( $\uparrow$ )	ChestMNIST ( $\uparrow$ )	OCTDL-IN ( $\uparrow$ )	OCTDL-OUT ( $\uparrow$ )	ChestMNIST ( $\uparrow$ )
RN (Ours)	<b>0.777 <math>\pm</math> 0.086</b>	<b>0.846 <math>\pm</math> 0.070</b>	<b>0.919 <math>\pm</math> 0.097</b>	<b>0.783 <math>\pm</math> 0.082</b>	<b>0.866 <math>\pm</math> 0.049</b>	<b>1.000 <math>\pm</math> 0.000</b>
ResNet Entropy	0.742 $\pm$ 0.010	0.757 $\pm$ 0.027	0.664 $\pm$ 0.072	0.674 $\pm$ 0.031	0.386 $\pm$ 0.066	0.826 $\pm$ 0.060
ResNet MCD	<u>0.746 <math>\pm</math> 0.004</u>	<u>0.770 <math>\pm</math> 0.016</u>	0.678 $\pm$ 0.080	0.661 $\pm$ 0.034	0.392 $\pm$ 0.064	0.844 $\pm$ 0.060
ResNet DE	0.720 $\pm$ 0.020	0.703 $\pm$ 0.029	<u>0.728 <math>\pm</math> 0.055</u>	0.739 $\pm$ 0.009	0.497 $\pm$ 0.029	<u>0.900 <math>\pm</math> 0.028</u>
PN Epis	0.674 $\pm$ 0.021	0.686 $\pm$ 0.039	<u>0.654 <math>\pm</math> 0.050</u>	0.565 $\pm$ 0.052	0.483 $\pm$ 0.087	<u>0.470 <math>\pm</math> 0.145</u>
PN Alea	0.660 $\pm$ 0.026	0.669 $\pm$ 0.041	0.631 $\pm$ 0.065	0.625 $\pm$ 0.063	0.492 $\pm$ 0.025	0.500 $\pm$ 0.107
BNN	0.556 $\pm$ 0.011	0.540 $\pm$ 0.056	0.547 $\pm$ 0.059	<u>0.742 <math>\pm</math> 0.014</u>	0.734 $\pm$ 0.016	0.896 $\pm$ 0.015
EDL	0.533 $\pm$ 0.011	0.539 $\pm$ 0.023	0.619 $\pm$ 0.024	0.633 $\pm$ 0.012	0.669 $\pm$ 0.011	0.760 $\pm$ 0.013

471 recession (per the National Bureau of Economic Research), distributional uncertainty predominates,  
 472 demonstrating ReconstructionNet’s discriminative capability.



482 Figure 2: Aleatoric (Alea) and distributional (Dist) uncertainty over time on the Fund dataset. Un-  
 483 certainty values are min–max normalized, and error bars represent one standard deviation. In early  
 484 2020, distributional uncertainty overtook aleatoric uncertainty due to COVID-19–induced shifts,  
 485 while aleatoric uncertainty dropped as increasing class imbalance reduced label ambiguity.

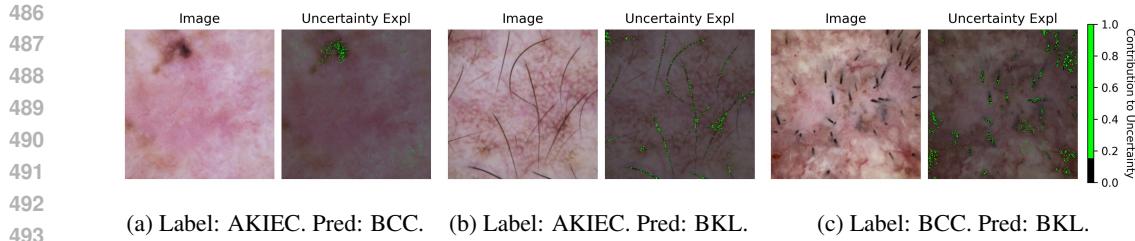


Figure 3: Uncertainty explanation illustration using images from the ISIC dataset. Positive attributions were min–max normalized, thresholded (pixel uncertainty  $> 0.15$  shown in green), and overlaid for clarity. The highlighted regions “explain” the prediction uncertainty. AKIEC refers to Actinic Keratosis, BCC to Basal Cell Carcinoma, and BKL to Benign Keratosis.

#### 4.4.3 UNCERTAINTY EXPLANATION PERFORMANCE

To highlight the practical value of our uncertainty explanations in identifying regions that contribute to misclassification, we qualitatively assess them on the ISIC dataset. In Figure 3a, the image was misclassified as BCC instead of AKIEC, with uncertainty concentrated on the mole at the top-left corner, likely because such dark, mole-like spots resemble features more typical of BCC than AKIEC (Lee, 2017). In Figures 3b and 3c, both images were misclassified as BKL, with uncertainty concentrated on the hairs covering a substantial portion of the lesion, as such occlusions can obscure diagnostic features. These visual explanations align with human intuition, demonstrating their effectiveness in pinpointing input features that confuse the model.

Table 4: Top-k accuracy of each uncertainty-explanation method in identifying perturbed features on the Covid and Diabetes datasets.

Method	Covid			Diabetes		
	Top-1 Acc	Top-3 Acc	Top-5 Acc	Top-1 Acc	Top-3 Acc	Top-5 Acc
IG	$0.163 \pm 0.003$	$0.355 \pm 0.017$	$0.436 \pm 0.020$	$0.065 \pm 0.008$	$0.221 \pm 0.005$	$0.446 \pm 0.022$
SHAP	$0.102 \pm 0.011$	$0.193 \pm 0.020$	$0.262 \pm 0.017$	$0.093 \pm 0.004$	$0.247 \pm 0.008$	$0.387 \pm 0.014$
RN (Ours)	<b><math>0.437 \pm 0.035</math></b>	<b><math>0.604 \pm 0.027</math></b>	<b><math>0.646 \pm 0.023</math></b>	<b><math>0.272 \pm 0.064</math></b>	<b><math>0.509 \pm 0.045</math></b>	<b><math>0.599 \pm 0.040</math></b>

We also evaluate the performance of our uncertainty explanations on tabular data using the covariate-shift experiment from Watson et al. (2023). For each test instance, we randomly perturb one feature by adding noise drawn from  $\mathcal{N}(0.5, 0.1)$  and assess whether an explanation ranks this perturbed feature among its top-k most uncertain features (Top-k Accuracy). This directly measures whether the explanation identifies the feature contributing to distributional uncertainty.

We compare ReconstructionNet (RN) explanations with: (1) Integrated Gradients (IG), a gradient-based method, and (2) KernelSHAP, a perturbation-based method, both applied to explain the entropy of ReconstructionNet’s predictive distribution (Iversen et al., 2024). Across the COVID and Diabetes datasets (Table 4), RN achieves consistently higher top-k accuracy for all  $k$ , demonstrating its effectiveness in localizing the source of uncertainty.

## 5 CONCLUSION

This paper proposed ReconstructionNet, a neural network for reliable uncertainty estimation alongside classification. Unlike models based on conditional probability, ReconstructionNet uses class-specific autoencoders to model the input–output joint distribution, predicting the class with maximal joint probability (or minimal reconstruction error). This approach quantifies aleatoric and distributional uncertainty while minimizing epistemic uncertainty in a single training session. Across five real-world datasets, ReconstructionNet achieved comparable or improved classification performance, with uncertainty estimates showing superior reliability, selectivity, robustness to false negatives, and strong OOD detection. Its inbuilt, cost-free explanations highlight input features contributing to uncertainty, with theoretical properties of (1) Implementation Invariance, (2) Sensitivity, and (3) Consistency. While ReconstructionNet performs best on large, balanced datasets, this limitation suggests future directions, such as quantifying epistemic uncertainty.

540 

## 6 REPRODUCIBILITY STATEMENT

541  
 542 For implementation details, see the code repository: <https://anonymous.4open.science/r/ReconstructionNet-4F8C/>. For tabular datasets, input features were normalized using min–max scaling prior to training. For image datasets, pixel intensities were scaled from [0, 255] to [0, 1].  
 543  
 544

545 

## REFERENCES

546  
 547 Alexander Amini, Wilko Schwarting, Ava Soleimany, and Daniela Rus. Deep eviden-  
 548 tial regression. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina  
 549 Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Sys-*  
 550 *tems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS*  
 551 *2020, December 6-12, 2020, virtual, volume 33, pp. 14927-14937, virtual, 2020. Curran*  
 552 *Associates, Inc.* URL <https://proceedings.neurips.cc/paper/2020/hash/aab085461de182608ee9f607f3f7d18f-Abstract.html>.  
 553  
 554  
 555 Javier Antorán, Umang Bhatt, Tameem Adel, Adrian Weller, and José Miguel Hernández-Lobato.  
 556 Getting a CLUE: A method for explaining uncertainty estimates. In *9th International Conference*  
 557 *on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021, Austria, 2021.*  
 558 OpenReview.net. URL <https://openreview.net/forum?id=XSLF1XFq5h>.  
 559  
 560 Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah.  
 561 Signature Verification using a "Siamese" Time Delay Neural Network. In *Ad-*  
 562 *vances in Neural Information Processing Systems*, volume 6, USA, 1993. Morgan-  
 563 Kaufmann. URL [https://proceedings.neurips.cc/paper\\_files/paper/1993/hash/288cc0ff022877bd3df94bc9360b9c5d-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/1993/hash/288cc0ff022877bd3df94bc9360b9c5d-Abstract.html).  
 564  
 565 Bertrand Charpentier, Daniel Zügner, and Stephan Günnemann. Posterior Network: Uncertainty  
 566 Estimation without OOD Samples via Density-Based Pseudo-Counts. In *Advances in Neural*  
 567 *Information Processing Systems*, volume 33, pp. 1356–1367, Online, 2020. Curran Associates,  
 568 Inc.  
 569  
 570 Bertrand Charpentier, Chenxiang Zhang, and Stephan Günnemann. Training, architecture, and prior  
 571 for deterministic uncertainty methods. *CoRR*, abs/2303.05796, 2023. doi: 10.48550/ARXIV.  
 572 2303.05796. URL <https://doi.org/10.48550/arXiv.2303.05796>.  
 573  
 574 Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic  
 575 minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.  
 576  
 577 Zhaomin Chen, Chai Kiat Yeo, Bu Sung Lee, and Chiew Tong Lau. Autoencoder-based network  
 578 anomaly detection. In *2018 Wireless telecommunications symposium (WTS)*, pp. 1–5. IEEE, 2018.  
 579  
 580 Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with  
 581 application to face verification. In *2005 IEEE computer society conference on computer vision*  
 582 *and pattern recognition (CVPR'05)*, volume 1, pp. 539–546, USA, 2005. IEEE.  
 583  
 584 Noel Codella, Veronica Rotemberg, Philipp Tschandl, M. Emre Celebi, Stephen Dusza, David Gut-  
 585 man, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, Harald Kittler, and  
 586 Allan Halpern. Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by  
 587 the International Skin Imaging Collaboration (ISIC), March 2019.  
 588  
 589 Marc Combalia, Noel CF Codella, Veronica Rotemberg, Brian Helba, Veronica Vilaplana, Ofer  
 590 Reiter, Cristina Carrera, Alicia Barreiro, Allan C Halpern, Susana Puig, et al. Bcn20000: Dermo-  
 591 scopic lesions in the wild. *arXiv preprint arXiv:1908.02288*, 2019.  
 592  
 593 Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE*  
 594 *Signal Processing Magazine*, 29(6):141–142, 2012.  
 595  
 596 Yukun Ding, Jinglan Liu, Jinjun Xiong, and Yiyu Shi. Revisiting the Evaluation of Uncertainty  
 597 Estimation and Its Application to Explore Model Complexity-Uncertainty Trade-Off. In *2020*  
 598 *IEEE CVPRW*, pp. 22–31, Seattle, WA, USA, June 2020. IEEE. ISBN 978-1-72819-360-1. doi:  
 599 10.1109/CVPRW50498.2020.00010.

594 Carl Doersch. Tutorial on Variational Autoencoders, 2021. URL <http://arxiv.org/abs/1606.05908>.

595

596

597 Xiuyi Fan. Position paper: Integrating explainability and uncertainty estimation in medical AI. In

598 *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN)*, 2025.

599 Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model

600 Uncertainty in Deep Learning. In *Proceedings of The 33rd ICML*, pp. 1050–1059, USA, June

601 2016. PMLR.

602

603 Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep

604 reconstruction-classification networks for unsupervised domain adaptation. In *European con-*

605 *ference on computer vision*, pp. 597–613. Springer, 2016.

606 James Harrison, John Willes, and Jasper Snoek. Variational bayesian last layers. *arXiv preprint*

607 *arXiv:2404.11599*, 2024.

608

609 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-

610 nition, 2015. URL <https://arxiv.org/abs/1512.03385>.

611 James Hinns, Xiuyi Fan, Siyuan Liu, Veera Raghava Kovvuri, Mehmet Orcun Yalcin, and

612 Markus Roggenbach. An initial study of machine learning underspecification using feature at-

613 tribution explainable ai algorithms: A covid-19 virus transmission case study. In *PRICAI 2021:*

614 *Trends in Artificial Intelligence: 18th Pacific Rim International Conference on Artificial Intelli-*

615 *gence, PRICAI 2021, Hanoi, Vietnam, November 8–12, 2021, Proceedings, Part I 18*, pp. 323–

616 335, Vietnam, 2021. Springer.

617 Xinyue Hu, Zhibin Duan, Bo Chen, and Mingyuan Zhou. Enhancing uncertainty estimation and

618 interpretability via bayesian non-negative decision layer. *arXiv preprint arXiv:2505.22199*, 2025.

619

620 Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning:

621 An introduction to concepts and methods. *Machine learning*, 110(3):457–506, 2021.

622 Pascal Iversen, Simon Witzke, Katharina Baum, and Bernhard Y. Renard. Identifying Drivers of

623 Predictive Aleatoric Uncertainty, 2024. URL <http://arxiv.org/abs/2312.07252>.

624

625 Laurent Valentin Jospin, Hamid Laga, Farid Boussaid, Wray Buntine, and Mohammed Bennamoun.

626 Hands-On Bayesian Neural Networks—A Tutorial for Deep Learning Users. *IEEE Computa-*

627 *tional Intelligence Magazine*, 17(2):29–48, May 2022. ISSN 1556-6048. doi: 10.1109/MCI.

628 2022.3155327.

629 Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L

630 Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diag-

631 noses and treatable diseases by image-based deep learning. *cell*, 172(5):1122–1131, 2018.

632 Lennard Korte, Li Rong Wang, and Xiuyi Fan. Confidence estimation in analyzing intravascular

633 optical coherence tomography images with deep neural networks. In *2024 IEEE Conference on*

634 *Artificial Intelligence (CAI)*, pp. 358–364, Singapore, 2024. IEEE.

635

636 Veera Raghava Reddy Kovvuri, Hsuan Fu, Xiuyi Fan, and Monika Seisenberger. Fund performance

637 evaluation with explainable artificial intelligence. *Finance Research Letters*, 58:104419, 2023.

638 Mikhail Kulyabin, Aleksei Zhdanov, Anastasia Nikiforova, Andrey Stepichev, Anna Kuznetsova,

639 Mikhail Ronkin, Vasilii Borisov, Alexander Bogachev, Sergey Korotkikh, Paul A Constable, et al.

640 Octdl: Optical coherence tomography dataset for image-based deep learning methods. *Scientific*

641 *data*, 11(1):365, 2024.

642 Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and Scalable Predictive

643 Uncertainty Estimation using Deep Ensembles. In *Advances in Neural Information Processing*

644 *Systems*, volume 30, pp. 6402–6413, USA, 2017. Curran Associates, Inc.

645

646 Lei Le, Andrew Patterson, and Martha White. Supervised autoencoders: Improving generalization

647 performance with unsupervised regularizers. *Advances in neural information processing systems*,

31, 2018.

648 Peter K. Lee. *Actinic Keratosis, Basal, and Squamous Cell Carcinoma*, chapter 17. McGraw-Hill  
 649 Education, New York, NY, 2017. URL [accessmedicine.mhmedical.com/content.aspx?aid=1177005687](http://accessmedicine.mhmedical.com/content.aspx?aid=1177005687).  
 650

651

652 Scott Lundberg and Su-In Lee. A unified approach to interpreting model predictions, 2017. URL  
 653 <https://arxiv.org/abs/1705.07874>.  
 654

655 Andrey Malinin and Mark Gales. Predictive Uncertainty Estimation via Prior Net-  
 656 works. In *Advances in Neural Information Processing Systems*, volume 31. Curran Asso-  
 657 ciates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/3ea2db50e62ceefceaf70a9d9a56a6f4-Abstract.html>.  
 658

659 Lu Mi, Hao Wang, Yonglong Tian, Hao He, and Nir N. Shavit. Training-Free Uncertainty Estimation  
 660 for Dense Regression: Sensitivity as a Surrogate. *Proceedings of AAAI*, 36(9):10042–10050,  
 661 2022. ISSN 2374-3468. doi: 10.1609/aaai.v36i9.21243.  
 662

663 Carlos Mougan and Dan Saattrup Nielsen. Monitoring Model Deterioration with Explainable Uncer-  
 664 tainty Estimation via Non-parametric Bootstrap. *Proceedings of the AAAI Conference on Artifi-  
 665 cial Intelligence*, 37(12):15037–15045, 2023. ISSN 2374-3468. doi: 10.1609/aaai.v37i12.26755.  
 666 URL <https://ojs.aaai.org/index.php/AAAI/article/view/26755>.  
 667

668 Mohammed Mustafa. Diabetes prediction dataset, version 1. <https://www.kaggle.com/datasets/iammustafatz/diabetes-prediction-dataset/data>, 2023.  
 669

670 Stephen Odaibo. Tutorial: Deriving the Standard Variational Autoencoder (VAE) Loss Function,  
 671 2019. URL <https://arxiv.org/abs/1907.08956>.  
 672

673 Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts.  
 674 *Distill*, 1(10):e3, 2016.  
 675

676 Janis Postels, Hermann Blum, Cesar Cadena, Roland Siegwart, Luc Van Gool, and Federico  
 677 Tombari. Quantifying aleatoric and epistemic uncertainty using density estimation in latent space.  
 678 *CoRR*, abs/2012.03082, 2020. URL <https://arxiv.org/abs/2012.03082>.  
 679

680 Janis Postels, Mattia Segù, Tao Sun, Luca Daniel Sieber, Luc Van Gool, Fisher Yu, and Federico  
 681 Tombari. On the practicality of deterministic epistemic uncertainty. In Kamalika Chaudhuri,  
 682 Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International  
 683 Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, vol-  
 684 ume 162 of *Proceedings of Machine Learning Research*, pp. 17870–17909, USA, 2022. PMLR.  
 685 URL <https://proceedings.mlr.press/v162/postels22a.html>.  
 686

687 Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential Deep Learning to Quantify Classifi-  
 688 cation Uncertainty. In *Advances in Neural Information Processing Systems*, volume 31, Canada,  
 689 2018. Curran Associates, Inc.  
 690

691 C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27  
 692 (3):379–423, 1948. ISSN 0005-8580. doi: 10.1002/j.1538-7305.1948.tb01338.x. URL <https://ieeexplore.ieee.org/abstract/document/6773024>.  
 693

694 Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learn-  
 695 ing. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob  
 696 Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Infor-  
 697 mation Processing Systems 30: Annual Conference on Neural Information Processing Systems  
 698 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 4077–4087, USA, 2017. Curran  
 699 Associates, Inc. URL <https://proceedings.neurips.cc/paper/2017/hash/cb8da6767461f2812ae4290eac7cbc42-Abstract.html>.  
 700

701 Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In  
 702 *International conference on machine learning*, pp. 3319–3328, Australia, 2017. PMLR.

702 Sing Yee Toh, Chang Cai, Li Rong Wang, Xiaoyin Bai, Joanne Ngeow, and Xiuyi Fan. The effect  
 703 of explainable ai and uncertainty quantification on medical students' perspectives of decision-  
 704 making ai: A cancer screening case study. In *Proceedings of the Extended Abstracts of the CHI*  
 705 *Conference on Human Factors in Computing Systems*, CHI EA '25, New York, NY, USA, 2025.  
 706 Association for Computing Machinery. ISBN 9798400713958. doi: 10.1145/3706599.3719791.  
 707 URL <https://doi.org/10.1145/3706599.3719791>.

708 Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The HAM10000 dataset, a large collection  
 709 of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 5(1):  
 710 180161, August 2018. ISSN 2052-4463. doi: 10.1038/sdata.2018.161.

711 Dennis Ulmer, Christian Hardmeier, and Jes Frellsen. Prior and Posterior Networks: A Survey  
 712 on Evidential Deep Learning Methods For Uncertainty Estimation. *Transactions on Machine*  
 713 *Learning Research*, 2023, 2023.

714 Uddesha Upadhyay, Shyamgopal Karthik, Yanbei Chen, Massimiliano Mancini, and Zeynep  
 715 Akata. BayesCap: Bayesian Identity Cap for Calibrated Uncertainty in Frozen Neural Networks.  
 716 In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner  
 717 (eds.), *Computer Vision – ECCV 2022*, pp. 299–317, Cham, 2022. Springer Nature Switzerland.  
 718 ISBN 978-3-031-19775-8. doi: 10.1007/978-3-031-19775-8\_18.

719 Hanjing Wang, Bashirul Azam Biswas, and Qiang Ji. Optimization-Based Uncertainty Attribution  
 720 Via Learning Informative Perturbations. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Rus-  
 721 sakovsky, Torsten Sattler, and Gü̈l Varol (eds.), *Computer Vision – ECCV 2024*, volume 15136,  
 722 pp. 237–253. Springer Nature Switzerland, Switzerland, 2025. ISBN 978-3-031-73228-7 978-3-  
 723 031-73229-4. doi: 10.1007/978-3-031-73229-4\_14. URL [https://link.springer.com/10.1007/978-3-031-73229-4\\_14](https://link.springer.com/10.1007/978-3-031-73229-4_14).

724 Li Rong Wang, Thomas C. Henderson, and Xiuyi Fan. An Uncertainty Estimation Model for Al-  
 725 gorithmic Trading Agent. In Soon-Geul Lee, Jinung An, Nak Young Chong, Marcus Strand,  
 726 and Joo H. Kim (eds.), *Intelligent Autonomous Systems 18*, pp. 459–465, Korea, 2024. Springer  
 727 Nature Switzerland. ISBN 978-3-031-44981-9. doi: 10.1007/978-3-031-44981-9\_39.

728 Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammad Bagheri, and Ronald M Sum-  
 729 mers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised  
 730 classification and localization of common thorax diseases. In *Proceedings of the IEEE conference*  
 731 *on computer vision and pattern recognition*, pp. 2097–2106, 2017.

732 David Watson, Joshua O'Hara, Niek Tax, Richard Mudd, and Ido Guy. Explaining predictive un-  
 733 certainty with information theoretic shapley values. *Advances in Neural Information Processing*  
 734 *Systems*, 36:7330–7350, 2023.

735 Claudio Zelenka, Andrea Göhring, Daniyal Kazempour, Maximilian Hünemörder, Lars Schmarje,  
 736 and Peer Kröger. A simple and explainable method for uncertainty estimation using attribute  
 737 prototype networks. In *Proceedings of the IEEE/CVF International Conference on Computer*  
 738 *Vision*, pp. 4570–4579, 2023a.

739 Claudio Zelenka, Andrea Göhring, Daniyal Kazempour, Maximilian Hünemörder, Lars Schmarje,  
 740 and Peer Kröger. A Simple and Explainable Method for Uncertainty Estimation using At-  
 741 tribute Prototype Networks. In *2023 IEEE/CVF International Conference on Computer Vi-  
 742 sion Workshops (ICCVW)*, pp. 4572–4581, France, 2023b. IEEE. ISBN 979-8-3503-0744-  
 743 3. doi: 10.1109/ICCVW60793.2023.00491. URL [https://ieeexplore.ieee.org/](https://ieeexplore.ieee.org/document/10350670/)  
 744 document/10350670/.

745

746

747

748

749 A APPENDIX

750

751 A.1 DATASET DETAILS

752

753 A.1.1 COVID-19 VIRUS TRANSMISSION DATASET

754

755 (Hinns et al., 2021) is a tabular binary classification dataset comprising 3,553 instances. Each in-  
 756 stance is characterized by 32 continuous features describing the United Kingdom's COVID-19 poli-

756 cies and daily case counts across 12 regions, covering the period from February 2020 to February  
 757 2021. The dataset is labelled with a binary class indicating whether  $R_t > 1$ , where  $R_t$  represents the  
 758 effective reproduction number. A value of  $R_t > 1$  signifies an increasing spread of COVID-19. We  
 759 randomly divided the dataset into three subsets: training (70%), validation (10%), and test (20%).  
 760

### 761 A.1.2 DIABETES DIAGNOSIS DATASET

763 (Mustafa, 2023) is a tabular binary classification dataset comprising 100,000 instances, each repre-  
 764 senting a patient. Each instance is described by eight features detailing the patient’s demographics  
 765 (e.g., age and gender), pre-existing conditions (e.g., prevalence of heart disease), and vital signs  
 766 (e.g., blood glucose level), and has a binary label indicating whether the patient suffers from dia-  
 767 betes (0 indicating the patient is diabetes-free and 1 indicating the patient suffers from diabetes). We  
 768 divided the dataset into three subsets for our experiments: training (80%), validation (10%), and test  
 769 (10%).  
 770

### 771 A.1.3 FUND PERFORMANCE EVALUATION DATASET

772 (Kovvuri et al., 2023) is a tabular binary classification dataset comprising 77,940 instances, designed  
 773 to predict whether a fund’s net asset value (NAV) exceeds its NAV from the previous quarter. Each  
 774 instance represents the state of one of 4,330 funds between March 2017 and June 2021, sampled  
 775 quarterly. Each instance is characterized by 18 continuous features, including macroeconomic in-  
 776 dicators (such as stock market returns (ST), exchange rate returns (EXR), and interest rates (IR)),  
 777 country-level equity investment percentages and net asset data for the fund (covering 11 countries  
 778 and “Other Country”), the Herfindahl–Hirschman index (HHI), and a past performance metric com-  
 779 puted as the sum of class labels from the past four quarters (“L4f Gain”). We split the dataset into  
 780 three subsets based on date: instances before March 2020 were assigned to the training set (66.7%),  
 781 instances between March 2020 and June 2020 to the validation set (11.1%), and instances after June  
 782 2020 to the test set (22.2%).  
 783

### 784 A.1.4 ISIC

785 (Codella et al., 2019; Tschandl et al., 2018) is a dermoscopic image dataset containing instances from  
 786 seven skin conditions: melanoma (MEL), melanocytic nevus (NV), basal cell carcinoma (BCC), ac-  
 787 tinic keratosis (AKIEC), benign keratosis (BKL), dermatofibroma (DF), and vascular lesion (VASC).  
 788 We followed the 2018 ISIC Challenge split, with 10,015 training, 193 validation, and 1,512 test im-  
 789 ages. For the Out-of-Distribution (OOD) detection experiments, we constructed three OOD datasets  
 790 with decreasing similarity to ISIC: (1) BCN-IN: 1,512 dermoscopic images of seen classes from the  
 791 BCN20000 dataset (Combalia et al., 2019), (2) BCN-OUT: 313 images from the unseen class Scar,  
 792 and (3) ChestMNIST: 1,512 images from the ChestMNIST dataset (Wang et al., 2017).  
 793

### 794 A.1.5 OCTMNIST

795 (Kermany et al., 2018) is a retinal OCT image dataset comprising four classes: 47% Normal, 34%  
 796 CNV (Choroidal Neovascularization), 11% DME (Diabetic Macular Edema), and 8% Drusen. We  
 797 split the dataset into three subsets: 97,477 training, 10,832 validation, and 1,000 test instances. For  
 798 the OOD detection experiments, we constructed three OOD datasets with decreasing similarity to  
 799 OCTMNIST: (1) **OCTDL-IN**: 618 OCTDL (Kulyabin et al., 2024) images from seen classes, (2)  
 800 **OCTDL-OUT**: 1,000 OCTDL images from unseen classes, and (3) **ChestMNIST**: 1,512 images  
 801 from the ChestMNIST dataset (Wang et al., 2017).  
 802

## 803 A.2 BASELINE DETAILS

804 We compared ReconstructionNet with recent state-of-the-art uncertainty estimates.

805 1. **Entropy** (Shannon, 1948) is derived from prediction probability.  
 806 2. **Monte Carlo Dropout (MCD)** (Gal & Ghahramani, 2016) We apply dropout ( $p = 0.5$ )  
 807 to the penultimate layer of the model and keep dropout active during inference to yield  $T$   
 808 predictions (where  $T = 100$  for the Covid and Fund datasets, and  $T = 10$  for all other  
 809 datasets), following the hyperparameters in (Gal & Ghahramani, 2016).

810

811 3. **Deep Ensemble (DE)** (Lakshminarayanan et al., 2017) uses five models, with the standard  
812 deviation of their predictions as the uncertainty estimate.

813 4. **Posterior Network (PN)** (Charpentier et al., 2020) computes aleatoric uncertainty as the  
814 inverse of the maximum prediction probability and epistemic uncertainty as the inverse of  
815 the maximum of the Dirichlet distribution parameters.

816 5. **Bayesian Neural Networks (BNN)** (Jospin et al., 2022) estimate uncertainty as the stan-  
817 dard deviation of  $T = 100$  predictions obtained with different weight samples.

818 6. **Evidential Deep Learning (EDL)** (Sensoy et al., 2018) estimates uncertainty using the  
819 entropy of the predicted probabilities.

820

821 **A.3 METRIC DETAILS**

822

823 **A.3.1 CORRELATION**

824

825 (Mi et al., 2022; Upadhyay et al., 2022) measures the reliability of the uncertainty estimate as the  
826 Pearson’s correlation coefficient between the uncertainty estimate and the prediction error (mea-  
827 sured as the absolute difference between the one-hot label and the predicted probability). A higher  
828 correlation indicates a more reliable estimate, as prediction errors are likely to be high when the  
829 model’s uncertainty is high, and vice versa.

830

831 **A.3.2 AURC (AREA UNDER RISK-COVERAGE CURVE)**

832

833 (Ding et al., 2020) quantifies the selectivity of the uncertainty estimate, indicating its usefulness in  
834 selective prediction i.e. making predictions only for confident instances. This is computed as the  
835 area under the risk-coverage curve, which plots the 0/1 loss (Risk) for instances with uncertainty  
836 scores within the  $\alpha$ %-percentile (Coverage) against the coverage. A selective uncertainty estimate  
837 would yield a low AURC, as it effectively reduces loss across all possible uncertainty thresholds.

838

839 **A.3.3 SIGMA-RISK SCORE**

840

841 Evaluates the resilience of the uncertainty estimate to false negatives (incorrect instances with low  
842 uncertainty), which can lead to significant costs if many are left undetected. It is computed as the  
843 0/1 loss of instances with normalized uncertainty less than  $\sigma = \{0.1, 0.2, 0.3, 0.4\}$ . A lower  $\sigma$ -risk  
844 score indicates greater robustness to false negatives. To ensure robustness to outliers, we apply min-  
845 max normalization with the minimum and maximum values computed after outlier removal with the  
846 interquartile range method.

847

848 **A.3.4 OOD DETECTION**

849

850 (Lakshminarayanan et al., 2017; Postels et al., 2020; Malinin & Gales, 2018) Evaluates how effec-  
851 tively the uncertainty estimate can distinguish between in-distribution and OOD instances. This is  
852 quantified using AUROC (Area Under the Receiver Operating Characteristic Curve), where the true  
853 label indicates whether an instance is OOD and the target score is the uncertainty estimate. A higher  
854 AUROC reflects a more reliable uncertainty measure capable of identifying OOD inputs.

855

856 **A.4 IMPLEMENTATION DETAILS**

857

858 The use of error weights  $\omega$  allows the model to decouple the modeling of joint prediction proba-  
859 bilities from the classification task, enabling simultaneous optimization of both objectives. It also  
860 dynamically scales errors across classes to address variations in reconstruction difficulty and adjusts  
861 the contribution of each feature to the prediction, acknowledging that not all reconstruction errors  
862 equally impact the final outcome.

863 For image datasets, to reduce the size of the ReconstructionNet model, we trained class-specific  
864 decoders sharing a common encoder instead of training separate class-specific autoencoders. The  
865 encoder was based on a ResNet18 (He et al., 2015) backbone, and the decoders were implemented  
866 as inverted ResNet18 backbones, replacing convolutional layers with transposed convolutions.

864 Hyperparameters for all models were determined using grid search on the validation set. For Re-  
 865 constructionNet, we tuned the compression ratio (which determines the latent vector length as its  
 866 product with the feature count) along with the number of encoder and decoder layers, the width  
 867 of intermediate layers, and the loss function parameter  $\beta$  ( $\beta = 1.25, 2, 0.75$  for Covid, Diabetes,  
 868 and Fund datasets). Both the MLP and ReconstructionNet models were trained on the training set  
 869 with oversampling using SMOTE (Chawla et al., 2002). All models were trained using the Adam  
 870 optimizer with early stopping. Each experiment was repeated five times, and we report the mean  
 871 and standard deviation of each metric.

872

## A.5 VERIFYING UNCERTAINTY EXPLANATION PROPERTIES WITH TOY EXAMPLES

873

## A.5.1 IMPLEMENTATION INVARIANCE

874

875 We illustrate the property of implementation invariance with a simple example. Suppose we have  
 876 two ReconstructionNet models,  $f$  and  $f'$ , that differ in architecture:  
 877

878

- $f$ : Each class autoencoder is shallow and linear;
- $f'$ : Each class autoencoder is deeper with nonlinearities.

879

If for every input  $\mathbf{x}$ , the two ReconstructionNet models yield identical prediction probabilities  $\hat{p}$  for  
 880 all classes, then:

881

1. The predicted classes  $c^* = \arg \max_j \hat{p}_j$  are identical,
2. The weighted reconstruction errors  $\epsilon$  are identical,
3. The uncertainty explanations  $\zeta(\mathbf{x})$ , which are the feature-wise weighted reconstruction er-  
 882 rors of the predicted classes, are identical.

883

Thus, ReconstructionNet’s uncertainty explanations are *implementation-invariant* under this defini-  
 884 tion of functional equivalence.

885

## A.5.2 SENSITIVITY

886

We illustrate the sensitivity property with a simple example. Consider a dataset where each input  $\mathbf{x}$   
 887 has three features,  $(x^1, x^2, x^3)$ . Suppose that for the predicted class  $c^*$ :

888

- The weighted reconstruction errors depend only on  $x^1$  and  $x^2$ ;
- The weight for  $x^3$  is zero:  $w_{c^*}^3 = 0$ , indicating that changes in  $x^3$  do not affect the class  
 889 prediction.

890

The uncertainty explanation for input  $\mathbf{x}$  is given by:

891

$$\zeta(\mathbf{x}) = [w_{c^*}^1 \|x^1 - \hat{x}_{c^*}^1\|^2 \quad w_{c^*}^2 \|x^2 - \hat{x}_{c^*}^2\|^2 \quad w_{c^*}^3 \|x^3 - \hat{x}_{c^*}^3\|^2] \quad (15)$$

892

Since  $w_{c^*}^3 = 0$ , the explanation assigns zero attribution to  $x^3$ , which is irrelevant to the predic-  
 893 tion. Hence, ReconstructionNet’s uncertainty explanations satisfy the *sensitivity* property: features  
 894 irrelevant to the prediction receive zero attribution in the weighted reconstruction error.

895

## A.5.3 CONSISTENCY

896

We illustrate the consistency property with a simple example. Consider an input  $\mathbf{x} = (x^1, x^2, x^3)$ ;  
 897 For the predicted class  $c^*$ , the uncertainty explanation is:

898

$$\zeta(\mathbf{x}) = [w_{c^*}^1 \|x^1 - \hat{x}_{c^*}^1\|^2 \quad w_{c^*}^2 \|x^2 - \hat{x}_{c^*}^2\|^2 \quad w_{c^*}^3 \|x^3 - \hat{x}_{c^*}^3\|^2] \quad (16)$$

899

Now consider a modification  $f'$  of the model  $f$  where feature  $x^2$  is made more important for the  
 900 prediction (i.e., its weight increases in the softmax over reconstruction errors).

901

- The reconstruction errors  $\|x^j - \hat{x}_{c^*}^j\|^2$  remain unchanged, since the autoencoders are un-  
 902 changed.
- The weight for  $x^2$  increases:  $w_{c^*}^{2'} > w_{c^*}^2$ .

903

918 Then the uncertainty attribution for  $x_2$  becomes:  
 919

$$\zeta(\mathbf{x})^{2'} = w_{c^*}^{2'} \|x_2 - \hat{x}_{c^*}^{2'}\|^2 \geq w_{c^*}^2 \|x^2 - \hat{x}_{c^*}^2\|^2 = \zeta(\mathbf{x})^2. \quad (17)$$

920 Thus, increasing a feature’s contribution to the prediction does not decrease its uncertainty attribution-  
 921 tion, demonstrating that ReconstructionNet’s uncertainty explanations satisfy the *consistency* prop-  
 922 erty.  
 923

## 924 A.6 ADDITIONAL UNCERTAINTY EXPLANATION EXAMPLES

925 **MNIST:** The examples illustrate how the model’s uncertainty explanations correspond to ambiguous  
 926 digit features that resemble other classes. In Figure 4a, the model predicted a 2 instead of a 0,  
 927 showing uncertainty around the missing connection in the 0. In Figure 4b, it predicted a 6 instead of  
 928 a 0, with uncertainty focused on the tail of the 6. In Figure 4c, the model again predicted a 6 instead  
 929 of a 0, highlighting uncertainty on the right curve that resembles a 0. In Figure 4d, it predicted a  
 930 9 instead of a 4, with uncertainty at the connecting bump of the 4 (which is atypical in a 9). In  
 931 Figure 4e, the model predicted a 5, with uncertainty concentrated on the top tail of the digit. Finally,  
 932 in Figure 4f, it predicted a 3 instead of an 8, highlighting the lower-left connecting curve, whose  
 933 removal would make the digit resemble a 3 more strongly.  
 934

935 **ISIC:** The examples demonstrate how the model’s uncertainty explanations capture visual features  
 936 that increase ambiguity in lesion classification, leading to misclassifications. In Figures 4g and 4h,  
 937 the model misdiagnoses the dermoscopic image, with the uncertainty explanations highlighting a  
 938 scab-like region. In Figures 4i, 4j and 4k, images were misclassified as MEL (melanoma) instead of  
 939 BCC or AKIEC, with uncertainty concentrated on the uneven pigmentation of the lesion, a feature  
 940 often associated with melanoma. In Figure 4l, the image was misclassified as AKIEC with uncer-  
 941 tainty explanations highlighting the presence of hairs covering a substantial portion of the lesion.  
 942

943 **OCTMNIST** These cases illustrate how the model’s uncertainty highlights regions that contribute  
 944 to misclassification in OCT images. In Figure 4m, the model classified the OCT image as DME,  
 945 with uncertainty concentrated around one of the fluid pockets. In Figure 4n, the model predicted  
 946 CNV instead of DME, again highlighting a fluid pocket. This reflects the ambiguity introduced by  
 947 fluid pockets, which are common to both conditions. Figures 4o, 4p, 4q, and 4r are Drusen images  
 948 misclassified as other conditions, with uncertainty explanations highlighting the deposits (uneven  
 949 bumps in the OCT images), which are key features of Drusen and major contributors to the model’s  
 950 uncertainty in each incorrect prediction.  
 951

## 952 A.7 COMPARISON WITH LLMs

953 Figure 5 compares uncertainty explanations from ReconstructionNet and a publicly available, off-  
 954 the-shelf large language model (LLM; Microsoft Copilot). We observe that ReconstructionNet pro-  
 955 duces more selective explanations, highlighting only small regions that contribute to uncertainty,  
 956 whereas Copilot often highlights large areas (notably, in Figure 5b the entire digit is highlighted).  
 957 Moreover, the regions identified by ReconstructionNet are more accurate. For example, in Figure 5a,  
 958 where the digit 4 is misclassified as 7, ReconstructionNet correctly highlights the additional vertical  
 959 line on the right, while Copilot highlights the connecting region that is common to both 4 and 7  
 960 and should not confuse the model. Similarly, in Figure 5c, where a 6 is misclassified as 5, Recon-  
 961 structionNet pinpoints the extra vertical line on the bottom left that distinguishes 5 from 6, whereas  
 962 Copilot highlights the top-left vertical line, which can occur in both digits.  
 963

## 964 A.8 RELATION TO CONTRASTIVE LEARNING AND SIAMESE NETWORKS

965 Note that, although the training procedure of ReconstructionNet shares similarities with contrastive  
 966 learning (Chopra et al., 2005), as both aim to minimize the distance between instances of the same  
 967 class while maximizing the distance between instances of different classes, there are key differences  
 968 in how these distances are computed. In contrastive learning, the distance is explicitly calculated  
 969 as the L2-norm between the embeddings of a pair of instances. In ReconstructionNet, the dis-  
 970 tance between an instance and the instances of a particular class is represented as the reconstruc-  
 971 tion error of the corresponding class autoencoder. This distinction in distance measurement leads to  
 972 differences in the training loss. While contrastive learning incorporates the L2-norm distance be-  
 973 tween pairs of instances, ReconstructionNet uses class reconstruction error for its loss function.  
 974

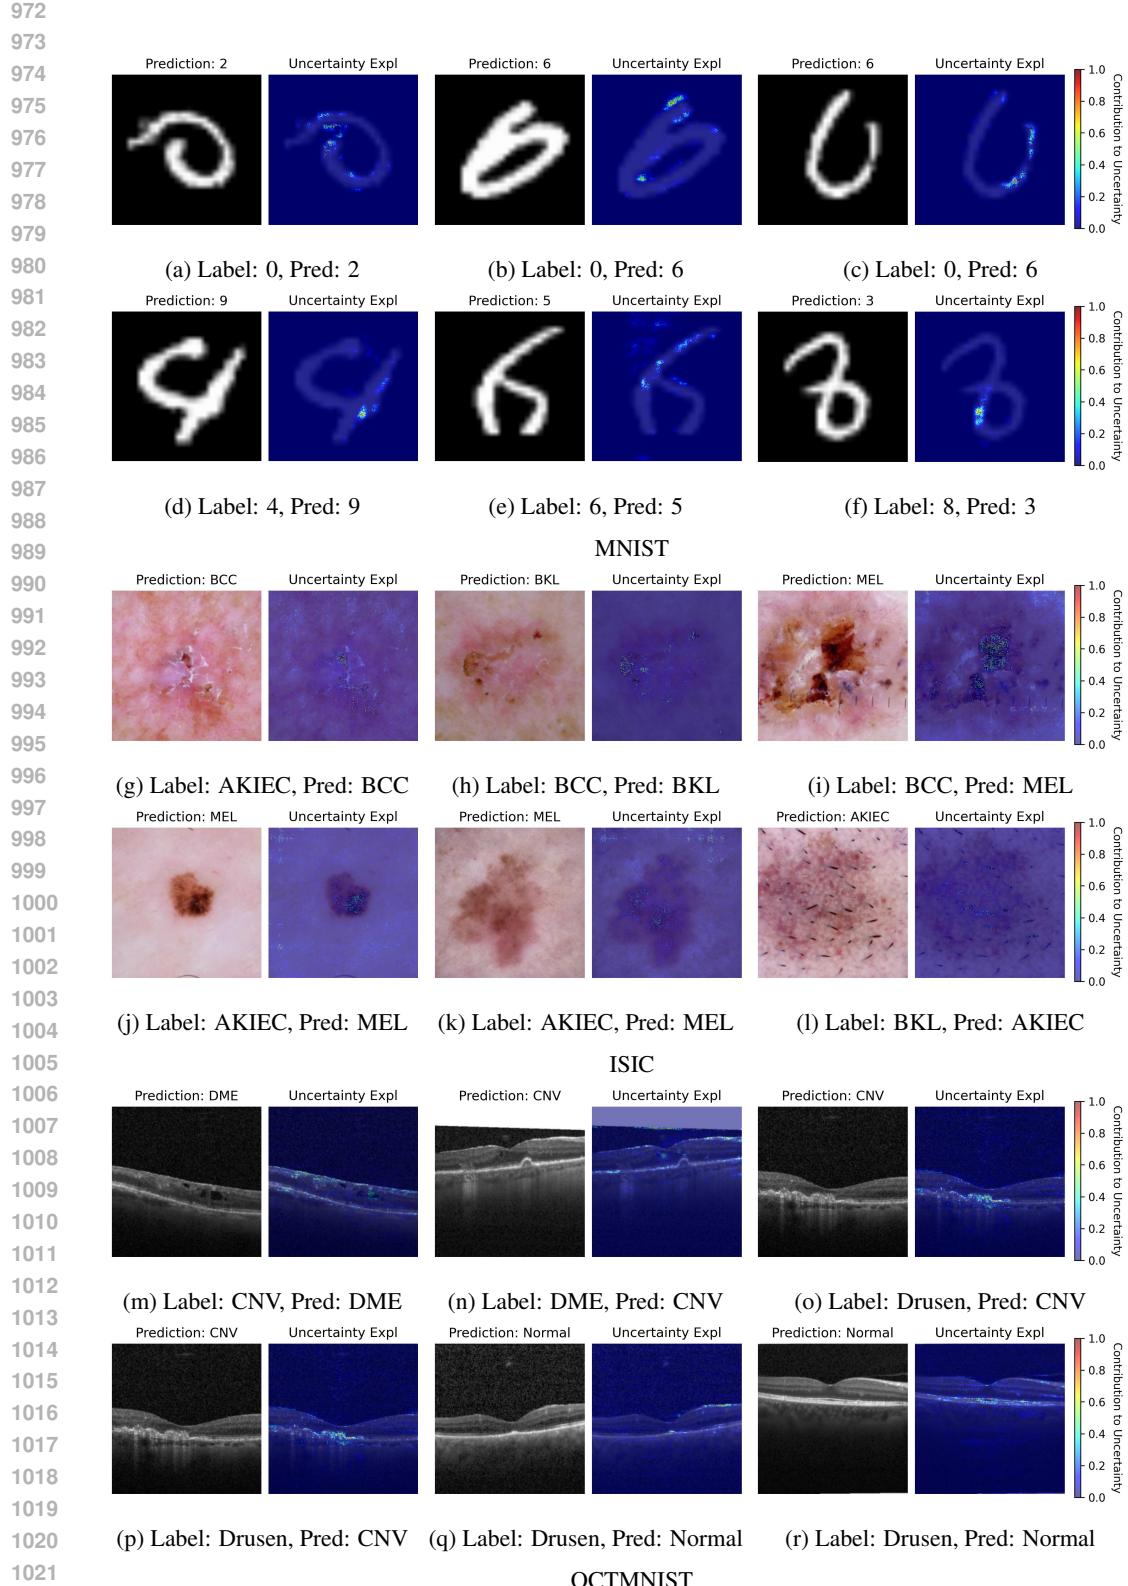


Figure 4: Uncertainty explanation illustration for MNIST, ISIC and OCTMNIST datasets. Positive uncertainty explanations were min-max normalized, gamma-corrected, and overlaid on the images for clarity. The highlighted regions “explain” the prediction uncertainty.

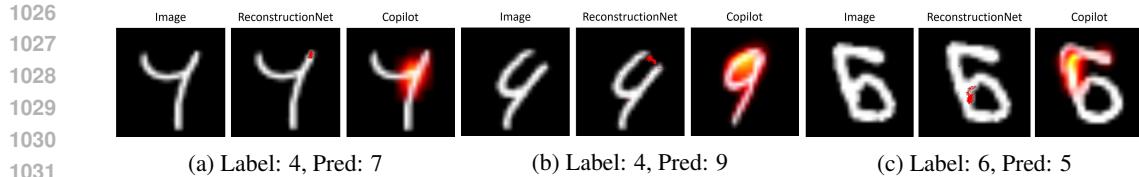


Figure 5: Comparison of uncertainty explanations from ReconstructionNet and Microsoft Copilot on the MNIST dataset. Positive uncertainty explanations from ReconstructionNet were min–max normalized, with pixels having attribution greater than 0.2 shown in red and overlaid on the images for clarity. Microsoft Copilot explanations were generated using the prompt: “This image has been misclassified as [image class prediction]. Generate an image highlighting the regions contributing to prediction uncertainty in red.” The highlighted regions are intended to “explain” the prediction uncertainty.

Furthermore, in contrastive learning, the distances between instances of different classes are explicitly maximized and included in the loss function. In contrast, ReconstructionNet exclusively trains each class autoencoder with instances belonging to its specific class. The architecture of ReconstructionNet shares similarities with Siamese networks (Bromley et al., 1993) in its use of sub-networks. However, unlike Siamese networks, which share weights across sub-networks, ReconstructionNet allows each class autoencoder to have distinct weights.