

Evidential Concept Embedding Models: Towards Reliable Concept Explanations for Skin Disease Diagnosis

Yibo Gao, Zheyao Gao, Xin Gao, Yuanye Liu, Bomin Wang,
Xiahai Zhuang

Fudan University
<https://zmiclalab.github.io>

Abstract. Due to the high stakes in medical decision-making, there is a compelling demand for interpretable deep learning methods in medical image analysis. Concept Bottleneck Models (CBM) have emerged as an active interpretable framework incorporating human-interpretable concepts into decision-making. However, their concept predictions may lack reliability when applied to clinical diagnosis, impeding concept explanations' quality. To address this, we propose an evidential Concept Embedding Model (evi-CEM), which employs evidential learning to model the concept uncertainty. Additionally, we offer to leverage the concept uncertainty to rectify concept misalignments that arise when training CBMs using vision-language models without complete concept supervision. With the proposed methods, we can enhance concept explanations' reliability for both supervised and label-efficient settings. Furthermore, we introduce concept uncertainty for effective test-time intervention. Our evaluation demonstrates that evi-CEM achieves superior performance in terms of concept prediction, and the proposed concept rectification effectively mitigates concept misalignments for label-efficient training. Our code is available at <https://github.com/obiyoag/evi-CEM>.

Keywords: Concept explanations · Skin disease diagnosis · Vision language models.

1 Introduction

Black-box deep learning methods have demonstrated significant efficacy for medical images [9,20]. Nonetheless, the rigorous trustworthiness requirements inherent to the healthcare domain have stimulated research efforts towards explainable methods [22]. Relying solely on explaining black box models will likely sustain problematic clinical practices and lead to adverse medical outcomes [18]. Hence, it is crucial to develop active interpretable approaches [28] for medical images, which are inherently designed to be interpretable [6,5,2].

Xiahai Zhuang is the corresponding author.

This work was funded by the National Natural Science Foundation of China (grant No. 62372115, 61971142 and 62111530195).




Fig. 1: (a) An illustration of CBMs. (b) Over-confident examples of CEMs. The orange and blue bars are used to visually represent the probabilities supporting and opposing the concepts, respectively.

Recently, Concept Bottleneck Models (CBM) [14] have emerged as a promising framework for achieving active interpretability. As shown in Fig. 1a, CBM integrates human-interpretable concepts into the decision-making process [29,7] by decomposing the task into concept prediction and task prediction stages. Images are first mapped to the concepts, and then a classifier is employed to predict task labels based on the concepts. Building upon CBM, Concept Embedding Models (CEM) [4] overcome the trade-off between accuracy and interpretability by learning concept embeddings and achieving improved performance compared to CBM. However, complete concept supervision is challenging to acquire in real-world clinical scenarios. Some works utilize foundation models to train CBMs in a label-efficient manner [25,16,23]. This line of research involves querying useful visual concepts of tasks from large language models like GPT-3 or ChatGPT and then using vision language models (VLM), such as CLIP [17], to assess the presence of these concepts in an image.

CBMs emulate the clinical practice of medical professionals who initially evaluate symptoms before diagnosing diseases, which is promising for interpretable medical image analysis [24]. Nevertheless, their concept predictions may be unreliable for clinical diagnosis, which hinders the quality of concept explanations. The unreliability lies in two aspects. Firstly, high certainty for less obvious concepts can lead to over-confidence in predictions as depicted in Fig. 1b, which may not be accurate during inference. Secondly, in label-efficient training, VLMs can be misaligned with certain concepts [26,27], due to their training paradigm. Therefore, CBMs trained with VLMs suffer from inaccurate concept prediction, which impacts models' reliability and interpretability.

To enhance the reliability of concept explanations, we propose evidential-CEM (evi-CEM) based on evidential deep learning [19]. By modeling concept labels and evidence with Binomial and Beta distributions, we could calibrate the concept predictions and quantify their corresponding uncertainty by em-




Fig. 2: The architecture of Evidential Concept Embedding Model (evi-CEM). The model primarily comprises a backbone network $\Psi(\cdot)$, an evidential concept bottleneck layer (ECBL) and a task predictor $f(\cdot)$.

employing the variational method to minimize the Binomial likelihood. Besides, for label-efficient training, the misaligned concepts with higher uncertainty could be identified. Then, a small set of positive and negative examples is gathered for the misaligned concepts to learn concept activation vectors (CAVs) [11]. The concept misalignments could be rectified by incorporating the CAVs into label-efficient training. Furthermore, we suggest employing evi-CEM for uncertainty-aware intervention, wherein uncertainty serves as a metric for selecting concepts that warrant more significant intervention.

Our contributions can be summarized as following:

- We propose a concept-based model named evi-CEM, which utilizes evidential deep learning to mitigate the over-confidence problem for concept prediction.
- With concept uncertainty, we alleviate concept misalignments that arise in the label-efficient training of CBMs by concept rectification.
- We introduce uncertainty-aware intervention by incorporating uncertainty into the process to improve the intervention efficiency.
- We evaluate the effectiveness of the proposed methods on a public dataset. The experimental results demonstrate that our approach improves the quality of concept explanations for skin disease diagnosis.

2 Method

2.1 Evidential Concept Embedding Model

Model architecture: Given an image $\mathbf{x}$ and a set of K target concepts $\mathcal{C} = \{C_1, C_2, \dots, C_K\}$, the goal of evi-CEM is to predict the true label of target concepts $\mathbf{c} = [c_1, c_2, \dots, c_K]^T$ and provide an accurate diagnosis $\mathbf{y}$. As depicted in Fig. 2, a backbone network is firstly employed to extract the image feature

$\mathbf{h} = \Psi(\mathbf{x})$, which is then fed into the evidential concept bottleneck layer (ECBL). In ECBL, following [4], each concept C_k are associated with two embeddings generated by two linear layers $\mathbf{c}_k^+ = \Phi_k^+(\mathbf{h})$ (indicating the presence of C_k) and $\mathbf{c}_k^- = \Phi_k^-(\mathbf{h})$ (indicating the absence of C_k). The positive and negative evidence of concept C_k are derived using two functions:

$$\begin{aligned}\alpha_k &= \alpha([\mathbf{c}_k^+, \mathbf{c}_k^-]^T) = \text{ReLU}(W_\alpha[\mathbf{c}_k^+, \mathbf{c}_k^-]^T + \mathbf{b}_\alpha) + 1, \\ \beta_k &= \beta([\mathbf{c}_k^+, \mathbf{c}_k^-]^T) = \text{ReLU}(W_\beta[\mathbf{c}_k^+, \mathbf{c}_k^-]^T + \mathbf{b}_\beta) + 1,\end{aligned}$$

where W and $\mathbf{b}$ represent the weight and bias of linear layers. Then, the concept embedding $\mathbf{c}_k$ can be constructed as

$$\mathbf{c}_k = \frac{\alpha_k}{\alpha_k + \beta_k} \mathbf{c}_k^+ + \frac{\beta_k}{\alpha_k + \beta_k} \mathbf{c}_k^-.$$

Finally, the concatenation of the different concept embeddings is fed into the task predictor to obtain the final prediction $\mathbf{y} = f([\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K]^T)$.

Uncertainty modelling: According to the principles of subjective logic theory [10], each concept C_k is characterized by a binomial opinion represented as $\omega_k = (b_k, d_k, u_k, a_k)$. This binomial opinion encompasses belief mass (b_k), disbelief mass (d_k), uncertainty mass (u_k), and the base rate (prior probability of C_k without any evidence) denoted as a_k . It satisfies the additivity requirement $b_k + d_k + u_k = 1$, and the projected probability is calculated as $p_k = b_k + a_k \cdot u_k$. The binomial opinion ω_k is equivalent to a Beta distribution under the bijective mapping: $p_k \sim \text{Beta}(\alpha_k, \beta_k)$, where α_k and β_k are viewed as positive and negative evidence of the concept C_k , respectively. The uncertainty mass can be computed with the evidence as $u_k = 2/(\alpha_k + \beta_k)$.

Training paradigm: Since the concept label c_k follows Binomial distribution $c_k \sim \text{Bin}(c_k|p_k)$, the marginal log likelihood $\log p(c_k|\mathbf{x})$ has an Evidence Lower BOund (ELBO) as

$$\log p(c_k|\mathbf{x}) \geq \mathbb{E}_{q(p_k|\mathbf{x})} [\log p(c_k|p_k)] - \text{KL}(q(p_k|\mathbf{x})||p(p_k|\mathbf{x})),$$

where $q(p_k|\mathbf{x})$ is the variational distribution $\mathcal{B}(\alpha_k, \beta_k)$. Minimizing the negative ELBO, we can derive the variational loss for concept C_k , given by

$$\begin{aligned}\mathcal{L}_{Beta}^k &= \psi(\alpha_k + \beta_k) + c_k \left[\log \beta_k + \frac{1 - \beta_k}{\beta_k} - \psi(\alpha_k) \right] \\ &\quad + (1 - c_k) \left[\log \alpha_k + \frac{1 - \alpha_k}{\alpha_k} - \psi(\beta_k) \right].\end{aligned}$$

Here, $\psi(\cdot)$ denotes the digamma function, and the detailed derivation can be found in the supplementary material. The total loss for training evi-CEM is a combination of the cross-entropy task loss and the variational concept losses:

$$\mathcal{L} = \mathcal{L}_{ce}(\mathbf{y}, \hat{\mathbf{y}}) + \lambda \sum_{k=1}^K \mathcal{L}_{Beta}^k(\alpha_k, \beta_k, c_k), \quad (1)$$




Fig. 3: The process of ECBL pretraining in conjunction with VLMs. The image and text encoders of VLMs remain frozen to extract embeddings for concept estimation using Eq. (2). The ECBL is trained with the estimated concept labels to minimize variational concept loss.

where $\hat{y}$ denotes the ground-truth task label, and λ controls the weight balance between the task and concept losses.

2.2 Concept Rectification for Label-efficient Training

Without concept supervision, many works employ VLMs to assess the presence of the target concepts $\mathcal{C}$ in images. However, the VLMs can be misaligned for certain concepts, undermining the reliability of concept explanations. This section describes the procedure to rectify concept misalignments for VLMs using the proposed ECBL. The entire process consists of three stages.

Pretraining ECBL with VLMs: VLMs have an image and a text encoder to map images and texts into a shared embedding space. As shown in Fig. 3, given an image $\mathbf{x}$, we firstly utilize the image encoder to extract the image embedding, denoted as $\mathbf{i} = f_I(\mathbf{x})$. Then, we leverage the text encoder to compute the concept prompt embedding and the reference prompt embedding, represented as $\mathbf{e}_{kr}^t = f_T(\mathbf{t}_{kr}^t)$ and $\mathbf{e}_r = f_T(\mathbf{t}_r)$ respectively, where $\mathbf{t}_{kr}^t$ and $\mathbf{t}_r$ denotes the tokenized concept prompt and reference prompt for the t -th term and the r -th template respectively. With the computed embeddings, the probability c_k can be estimated as follows:

$$\tilde{c}_k = \frac{1}{T_k \times R} \sum_{t=1}^{T_k} \sum_{r=1}^R \frac{e^{\cos(\mathbf{i}, \mathbf{e}_{kr}^t)/\tau}}{e^{\cos(\mathbf{i}, \mathbf{e}_{kr}^t)/\tau} + e^{\cos(\mathbf{i}, \mathbf{e}_r)/\tau}}, \quad (2)$$

where T_k represents the number of terms associated with concept C_k , R denotes the number of templates, and τ is the temperature parameter. The soft concept label $\tilde{c}_k$ are utilized to pretrain ECBL minimizing $\sum_{k=1}^K \mathcal{L}_{Beta}^k(\tilde{c}_k, \alpha_k, \beta_k)$.

Learning CAVs of the misaligned concepts: After pretraining, we can identify the misaligned concepts, characterized by higher uncertainty, denoted as $\mathcal{C}_m = \{C_k | \forall k, 2/(\alpha_k + \beta_k) \geq \gamma\}$, where γ is the threshold. For each concept $C_k \in \mathcal{C}_m$, we gather embeddings for positive samples $\mathcal{P}_k = \{\Psi^*(\mathbf{x}_{pi}^k)\}_{i=1}^N$ and negative samples $\mathcal{N}_k = \{\Psi^*(\mathbf{x}_{ni}^k)\}_{i=1}^N$, where $\Psi^*(\cdot)$ is the pretrained backbone network. Notably, we only require $2N$ samples associated with the misaligned concepts rather than a densely annotated dataset. Following [11], we train an SVM to separate $\mathcal{P}_k$ and $\mathcal{N}_k$, obtaining the CAV for $C_k \in \mathcal{C}_m$, represented by the normal vector $\mathbf{w}_k$ to the classification boundary.

Incorporating CAVs into evi-CEM training: In the final stage, we incorporate the acquired CAVs into the training of evi-CEM in conjunction with VLMs to mitigate concept misalignment. The CAVs are utilized to adjust the probabilities of misaligned concepts estimated by VLMs, given by

$$c_k = \begin{cases} \tilde{c}_k \times H(\Psi^*(\mathbf{x})^T \mathbf{w}_k), & C_k \in \mathcal{C}_m \\ \tilde{c}_k, & C_k \notin \mathcal{C}_m \end{cases},$$

where $H(\cdot)$ denotes the unit step function. The evi-CEM is then trained with the rectified concept labels $[c_1, c_2, \dots, c_K]$ in accordance with Eq. (1).

2.3 Uncertainty-aware Intervention

Intervention plays a crucial role in CBMs by allowing users to modify intermediate concept predictions to correct the final prediction. The expert-model interactions can enhance the reliability of AI systems, particularly in clinical practice. However, involving senior doctors to review each concept can be expensive. Hence, we propose utilizing evi-CEM for uncertainty-aware intervention, which employs uncertainty as a metric to select concepts for intervention. Evi-CEM predicts the probability and uncertainty $[u_1, u_2, \dots, u_K]$ of target concepts. The uncertainty serves as an indicator, with the concepts exhibiting the highest uncertainty being selected for intervention $k' = \arg \max_k u_k$. Here, k' denotes the index of the target concept to be intervened.

3 Experiments

3.1 Experimental Setup

Dataset: The evaluation is conducted on Fitzpatrick17k (F17k) dataset [8], which consists of 16,523 clinical images used for skin disease diagnosis. The classification task involves distinguishing between malignant, benign, and non-neoplastic conditions. Among the F17k dataset, concept labels of 3,218 images are provided by the SkinCon [3] dataset. SkinCon is currently the most comprehensive dataset in the field of dermatology, containing 48 concepts annotated by board-certified dermatologists. We select 22 concepts (*i.e.*, $K = 22$) from

Table 1: Comparison with CBM variants under complete concept supervision (no VLMs are involved).

Method	Concept Metric			Diagnosis Metric		
	AUC	ACC	F1	AUC	ACC	F1
CBM [14]	69.15 $\pm$ 0.20	78.42 $\pm$ 0.50	54.46 $\pm$ 0.32	76.37 $\pm$ 1.13	78.37 $\pm$ 0.70	56.79 $\pm$ 0.45
CEM [4]	71.00 $\pm$ 0.64	78.96 $\pm$ 0.50	54.95 $\pm$ 0.83	78.63 $\pm$ 0.28	77.07 $\pm$ 1.44	58.02 $\pm$ 0.61
ProbCBM [13]	74.60 $\pm$ 0.87	88.49 $\pm$ 0.97	57.01 $\pm$ 1.51	73.00 $\pm$ 0.17	72.13 $\pm$ 1.00	54.03 $\pm$ 0.15
evi-CEM	80.45 $\pm$ 0.83	90.33 $\pm$ 0.37	64.99 $\pm$ 0.73	77.33 $\pm$ 1.34	76.50 $\pm$ 2.52	58.55 $\pm$ 2.27

the F17k dataset that are represented by at least 50 images for our analysis. Therefore, the F17k dataset is divided into a subset with concept annotations $\mathcal{D}_c$ (3,218 images) and a subset without concept annotations $\mathcal{D}_u$ (13,305 images). We split $\mathcal{D}_c$ into a training set, validation set, and test set according to the proportion of 60%, 20%, and 20%, respectively.

Implementation details: Our model uses ResNet-34 pretrained on the ImageNet dataset as the backbone network $\Psi(\cdot)$. We employ MONET [12] as the VLM, a foundation model pretrained with a massive collection of dermatological images and medical literature for label-efficient training. Regarding the hyperparameters, we use $\lambda = 1$, $\tau = 0.01$, $N = 50$ and $\gamma = 0.6$, and AdamW optimizer with a learning rate of 5×10^{-4} and weight decay of 0.01. The batch size is set to 128. We used PyTorch on NVIDIA Tesla V100 GPU with 16 GB memory for all experiments. We report three metrics (AUC, ACC, F1) for the concept prediction and diagnosis tasks. The concept metrics represent the average results across all concepts. All experiments are conducted using three different random seeds, and we present the mean and standard deviation of the metrics.

3.2 Results

Evi-CEM performance: To showcase the effectiveness of evi-CEM under concept supervision, we compare evi-CEM with other CBM variants on $\mathcal{D}_c$. The comparative methods include CBM [14], CEM [4] and ProbCBM [13]. ProbCBM was proposed to address the concept ambiguity issues by incorporating probabilistic embeddings. As shown in Table 1, the results indicate that both evi-CEM and ProbCBM demonstrate superior performance in concept prediction, which can be attributed to their prediction calibration. Notably, evi-CEM outperforms ProbCBM in both tasks, highlighting the efficacy of evidential modeling, which associates high uncertainty with incorrect predictions. The experimental results demonstrate that while evi-CEM achieves comparable results with CBM and CEM for diagnosis, it can generate more reliable concept explanations.

Concept rectification for label-efficient training: To evaluate the performance of the proposed concept rectification, we compare it with other label-efficient concept prediction methods. CONFES [21] and Reweight [15] are offered

Table 2: Comparison with CBM variants trained in label-efficient manner.

Method	AUC	Concept Metric			AUC	Diagnosis Metric	
		ACC	F1			ACC	F1
CEM+CONFES [21]	73.85 ± 0.29	79.39 ± 0.22	56.73 ± 0.28		89.25 ± 2.55	86.75 ± 0.95	75.92 ± 2.33
CEM+Reweight [15]	73.75 ± 0.43	76.76 ± 1.12	53.83 ± 1.07		89.46 ± 1.84	85.61 ± 1.44	74.88 ± 2.33
CLIP-IP-OMP [1]	65.00 ± 1.62	82.97 ± 0.41	56.91 ± 0.89		77.74 ± 0.39	77.61 ± 0.28	51.64 ± 0.84
LF-CBM [16]	60.45 ± 0.86	81.93 ± 0.69	55.60 ± 0.26		66.16 ± 1.72	75.19 ± 0.21	41.18 ± 0.15
evi-CEM	72.94 ± 0.45	79.13 ± 0.15	56.44 ± 0.19		89.44 ± 1.05	86.49 ± 1.12	77.00 ± 1.27
rectified-evi-CEM	78.55 ± 0.35	83.74 ± 0.12	58.98 ± 0.18		90.11 ± 0.98	86.91 ± 0.50	77.53 ± 0.50




Fig. 4: (a) A qualitative result of uncertainty-aware intervention. $\hat{c}$ and u denotes the concept prediction and its uncertainty respectively. (b) Test diagnosis AUC of interventions on evi-CEM and their corresponding standard deviations.

for noisy learning, which is adapted for concept prediction with VLMs in this study. CLIP-IP-OMP [1] utilizes orthogonal matching pursuit and VLMs to select the most informative concepts in images. LF-CBM [16] uses CLIP to learn the projection from feature space to concept space. All the methods are implemented using $\mathcal{D}_u$ and evaluated with the test set of $\mathcal{D}_c$. As Table 2 presents, rectified-evi-CEM achieves the best results among all the methods for both concept prediction and diagnosis. Comparing evi-CEM and rectified-evi-CEM, it is evident that concept rectification could identify misaligned concepts with ECBL and correct them using CAVs. In contrast, other VLM-based methods are undermined by concept misalignment issues and cannot obtain accurate and reliable concept explanations.

Uncertainty-aware intervention: In order to validate the effectiveness of uncertainty-aware intervention, we carry out a test-time intervention experiment using the test set of $\mathcal{D}_c$. We intervene for the same evi-CEM in both an uncertainty-aware and random manner. The results are plotted in Fig. 4b. As the number of intervened concepts increases, we observe that the diagnosis AUC of uncertainty-aware intervention exhibits faster improvement than random intervention. This illustrates that uncertainty brings about more significant interventions. Fig. 4a illustrates a qualitative result of uncertainty-aware intervention.

After intervening the concept with the highest uncertainty, the diagnosis label is corrected from non-neoplastic to benign.

4 Conclusion

This work proposes evi-CEM for reliable concept explanations by quantifying concept uncertainty with evidential learning. Furthermore, we propose concept rectification based on evi-CEM to correct concept misalignments. The misalignments originated from VLMs, could be identified and rectified via concept uncertainty. Experimental results show that our approaches improve the reliability of concept explanations, whether trained with concept supervision or VLMs.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chattopadhyay, A., Pilgrim, R., Vidal, R.: Information Maximization Perspective of Orthogonal Matching Pursuit with Applications to Explainable AI. In: Thirty-Seventh Conference on Neural Information Processing Systems (2023)
2. Chattopadhyay, A., Slocum, S., Haeffele, B.D., Vidal, R., Geman, D.: Interpretable by Design: Learning Predictors by Composing Interpretable Queries. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(6), 7430–7443 (2023)
3. Daneshjou, R., Yuksekgonul, M., Cai, Z.R., Novoa, R.A., Zou, J.: SkinCon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. In: Neural Information Processing Systems (2022)
4. Espinosa Zarlenga, M., Barbiero, P., Ciravegna, G., Marra, G., Giannini, F., Diligenti, M., Shams, Z., Precioso, F., Melacci, S., Weller, A., Lió, P., Jamnik, M.: Concept Embedding Models: Beyond the Accuracy-Explainability Trade-Off. In: Advances in Neural Information Processing Systems. vol. 35, pp. 21400–21413 (2022)
5. Gao, S., Zhou, H., Gao, Y., Zhuang, X.: Joint Modeling of Image and Label Statistics for Enhancing Model Generalizability of Medical Image Segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. pp. 360–369 (2022)
6. Gao, S., Zhou, H., Gao, Y., Zhuang, X.: BayeSeg: Bayesian modeling for medical image segmentation with interpretable generalizability. *Medical Image Analysis* **89**, 102889 (2023)
7. Gao, Y., Wang, H., Liu, Z.: An end-to-end atrial fibrillation detection by a novel residual-based temporal attention convolutional neural network with exponential nonlinearity loss. *Knowledge-Based Systems* **212**, 106589 (2021)
8. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1820–1828 (2021)
9. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)

10. Jøsang, A.: Subjective Logic: A Formalism for Reasoning Under Uncertainty. Springer (2016)
11. Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., Sayres, R.: Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In: Proceedings of the 35th International Conference on Machine Learning. pp. 2668–2677 (2018)
12. Kim, C., Gadgil, S.U., DeGrave, A.J., Cai, Z.R., Daneshjou, R., Lee, S.I.: Fostering transparent medical image AI via an image-text foundation model grounded in medical literature (2023)
13. Kim, E., Jung, D., Park, S., Kim, S., Yoon, S.: Probabilistic Concept Bottleneck Models. In: Proceedings of the 40th International Conference on Machine Learning. pp. 16521–16540 (2023)
14. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept Bottleneck Models. In: Proceedings of the 37th International Conference on Machine Learning. pp. 5338–5348 (2020)
15. Liu, T., Tao, D.: Classification with Noisy Labels by Importance Reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(03), 447–461 (2016)
16. Oikarinen, T., Das, S., Nguyen, L.M., Weng, T.W.: Label-Free Concept Bottleneck Models. In: International Conference on Learning Representations (2023)
17. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning Transferable Visual Models From Natural Language Supervision. In: International Conference on Machine Learning. pp. 8748–8763 (2021)
18. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* **1**(5), 206–215 (2019)
19. Sensoy, M., Kaplan, L., Kandemir, M.: Evidential Deep Learning to Quantify Classification Uncertainty. In: Advances in Neural Information Processing Systems. vol. 31 (2018)
20. Tang, Y., Yang, D., Li, W., Roth, H.R., Landman, B., Xu, D., Nath, V., Hatamizadeh, A.: Self-Supervised Pre-Training of Swin Transformers for 3D Medical Image Analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20730–20740 (2022)
21. Torkzadehmahani, R., Nasirigerdeh, R., Rueckert, D., Kaissis, G.: Label Noise-Robust Learning using a Confidence-Based Sieving Strategy. *Transactions on Machine Learning Research* (2022)
22. van der Velden, B.H.M., Kuijf, H.J., Gilhuijs, K.G.A., Viergever, M.A.: Explainable artificial intelligence (XAI) in deep learning-based medical image analysis. *Medical Image Analysis* **79**, 102470 (2022)
23. Yan, A., Wang, Y., Zhong, Y., Dong, C., He, Z., Lu, Y., Wang, W.Y., Shang, J., McAuley, J.: Learning Concise and Descriptive Attributes for Visual Recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3090–3100 (2023)
24. Yan, A., Wang, Y., Zhong, Y., He, Z., Karypis, P., Wang, Z., Dong, C., Gentili, A., Hsu, C.N., Shang, J., McAuley, J.: Robust and Interpretable Medical Image Classifiers via Concept Bottleneck Models (2023)
25. Yang, Y., Panagopoulou, A., Zhou, S., Jin, D., Callison-Burch, C., Yatskar, M.: Language in a Bottle: Language Model Guided Concept Bottlenecks for Interpretable Image Classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19187–19197 (2023)

26. Yuksekgonul, M., Wang, M., Zou, J.: Post-hoc Concept Bottleneck Models. In: International Conference on Learning Representations (2023)
27. Yun, T., Bhalla, U., Pavlick, E., Sun, C.: Do Vision-Language Pretrained Models Learn Composable Primitive Concepts. *Transactions on Machine Learning Research* (2023)
28. Zhang, Y., Tiño, P., Leonardis, A., Tang, K.: A Survey on Neural Network Interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence* **5**(5), 726–742 (2021)
29. Zhao, X., Zhang, D., Liyuan, H., Zhang, T., Xu, B.: Ode-based recurrent model-free reinforcement learning for pomdps. In: *Advances in Neural Information Processing Systems*. vol. 36, pp. 65801–65817 (2023)