
Stochastic Concept Bottleneck Models

Moritz Vandenhirtz*, Sonia Laguna*, Ričards Marcinkevičs, Julia E. Vogt
Department of Computer Science
ETH Zurich
Switzerland

Abstract

Concept Bottleneck Models (CBMs) have emerged as a promising interpretable method whose final prediction is based on intermediate, human-understandable concepts rather than the raw input. Through time-consuming manual interventions, a user can correct wrongly predicted concept values to enhance the model’s downstream performance. We propose *Stochastic Concept Bottleneck Models* (SCBMs), a novel approach that models concept dependencies. In SCBMs, a single-concept intervention affects all correlated concepts, thereby improving intervention effectiveness. Unlike previous approaches that model the concept relations via an autoregressive structure, we introduce an explicit, distributional parameterization that allows SCBMs to retain the CBMs’ efficient training and inference procedure. Additionally, we leverage the parameterization to derive an effective intervention strategy based on the confidence region. We show empirically on synthetic tabular and natural image datasets that our approach improves intervention effectiveness significantly. Notably, we showcase the versatility and usability of SCBMs by examining a setting with CLIP-inferred concepts, alleviating the need for manual concept annotations.

1 Introduction

In today’s world, machine learning plays a crucial role in making important decisions, from healthcare to finance and law. However, as these algorithms become more complex, understanding how they arrive at their decisions becomes increasingly challenging. This lack of interpretability is a significant concern, especially in situations where trustworthiness, transparency, and accountability are paramount (Lipton, 2016; Doshi-Velez & Kim, 2017). Recent studies have focused on Concept Bottleneck Models (CBMs) (Koh et al., 2020; Havasi et al., 2022; Shin et al., 2023), a class of models that predict human-understandable concepts upon which the final target prediction is based. CBMs offer interpretability since a user can inspect the predicted concept values to understand how the model arrives at its final target prediction. Moreover, if they disagree with a concept prediction, they can intervene by adjusting it to the right value, which in turn affects the target prediction.

For example, consider the yellow warbler in Figure 1 (a), where a user might notice that the binary concept ‘yellow primary color’ is mispredicted. Upon this realization, they can intervene on the CBM by setting its value to 1, which increases the probability of the class yellow warbler. This way of interacting allows any untrained user to engage with the model to increase its predictive performance.

However, if the user input is that the primary color is yellow, should not the likelihood of a yellow crown increase too? This adaptation would increase the predicted likelihood of the correct class even more, as yellow warblers are characterized by their fully yellow body. Currently, vanilla CBMs do not exhibit this behavior as they do not use the intervened-on concepts to update their remaining concept predictions. This indicates that they suboptimally adapt to the additional knowledge gained.

*Equal contribution. Correspondence to {moritz.vandenhirtz,slaguna}@inf.ethz.ch

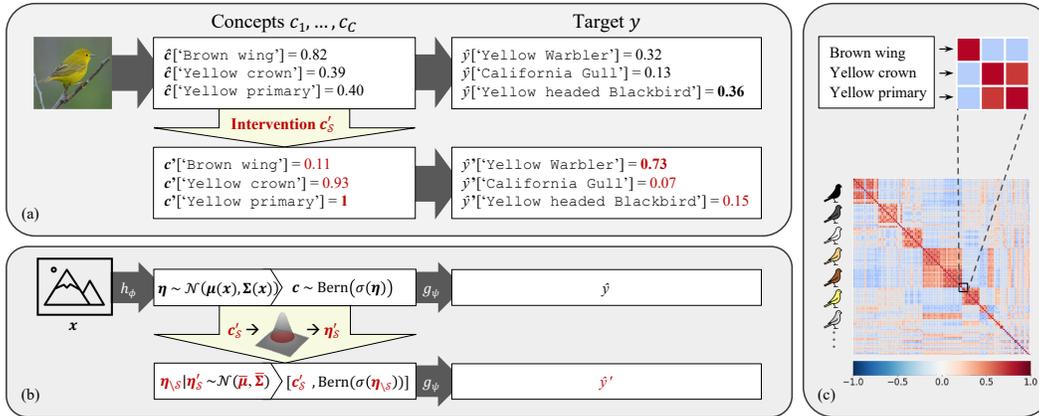


Figure 1: Overview of the proposed method for the CUB dataset. (a) A user intervenes on the concept of ‘primary color: yellow’. Unlike CBMs, our method then uses this information to adjust the predicted probability of correlated concepts, thereby affecting the target prediction. (b) Schematic overview of the intervention procedure. A user’s intervention c'_S is used to infer the logits η'_S of the remaining concepts. (c) Visualization of the learned global dependency structure as a correlation matrix for the 112 concepts of CUB (Wah et al., 2011). Characterization of concepts on the left.

To this end, we propose to extend the concept predictions with the modeling of their dependencies, as depicted in Figure 1.

The proposed approach captures the concept dependencies by modeling the concept logits with a learnable non-diagonal normal distribution, which enables efficient, scalable computing of the effect of interventions on other concepts. By integrating concept correlations, we reduce the time and effort of having to laboriously intervene on many correlated variables and increase the efficacy of interventions on the downstream prediction. Thanks to the explicit distributional assumptions, the model is trained end-to-end, retaining the training and inference speed of classic CBMs as well as the benefits of training the concept and target predictor jointly. Moreover, we show that our method excels when querying user interventions based on predicted concept uncertainty (Shin et al., 2023), further highlighting the practical utility of our approach as such policies spare users from manually sifting through the concepts to identify necessary interventions. Lastly, based on the distributional concept parameterization, we propose a novel approach for computing dependency-aware interventions through the likelihood-based confidence region.

Contributions This work contributes to the line of research on concept bottleneck models in several ways. (i) We propose to capture and model concept dependencies with a multivariate normal distribution. (ii) We derive a novel intervention strategy based on the confidence region of the normal distribution that incorporates concept correlations. Using the learned concept dependencies during the intervention procedure allows for stronger interventional effectiveness. (iii) We provide a thorough empirical assessment of the proposed method on synthetic tabular and natural image data. Additionally, we combine our method with concept discovery where we alleviate the need for annotations by using CLIP-inferred concepts. In particular, we show the proposed method (a) discovers meaningful, interpretable patterns in the form of concept dependencies, (b) allows for fast, scalable inference, and (c) outperforms related work with respect to intervention effectiveness thanks to the proposed concept modeling and intervention strategy.

2 Background & Related Work

Concept bottleneck models (Koh et al., 2020; Lampert et al., 2009; N. Kumar et al., 2009) are typically trained on data points (x, c, y) , comprising the covariates $x \in \mathcal{X}$, target $y \in \mathcal{Y}$, and C annotated binary concepts $c \in \mathcal{C}$. Consider a neural network f_θ parameterized by θ and a slice $\langle g_\psi, h_\phi \rangle$ (Leino et al., 2018) s.t. $\hat{y} := f_\theta(x) = g_\psi(h_\phi(x))$. CBMs enforce a concept bottleneck $\hat{c} := h_\phi(x)$ such that the model’s final output depends on the covariates x solely through the predicted concepts \hat{c} .

While Koh et al. (2020) propose the *soft* CBM, where the concept logits parameterize the bottleneck, Havasi et al. (2022) argue that such a representation leads to leakage, where additional unwanted information in the concept representation is used to predict the target (Margeloiu et al., 2021; Mahinpei et al., 2021). Thus, they parameterize the bottleneck by binarized concept predictions and call it the *hard* CBM. Then, Havasi et al. (2022) equip the hard CBM with an autoregressive structure of the form $c_i | \mathbf{x}, \mathbf{c}_{<i}$, which is supposed to learn the concept dependencies. As such, the implicit autoregressive modeling of concept dependencies by Havasi et al. (2022) is the most related to the current work. Complementary to our work, Heidemann et al. (2023) analyze how a CBM’s performance is affected by concept correlations. Unlike approaches that restrict the bottleneck to prevent leakage, Concept Embedding Models (CEM) (Espinosa Zarlenga et al., 2022) represent each concept with an embedding vector from which the concept probabilities can be inferred. E. Kim et al. (2023) model the embedding with a normal distribution, assuming a diagonal covariance matrix, which prevents them from capturing concept dependencies. Therefore, their intervention performance is not expected to differ from that of CEMs. Recent works explored how a CBM-like structure can be enforced even without a concept-annotated training set. Yuksekogonul et al. (2023) transform a pre-trained model into a CBM via a concept bank from concept activation vectors and multimodal models (B. Kim et al., 2018), while Oikarinen et al. (2023) query GPT-3 (Brown et al., 2020) for the concept set \mathcal{C} and assign the values of the concept activations to each datapoint \mathbf{x} with CLIP (Radford et al., 2021) similarities. Similarly, Panousis et al. (2023) uses CLIP to probabilistically discover a sparse set of concepts for each input, which could be used in our model for a fully probabilistic pipeline. Lastly, Marcinkevičs et al. (2024) instead relax the need for a concept labeled training set to a smaller validation set by fine-tuning a pre-trained model.

Intervenability (Marcinkevičs et al., 2024) is a crucial element of CBMs as it allows the user to correct wrongly predicted concepts \hat{c} to c' , which in turn affects the target prediction of the model \hat{y}' . If multiple concepts are intervened on sequentially, the order of interventions is important. To this end, Sheth et al. (2022) and Shin et al. (2023) explore multiple policies according to which the order of concepts is determined. Chauhan et al. (2023) propose to combine predefined policies with learnable weighting parameters, while Espinosa Zarlenga et al. (2024) learn the policy itself. Concurrently, Singhi et al. (2024) learn a realignment module to align concept predictions. Steinmann et al. (2023) argue that instance-specific interventions are costly and store previous interventions in a memory to automatically reapply them for similar data points. Lastly, Collins et al. (2023) explore the advantages of including uncertainty rather than treating humans as oracles.

Our work models concept dependencies by parameterizing the bottleneck with a distribution. In a similar vein, Variational Autoencoders (Kingma & Welling, 2014) parameterize the bottleneck with a normal distribution to model and generate new data. Stochastic Segmentation Networks (Monteiro et al., 2020) parameterize the logits of a segmentation map with a non-diagonal normal distribution to capture the spatial correlations of pixels and model the aleatoric uncertainty. The modeling of uncertainty with a distribution is also explored by Bayesian Neural Networks (Neal, 1995) that learn a probability distribution over the neurons of a neural network.

3 Methods

We propose Stochastic Concept Bottleneck Models¹ (SCBM), a novel concept-based method that relaxes the implicit CBM assumption of independent concepts. SCBM captures the concept dependencies by learning their multivariate distribution. As a result, interventions become more effective and scalable, as a single intervention can influence multiple correlated concepts. A schematic overview of the proposed method is depicted in Figure 1 (b).

3.1 Model Formulation

To capture the concept dependencies, we model the concept logits $\boldsymbol{\eta}$ with a learned multivariate normal distribution. Modeling logits with a normal distribution has proven to be effective in the context of segmentation (Monteiro et al., 2020). While Monteiro et al. (2020) use it to capture the spatial dependencies of pixels, we, instead, model the relations between concepts, where the properties of the normal distribution will prove useful. A neural network is trained to predict the distribution’s parameters $\boldsymbol{\eta} | \mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma}(\mathbf{x}))$, where $\boldsymbol{\mu}(\mathbf{x}) \in \mathbb{R}^{\mathcal{C}}$, and $\boldsymbol{\Sigma}(\mathbf{x}) \in \mathbb{R}^{\mathcal{C} \times \mathcal{C}}$. Thus, the traditional

¹The code is available here: <https://github.com/mvandenhi/SCBM>.

assumption of independent concepts $c_i \perp\!\!\!\perp c_j \mid \mathbf{x}$, $\forall i \neq j$ is relaxed to $c_i \perp\!\!\!\perp c_j \mid \boldsymbol{\eta}$, $\forall i \neq j$, where the assumed normal distribution induces linear concept dependencies. The inductive bias of linearity is useful in practice as it is more robust to overfitting and computationally more scalable with respect to C compared to its nonlinear alternative (Havasi et al., 2022), as we will show in Section 5.

To learn the distribution, we minimize the negative log-likelihood

$$-\log p(\mathbf{c} \mid \mathbf{x}) = -\log \int p(\mathbf{c} \mid \boldsymbol{\eta}) p_\phi(\boldsymbol{\eta} \mid \mathbf{x}) d\boldsymbol{\eta}, \quad (1)$$

where ϕ are the parameters of a neural network that predicts the distribution $\boldsymbol{\eta} \mid \mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma}(\mathbf{x}))$. This integral is intractable due to the softmax operation applied in $p(\mathbf{c} \mid \boldsymbol{\eta})$. Thus, the integral is approximated by M Monte Carlo samples

$$-\log \int p(\mathbf{c} \mid \boldsymbol{\eta}) p_\phi(\boldsymbol{\eta} \mid \mathbf{x}) d\boldsymbol{\eta} \approx -\log \frac{1}{M} \sum_{m=1}^M p(\mathbf{c} \mid \boldsymbol{\eta}^{(m)}), \quad \boldsymbol{\eta}^{(m)} \mid \mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma}(\mathbf{x})). \quad (2)$$

In order to learn ϕ , we make use of the parameterization as normal distribution and employ the reparameterization trick $\boldsymbol{\eta}^{(m)} \mid \mathbf{x} = \boldsymbol{\mu}(\mathbf{x}) + \mathbf{L}(\mathbf{x})\boldsymbol{\epsilon}^{(m)}$, $\mathbf{L}(\mathbf{x})\mathbf{L}(\mathbf{x})^T = \boldsymbol{\Sigma}(\mathbf{x})$, $\boldsymbol{\epsilon}^{(m)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ such that gradients can be computed with respect to the parameters. Lastly, we incorporate the new relaxed conditional independence assumption

$$\log p(\mathbf{c} \mid \boldsymbol{\eta}) = \log \prod_{i=1}^C p(c_i \mid \eta_i) = \sum_{i=1}^C \log p(c_i \mid \eta_i), \quad (3)$$

where $p(c_i \mid \eta_i)$ describes a Bernoulli distribution parameterized by the sigmoid-transformed logits $\sigma(\eta_i)$. Combining the above considerations results in the following reformulation of the negative log-likelihood:

$$\begin{aligned} -\log p(\mathbf{c} \mid \mathbf{x}) &\approx -\log \frac{1}{M} \sum_{m=1}^M p(\mathbf{c} \mid \boldsymbol{\eta}^{(m)}) \\ &\propto -\log \sum_{m=1}^M \exp \sum_{i=1}^C \log p(c_i \mid \eta_i^{(m)}) \\ &= -\log \sum_{m=1}^M \exp \sum_{i=1}^C \left[-\text{BCE}(c_i, \sigma(\eta_i^{(m)})) \right], \end{aligned} \quad (4)$$

where BCE stands for Binary Cross Entropy, and the logsumexp trick is used for numerical stability.

The distribution-based modeling procedure allows for efficient sampling, thus, enabling SCBM to train concept and target predictors jointly, sequentially, or independently. In contrast, the autoregressive alternative (Havasi et al., 2022) requires independent training due to the computational complexity. We adopt a joint training scheme to obtain the benefits of end-to-end learning where concept and target predictors can adjust to each other. To prevent leakage, we follow Havasi et al. (2022) and train the model with the hard $\{0, 1\}$ concept values as bottleneck rather than the logits used in the original CBM (Koh et al., 2020). To this end, we employ the straight-through Gumbel-Softmax trick (Jang et al., 2017; Maddison et al., 2017) that approximates Bernoulli samples while being differentiable. The target predictor g_ψ is then learned by minimizing the negative log-likelihood

$$\begin{aligned} -\log p(y \mid \mathbf{x}) &= -\log \sum_{\mathbf{c} \in \mathcal{C}} p_\psi(y \mid \mathbf{c}) p(\mathbf{c} \mid \mathbf{x}) \\ &\approx -\log \frac{1}{M} \sum_{m=1}^M p_\psi(y \mid \mathbf{c}^{(m)}), \quad \mathbf{c}^{(m)} \sim p(\mathbf{c} \mid \mathbf{x}). \end{aligned} \quad (5)$$

Lastly, the learned dependencies are regularized by following Occam's razor and to prevent overfitting. We take inspiration from the Graphical Lasso (Friedman et al., 2008) and penalize the off-diagonal elements of the precision matrix $\boldsymbol{\Sigma}^{-1}$.

By combining concept, target, and precision loss with weighting factors λ_1 and λ_2 , we arrive at the final loss function

$$-\log \sum_{m=1}^M \exp \sum_{i=1}^C -\text{BCE}(c_i, \sigma(\eta_i^{(m)})) + \lambda_1 \text{CE} \left(y, \frac{1}{M} \sum_{m=1}^M g_\psi(\mathbf{c}^{(m)}) \right) + \lambda_2 \sum_{i \neq j} \boldsymbol{\Sigma}(\mathbf{x})_{i,j}^{-1}. \quad (6)$$

3.2 Covariance Learning

The introduced amortized covariance matrix $\Sigma(\mathbf{x})$ provides the flexibility to tailor its predicted concept dependencies to each data point, making it adaptable to many data-generating mechanisms. For example, in the commonly used CUB (Wah et al., 2011; Koh et al., 2020), it can learn the class-wise concept structure present in the dataset. The explicit dependency representation inferred by the learned covariance matrix is useful as it provides insights into the learned correlations among the concepts, which is important for understanding and interpreting the model behavior.

However, an amortized covariance matrix comes at the price of not being able to visualize and interpret a unified concept structure on a dataset level. Depending on the need of the application, such a global structure might be preferable. Thus, we propose a variation of SCBM, where the covariance matrix is not *amortized* ($\Sigma(\mathbf{x})$), but learned *globally* (Σ). An example of the global concept structure learned on CUB is shown in Figure 1 (c). This variation has the inductive bias of assuming a constant covariance matrix, whose utility depends on the underlying data-generating mechanism. We recommend using the more flexible, amortized version by default and only utilizing a global covariance if the strong assumption of fixed dependencies is reasonable. We will explore this empirically in more detail in Section 5.

3.3 Interventions

A distinguishing property of CBM-like methods is the user’s capacity to correct wrongly predicted concepts, which in turn affects the target prediction (Marcinkevičs et al., 2024). For a big concept set, this intervention procedure can become quite laborious as a user has to inspect and manually intervene on each concept separately. SCBMs are designed to alleviate this need by utilizing the learned concept dependencies such that a single intervention affects all related concepts as modeled by the multivariate normal distribution.

The parameterization as a multivariate normal distribution allows for a quick, scalable intervention procedure. Given a set $\mathcal{S} \subset \{1, \dots, C\}$ of concept interventions, the effect on the remaining concepts $c_{\setminus \mathcal{S}}$ is computed via their logits $\eta_{\setminus \mathcal{S}}$ by conditioning on the intervention logits $\eta'_{\mathcal{S}}$, utilizing the known properties of the normal distribution

$$\begin{aligned} \eta_{\setminus \mathcal{S}} \mid \mathbf{x}, \eta'_{\mathcal{S}} &\sim \mathcal{N}(\bar{\boldsymbol{\mu}}(\mathbf{x}), \bar{\boldsymbol{\Sigma}}(\mathbf{x})), \\ \bar{\boldsymbol{\mu}} &= \boldsymbol{\mu}_{\setminus \mathcal{S}} + \boldsymbol{\Sigma}_{\setminus \mathcal{S}, \mathcal{S}} \boldsymbol{\Sigma}_{\mathcal{S}, \mathcal{S}}^{-1} (\eta'_{\mathcal{S}} - \boldsymbol{\mu}_{\mathcal{S}}), \\ \bar{\boldsymbol{\Sigma}} &= \boldsymbol{\Sigma}_{\setminus \mathcal{S}, \setminus \mathcal{S}} - \boldsymbol{\Sigma}_{\setminus \mathcal{S}, \mathcal{S}} \boldsymbol{\Sigma}_{\mathcal{S}, \mathcal{S}}^{-1} \boldsymbol{\Sigma}_{\mathcal{S}, \setminus \mathcal{S}}. \end{aligned} \tag{7}$$

In standard CBMs, an intervention affects only the concepts on which the user intervenes. As such, Koh et al. (2020) set η'_i to the 5th percentile of the training distribution if $c_i = 0$ and the 95th percentile if $c_i = 1$. While this strategy is effective for SCBMs too, see Appendix C.5, the modeling of the concept dependencies warrants a more thorough analysis of the *intervention strategy*. We present two desiderata, which our intervention strategy should fulfill.

i) $p(c_i \mid \eta'_i) \geq p(c_i \mid \mu_i)$

The likelihood of the intervened-on concept c_i should always increase after the intervention. If SCBMs used the same strategy as CBMs, it could happen that the initially predicted μ_i was more extreme than the selected training percentile. Then, the interventional shift $\eta'_i - \mu_i$ in Eq. 7 would point in the wrong direction. This would cause $\eta_{\setminus \mathcal{S}}$ to shift incorrectly.

ii) $|\eta'_i - \mu_i|$ should not be “too large”.

We posit that the interventional shift should stay within a reasonable range of values. Otherwise, the effect on $\eta_{\setminus \mathcal{S}}$ would be unreasonably large such that the predicted $\boldsymbol{\mu}_{\setminus \mathcal{S}}$ would be completely disregarded.

To fulfill these desiderata, we take advantage of the explicit distributional representation: the likelihood-based confidence region of μ_i provides a natural way of specifying the region of possible $\eta'_{\mathcal{S}}$ that fulfill our desiderata. Informally, a confidence region captures the region of plausible values for a parameter of a distribution. Note that the confidence region takes concept dependencies into account when describing the area of possible $\eta'_{\mathcal{S}}$. To determine the specific point within this region, we search for the values $\eta'_{\mathcal{S}}$, which maximize the log-likelihood of the known, intervened-on concepts $c_{\mathcal{S}}$, implicitly focusing on concepts that the model predicts poorly:

$$\begin{aligned}
\eta'_S &= \arg \max_{\eta_S} \log p(c_S | \eta_S) \\
\text{s.t. } &-2(\log p(\eta_S | \mu_S, \Sigma_{S,S}) - \log p(\mu_S | \mu_S, \Sigma_{S,S})) \leq \chi_{d,1-\alpha}^2 \quad (8) \\
&\eta'_i - \mu_i \geq 0 \text{ if } c_i = 1, \quad \forall i \in S \\
&\eta'_i - \mu_i \leq 0 \text{ if } c_i = 0, \quad \forall i \in S,
\end{aligned}$$

where $d = |S|$. The first inequality describes the confidence region. It is based on the logarithm of the likelihood ratio, which, after multiplying with -2 , asymptotically follows a χ^2 distribution (Silvey, 1975). The last two inequalities restrict the region to the desired direction. Note that η'_S is computed to determine the conditional effect of the interventions on $\eta_{\setminus S}$ using Equation 7. When predicting \hat{y}' under interventions, the logits $\eta_{\setminus S}$ are then used for sampling the binary concept values $c_{\setminus S}$ while the intervened-on concepts c'_S are directly set to their known, binary value.

4 Experimental Setup

Datasets and Evaluation We perform experiments on a variety of datasets to showcase the validity of our method. Inspired by Marcinkevičs et al. (2024), we introduce a synthetic tabular dataset with a data-generating mechanism that contains fixed concept dependencies we can regulate. In particular, the concept logits η are sampled from a randomly initialized positive definite covariance matrix and generate x . Binary concept values c are inferred from η and generate the target y . We refer to Appendix A.1 for a more detailed description.

As a natural image classification benchmark, we evaluate on the Caltech-UCSD Birds-200-2011 dataset (Wah et al., 2011), comprised of bird photographs from 200 distinct classes. It includes 112 concepts, such as wing color and beak shape, shared across the same class instances as revised in the original CBM work (Koh et al., 2020). Additionally, we explore another natural image classification task on CIFAR-10 (Krizhevsky et al., 2009) with 10 classes. To mitigate the concept annotations requirement, the concepts are synthetically acquired in a similar fashion to the concept discovery literature. We adopt the 143 concept classes generated via GPT-3 (Brown et al., 2020) in prior work (Oikarinen et al., 2023). To obtain the binary concept values, we use the CLIP model (Radford et al., 2021) to compute the similarity between each instance of an image with the text embedding of a specific concept and compare it to the similarity of its negative counterpart, i.e. *not* the concept. Appendix A.2 contains further details about the natural image datasets.

To compare methods, we evaluate the model performance based on the concept and target accuracy. We compute test performance before and after intervening on an increasing number of concepts. The order of concepts in the intervention is determined by an uncertainty-based policy (Shin et al., 2023) that selects the concept whose predicted probability is closest to 0.5. We also show results for a random policy in Appendix C.3. Additionally, we evaluate the calibration of the predicted concept uncertainties that are being used for the uncertainty-based policy, with the Brier score (Brier, 1950) and the Expected Calibration Error (Naeini et al., 2015; A. Kumar et al., 2019).

Baselines We evaluate the performance of our method in comparison with state-of-the-art models. Namely, we focus on the vanilla concept bottleneck model (CBM) by Koh et al. (2020) in its *hard* version (Havasi et al., 2022), trained jointly using the straight-through Gumbel-Softmax trick (Jang et al., 2017; Maddison et al., 2017), as a sensical baseline to our binary modeling of concepts. Additionally, we explore the concept embedding model (CEM) by Espinosa Zarlenga et al. (2022) that learns two concept embeddings, \hat{c}_i^+ and \hat{c}_i^- . These representations are used to predict the final concept probability with a learnable scoring function $\hat{p}_i = s(\hat{c}_i^+, \hat{c}_i^-) = \sigma(\mathbf{W}_s[\hat{c}_i^+, \hat{c}_i^-]^T + \mathbf{b}_s)$ and are then combined into a final concept embedding $\hat{c}_i = (\hat{p}_i \hat{c}_i^+ + (1 - \hat{p}_i) \hat{c}_i^-)$ that is passed to the target predictor. Interventions are modeled by altering the concept probabilities \hat{p}_i . Note that Espinosa Zarlenga et al. (2022) optimize for intervention performance during training, which we omit, to ensure a fair comparison where no method was explicitly trained for intervention performance. Finally, we evaluate the autoregressive CBM structure proposed by Havasi et al. (2022), where concept dependencies are learned with an autoregressive structure. Here, each concept c_i is predicted with a separate MLP that takes as input a latent representation of the input $f_\theta(x)$ and all previous concepts c_1, \dots, c_{i-1} . To obtain a good initialization of the autoregressive structure, it is pretrained

Table 1: Test-set concept and target accuracy (%) prior to interventions. Results are reported as averages and standard deviations of model performance across ten seeds. For each dataset and metric, the best-performing method is **bolded** and the runner-up is underlined.

Dataset	Method	Concept Accuracy	Target Accuracy
Synthetic	Hard CBM	61.42 ± 0.07	58.38 ± 0.39
	CEM	61.42 ± 0.12	58.01 ± 0.49
	Autoregressive CBM	<u>62.17</u> ± 0.11	59.60 ± 0.62
	Global SCBM	61.57 ± 0.05	58.39 ± 0.53
	Amortized SCBM	62.41 ± 0.20	<u>58.96</u> ± 0.38
CUB	Hard CBM	94.97 ± 0.07	67.72 ± 0.57
	CEM	95.12 ± 0.07	<u>69.60</u> ± 0.30
	Autoregressive CBM	95.33 ± 0.07	<u>69.24</u> ± 0.44
	Global SCBM	94.99 ± 0.09	68.19 ± 0.63
	Amortized SCBM	<u>95.22</u> ± 0.09	69.87 ± 0.56
CIFAR-10	Hard CBM	85.51 ± 0.04	69.73 ± 0.29
	CEM	85.12 ± 0.14	72.24 ± 0.33
	Autoregressive CBM	85.31 ± 0.06	68.88 ± 0.47
	Global SCBM	<u>85.86</u> ± 0.04	70.74 ± 0.29
	Amortized SCBM	86.00 ± 0.03	<u>71.66</u> ± 0.25

for 50 epochs. As the Monte Carlo sampling from the autoregressive structure is time-consuming, the target predictor g_ψ is trained independently using the ground-truth concepts as input. At intervention time, a normalized importance sampling algorithm is used to estimate the concept distribution.

Implementation Details The model architectures comprise a backbone for concept prediction followed by a linear layer as head for an interpretable target prediction. More details can be found in Appendix B. To ensure the positive definiteness of the concept covariance matrix Σ , we parameterize it via its Cholesky decomposition $\Sigma = \mathbf{L}\mathbf{L}^\top$. Thus, we directly predict the lower triangular Cholesky matrix \mathbf{L} . We will evaluate two options for SCBMs: using a *global* (Σ) or an *amortized* covariance matrix ($\Sigma(x)$). For the amortized version, we set the weighting terms λ_1 and λ_2 of Equation 6 to 1. For the global version, we initialize it with the estimated empirical covariance matrix and set $\lambda_2 = 0$, as we did not observe big differences when varying λ_2 . In Appendix C.4, we provide an ablation study, demonstrating that SCBMs are not very sensitive to the choice of λ_2 . At intervention time, we solve the optimization problem based on the 99%-confidence region with the SLSQP algorithm (Kraft, 1988). In Appendix C.6, we provide an ablation with different confidence levels.

5 Results

Test performance In Table 1, we report the results of the concept and target accuracy prior to interventions. Overall, SCBM performs on par with the baseline methods, with no clear outperforming or underperforming technique throughout the datasets. In Appendix C.7, we show that other metrics lead to the same interpretation. This shows that the additional overhead of learning the concept dependencies does not negatively affect the predictive performance. We note that the amortized covariance variant consistently surpasses the globally learned matrix due to its ability to adjust the predicted concept dependency structure and uncertainty on an instance level. On the other hand, the global variant offers a unified understanding of the concept correlations, an example of which is presented in Figure 1 (c). Notably, in CIFAR-10, even though the concept performance of CEM is the worst of all methods, it has the best target performance. This might suggest the presence of leakage in CEM’s embeddings, as in CIFAR-10, the concept set alone is not sufficient to predict the target, and learning

Table 2: Relative time it takes for one epoch in the CUB dataset when training on the training set, or evaluating on the test set, respectively.

Method	Training	Inference
Hard CBM	5x	1x
CEM	5x	1x
Autoregressive CBM	5x	15x
Global SCBM	5x	1x
Amortized SCBM	5x	1x

additional information might be useful. In Table 2, we show the time it takes for training and testing of the methods. It is evident that the autoregressive CBM of Havasi et al. (2022) suffers from a slow sampling process due to its autoregressive structure, while SCBMs retain the efficiency of CBMs.

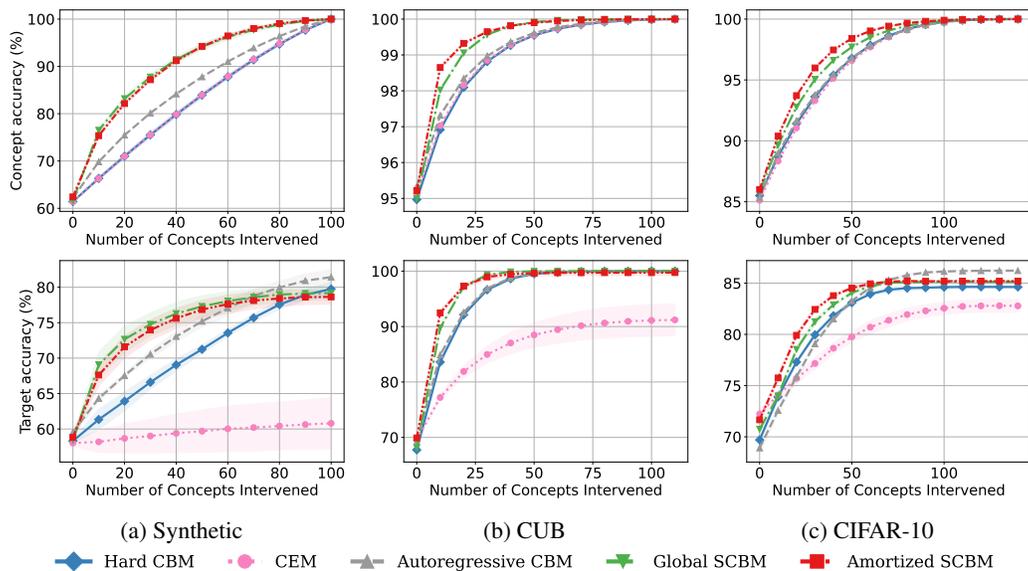


Figure 2: Performance after intervening on concepts in the order of highest predicted uncertainty. Concept and target accuracy (%) are shown in the first and second rows, respectively. Results are reported as averages and standard deviations of model performance across ten seeds.

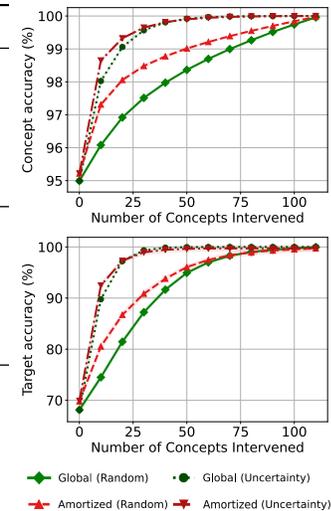
Interventions In this paragraph, we analyze the intervention performance of SCBMs and their baseline models, focusing on their effectiveness in modeling concept dependencies and improving target accuracy. Figure 2 shows the intervention curves across ten seeds, where the performance is measured based on the concept and target accuracy. The order of concepts to intervene on is determined by an uncertainty-based policy that makes use of the predicted probabilities. In Appendix C.3, we present the intervention performance if concepts were selected randomly. The intervention curves in the first row show that SCBMs are superior in modeling the concept dependencies, as evidenced by their significantly steeper intervention curves compared to the baseline methods. Furthermore, the second row of Figure 2 indicates that the strong concept modeling translates to a significant improvement in downstream performance, partly thanks to the intervention strategy introduced in Section 3.3. We note that especially for the most practical scenario of only a small number of interventions, SCBMs outperform their counterparts. Comparing the SCBM variants, the natural image datasets show an overall better intervention performance with the amortized covariance matrix, following the trend of Table 1, as it can capture the instance-wise correlation structure of the data. Only in the synthetic dataset, where the data-generating covariance matrix is fixed, does the global SCBM slightly outperform the amortized one. Thus, we advocate for the usage of the global variant only if the underlying assumption of a fixed covariance is reasonable. Lastly, the success of SCBMs on CIFAR-10, with CLIP-based concepts, shows our proposed method can work without human-annotated concepts. To strengthen this point and also showcase the scalability of our method, in Appendix C.1, we provide results on CIFAR-100 with 892 concepts, where our SCBMs also strongly outperform baselines.

Analyzing the performance of the autoregressive CBM, which also captures concept dependencies, we observe that they expectedly have a better intervention performance than the hard vanilla CBM, which does not take correlations into account. However, it becomes evident that, compared to the concept performance of SCBMs, their autoregressive structure does not capture the dependencies to the full extent. This shows in the target accuracy, where they only match or outperform SCBMs towards the full set of intervened concepts. We attribute the better performance on the full intervention set to the independent training procedure utilized by autoregressive CBMs, which comes at the cost of lower test performance in CIFAR-10. Arguably, in a realistic use-case, such a high number of instance-level interventions is not sensible, and if it were, SCBMs could also be trained independently. Finally, the CEM shows reduced intervention performance as the expressive concept embeddings, which are prone to information leakage, seem to suboptimally adapt to the injected concept information.

Table 3: Test-set calibration (%) of concept predictions. Results are reported as averages and standard deviations of model performance across ten seeds. For each dataset and metric, the best-performing method is **bolded** and the runner-up is underlined. Lower is better.

Dataset	Method	Brier	ECE
Synthetic	Hard CBM	28.79 ± 0.09	22.38 ± 0.15
	CEM	29.32 ± 0.08	23.55 ± 0.09
	Autoregressive CBM	24.84 ± 0.32	13.54 ± 0.49
	Global SCBM	27.73 ± 0.09	20.10 ± 0.14
	Amortized SCBM	<u>25.58</u> ± 0.20	<u>15.57</u> ± 0.55
CUB	Hard CBM	3.93 ± 0.05	2.44 ± 0.06
	CEM	4.04 ± 0.05	3.25 ± 0.07
	Autoregressive CBM	<u>3.75</u> ± 0.05	2.73 ± 0.05
	Global SCBM	3.87 ± 0.06	<u>2.33</u> ± 0.09
	Amortized SCBM	3.64 ± 0.07	1.85 ± 0.08
CIFAR-10	Hard CBM	10.42 ± 0.05	4.93 ± 0.17
	CEM	11.06 ± 0.16	7.11 ± 0.39
	Autoregressive CBM	10.70 ± 0.05	6.07 ± 0.10
	Global SCBM	<u>9.95</u> ± 0.02	<u>2.88</u> ± 0.11
	Amortized SCBM	9.84 ± 0.02	2.22 ± 0.12

Figure 3: Intervention performance of SCBMs measured in concept and target accuracy (%) on CUB for random and uncertainty-based policy.



Modeling the concept distribution A cornerstone of SCBMs is the explicit, distributional parameterization of concepts. This helps in understanding the data correlations and allows for visualization, as the example seen in Figure 1 (c). The explicit probabilistic modeling results in improved concept uncertainty estimates compared to the baseline CBM counterparts, as shown in Table 3, where lower metrics imply better estimates. This proves useful for interventions, where the uncertainty estimates can be leveraged for the choice of concepts to intervene on, improving the target prediction more effectively and reducing the need for manual user inspection. In Figure 3, we compare the performance of randomly intervening versus intervening based on the predicted uncertainty. We observe that there is a big gap between the two policies, indicating the usefulness of the estimated probabilities. Nevertheless, note that intervening at random remains successful and supports the observations made in the previous paragraph, as shown in Appendix C.3.

6 Conclusion

In this paper, we introduced SCBMs, a new concept-based method that models concept dependencies with a multivariate normal distribution. We proposed a novel, effective intervention strategy that takes concept correlations into account and is based on the confidence region inferred from the distributional parameterization. We showed that our modeling approach retains CBMs' training and inference speed, thus, being able to harness the benefits of end-to-end concept and target training. Additionally, the explicit parameterization offers the user a clearer understanding of the learned concept dependencies, providing deeper insights into how predictions and interventions are made. Empirically, we demonstrated that by modeling the concept dependencies, SCBMs offer a substantial improvement in intervention effectiveness, in concept as well as target accuracy, compared to related work. We showed that our method excels when iteratively intervening on the most uncertain concept predictions, sparing users from having to manually search through the concept set to identify necessary interventions. Additionally, our results indicate that learning the concept correlations does not decrease performance prior to interventions, in many cases even improving the performance over the baselines. Finally, the versatility of SCBMs is highlighted through their superior performance on CIFAR-10 and CIFAR-100, where concept values are CLIP-based rather than human-annotated.

Limitations & Future Work This work opens multiple new research avenues. A natural extension is to go beyond binary concepts, such as continuous domains with their corresponding adaptations of modeling the concept distribution. Additionally, addressing the quadratic memory complexity of the covariance matrix is essential for scaling to larger concept sets. Our proposed intervention strategy accounts for model uncertainty, but further research is needed to accommodate user uncertainty, as human interventions are not always the ground truth. This work allows the editing of the learned dependency structure by adjusting the entries of the predicted covariance matrix, which could be explored. Lastly, to model additional information and reduce leakage, Koh et al. (2020); Havasi et al. (2022) propose the adoption of a side channel. The complementary effectiveness of incorporating the side channel in the covariance structure could be explored in the context of SCBMs.

Acknowledgments and Disclosure of Funding

We thank Alexander Marx for the insightful discussions. MV and SL are supported by the Swiss State Secretariat for Education, Research, and Innovation (SERI) under contract number MB22.00047. RM is supported by the SNSF grant #320038189096.

References

- Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., ... others (2024). Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. In *Proceedings of the 29th acm international conference on architectural support for programming languages and operating systems, volume 2* (pp. 929–947). [Referenced on page 14]
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1), 1–3. Retrieved from [https://doi.org/10.1175/1520-0493\(1950\)078<0001:VOFEIT>2.0.CO;2](https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2) [Referenced on page 6]
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... others (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33, 1877–1901. [Referenced on page 3, 6]
- Chauhan, K., Tiwari, R., Freyberg, J., Shenoy, P., & Dvijotham, K. (2023). Interactive concept bottleneck models. In *Proceedings of the aaii conference on artificial intelligence* (Vol. 37, pp. 5948–5955). [Referenced on page 3]
- Collins, K. M., Barker, M., Zarlenga, M. E., Raman, N., Bhatt, U., Jamnik, M., ... Dvijotham, K. (2023). Human uncertainty in concept-based AI systems. In F. Rossi, S. Das, J. Davis, K. Firth-Butterfield, & A. John (Eds.), *Proceedings of the 2023 AAAI/ACM conference on ai, ethics, and society, AIES 2023, montréal, qc, canada, august 8-10, 2023* (pp. 869–889). ACM. [Referenced on page 3]
- Doshi-Velez, F., & Kim, B. (2017, March). *Towards A Rigorous Science of Interpretable Machine Learning* (No. arXiv:1702.08608). arXiv. doi: 10.48550/arXiv.1702.08608 [Referenced on page 1]
- Espinosa Zarlenga, M., Barbiero, P., Ciravegna, G., Marra, G., Giannini, F., Diligenti, M., ... others (2022). Concept embedding models: Beyond the accuracy-explainability trade-off. In *Advances in neural information processing systems* (Vol. 35, pp. 21400–21413). [Referenced on page 3, 6]
- Espinosa Zarlenga, M., Collins, K., Dvijotham, K., Weller, A., Shams, Z., & Jamnik, M. (2024). Learning to receive help: Intervention-aware concept embedding models. *Advances in Neural Information Processing Systems*, 36. [Referenced on page 3]
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3), 432–441. [Referenced on page 4]
- Havasi, M., Parbhoo, S., & Doshi-Velez, F. (2022). Addressing leakage in concept bottleneck models. In A. H. Oh, A. Agarwal, D. Belgrave, & K. Cho (Eds.), *Advances in neural information processing systems*. Retrieved from https://openreview.net/forum?id=tglniD_fn9 [Referenced on page 1, 3, 4, 6, 8, 10]
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 770–778). [Referenced on page 14]
- Heidemann, L., Monnet, M., & Roscher, K. (2023). Concept correlation and its effects on concept-based models. In *Proceedings of the ieee/cvf winter conference on applications of computer vision* (pp. 4780–4788). [Referenced on page 3]
- Jaccard, P. (1901). Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bull Soc Vaudoise Sci Nat*, 37, 547–579. [Referenced on page 17]
- Jang, E., Gu, S., & Poole, B. (2017). Categorical reparameterization with gumbel-softmax. In *5th international conference on learning representations, ICLR 2017, toulon, france, april 24-26, 2017, conference track proceedings*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=rkE3y85ee> [Referenced on page 4, 6]
- Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., & Sayres, R. (2018). Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In J. Dy & A. Krause (Eds.), *Proceedings of the 35th international conference on machine learning* (Vol. 80, pp. 2668–2677). PMLR. Retrieved from <https://proceedings.mlr.press/v80/kim18d.html> [Referenced on page 3]
- Kim, E., Jung, D., Park, S., Kim, S., & Yoon, S. (2023). Probabilistic concept bottleneck models. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, & J. Scarlett (Eds.), *Proceedings of the*

- 40th international conference on machine learning (Vol. 202, pp. 16521–16540). PMLR. Retrieved from <https://proceedings.mlr.press/v202/kim23g.html> [Referenced on page 3]
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Y. Bengio & Y. LeCun (Eds.), *3rd international conference on learning representations, ICLR 2015, san diego, ca, usa, may 7-9, 2015, conference track proceedings*. Retrieved from <http://arxiv.org/abs/1412.6980> [Referenced on page 14]
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational bayes. In Y. Bengio & Y. LeCun (Eds.), *2nd international conference on learning representations, ICLR 2014, banff, ab, canada, april 14-16, 2014, conference track proceedings*. Retrieved from <http://arxiv.org/abs/1312.6114> [Referenced on page 3]
- Koh, P. W., Nguyen, T., Tang, Y. S., Musmann, S., Pierson, E., Kim, B., & Liang, P. (2020). Concept bottleneck models. In H. D. III & A. Singh (Eds.), *Proceedings of the 37th international conference on machine learning* (Vol. 119, pp. 5338–5348). Virtual: PMLR. Retrieved from <https://proceedings.mlr.press/v119/koh20a.html> [Referenced on page 1, 2, 3, 4, 5, 6, 10, 14, 17]
- Kraft, D. (1988). A software package for sequential quadratic programming. *Forschungsbericht-Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*. [Referenced on page 7]
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images. [Referenced on page 6, 14]
- Kumar, A., Liang, P. S., & Ma, T. (2019). Verified uncertainty calibration. *Advances in Neural Information Processing Systems*, 32. [Referenced on page 6]
- Kumar, N., Berg, A. C., Belhumeur, P. N., & Nayar, S. K. (2009). Attribute and simile classifiers for face verification. In *2009 IEEE 12th international conference on computer vision* (pp. 365–372). Kyoto, Japan: IEEE. Retrieved from <https://doi.org/10.1109/ICCV.2009.5459250> [Referenced on page 2]
- Lampert, C. H., Nickisch, H., & Harmeling, S. (2009). Learning to detect unseen object classes by between-class attribute transfer. In *2009 IEEE conference on computer vision and pattern recognition*. Miami, FL, USA: IEEE. Retrieved from <https://doi.org/10.1109/CVPR.2009.5206594> [Referenced on page 2]
- Leino, K., Sen, S., Datta, A., Fredrikson, M., & Li, L. (2018). Influence-directed explanations for deep convolutional networks. In *2018 IEEE international test conference (ITC)*. IEEE. Retrieved from <https://doi.org/10.1109/test.2018.8624792> [Referenced on page 2]
- Lipton, Z. C. (2016, June). The Mythos of Model Interpretability. *Communications of the ACM*, 61(10), 35–43. doi: 10.48550/arxiv.1606.03490 [Referenced on page 1]
- Maddison, C. J., Mnih, A., & Teh, Y. W. (2017). The concrete distribution: A continuous relaxation of discrete random variables. In *5th international conference on learning representations, ICLR 2017, toulon, france, april 24-26, 2017, conference track proceedings*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=S1jE5L5g1> [Referenced on page 4, 6]
- Mahinpei, A., Clark, J., Lage, I., Doshi-Velez, F., & Pan, W. (2021). *Promises and pitfalls of black-box concept learning models*. Retrieved from <https://doi.org/10.48550/arXiv.2106.13314> (arXiv:2106.13314) [Referenced on page 3]
- Marcinkevičs, R., Laguna, S., Vandenhirtz, M., & Vogt, J. E. (2024). Beyond concept bottleneck models: How to make black boxes intervenable? In *Advances in neural information processing systems* (Vol. 37). [Referenced on page 3, 5]
- Marcinkevičs, R., Reis Wolfertstetter, P., Klimiene, U., Chin-Cheong, K., Paschke, A., Zerres, J., ... Vogt, J. E. (2024). Interpretable and intervenable ultrasonography-based machine learning models for pediatric appendicitis. *Medical Image Analysis*, 91, 103042. Retrieved from <https://www.sciencedirect.com/science/article/pii/S136184152300302X> [Referenced on page 6]
- Margeloiu, A., Ashman, M., Bhatt, U., Chen, Y., Jamnik, M., & Weller, A. (2021). *Do concept bottleneck models learn as intended?* Retrieved from <https://doi.org/10.48550/arXiv.2105.04289> (arXiv:2105.04289) [Referenced on page 3]

- Monteiro, M., Le Folgoc, L., Coelho de Castro, D., Pawlowski, N., Marques, B., Kamnitsas, K., ... Glocker, B. (2020). Stochastic segmentation networks: Modelling spatially correlated aleatoric uncertainty. In *Advances in neural information processing systems* (Vol. 33, pp. 12756–12767). [Referenced on page 3]
- Naeni, M. P., Cooper, G., & Hauskrecht, M. (2015). Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 29). [Referenced on page 6]
- Neal, R. M. (1995). *Bayesian learning for neural networks* (Doctoral dissertation, University of Toronto, Canada). Retrieved from https://librarysearch.library.utoronto.ca/permalink/01UTORONTO_INST/14bjeso/alma991106438365706196 [Referenced on page 3]
- Oikarinen, T., Das, S., Nguyen, L. M., & Weng, T.-W. (2023). Label-free concept bottleneck models. In *The 11th international conference on learning representations*. Retrieved from <https://openreview.net/forum?id=F1Cg47MNvBA> [Referenced on page 3, 6, 15]
- Panousis, K. P., Ienco, D., & Marcos, D. (2023). Sparse linear concept discovery models. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 2767–2771). [Referenced on page 3]
- Panousis, K. P., Ienco, D., & Marcos, D. (2024). Coarse-to-fine concept bottleneck models. In *Neurips 2024-38th annual conference on neural information processing systems*. [Referenced on page 17]
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... others (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning* (pp. 8748–8763). [Referenced on page 3, 6]
- Sheth, I., Rahman, A. A., Severyi, L. R., Havaii, M., & Kahou, S. E. (2022). Learning from uncertain concepts via test time interventions. In *Workshop on trustworthy and socially responsible machine learning, neurips 2022*. Retrieved from <https://openreview.net/forum?id=WVe3vok8Cc3> [Referenced on page 3]
- Shin, S., Jo, Y., Ahn, S., & Lee, N. (2023). A closer look at the intervention procedure of concept bottleneck models. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, & J. Scarlett (Eds.), *Proceedings of the 40th international conference on machine learning* (Vol. 202, pp. 31504–31520). PMLR. Retrieved from <https://proceedings.mlr.press/v202/shin23a.html> [Referenced on page 1, 2, 3, 6, 16]
- Silvey, S. (1975). *Statistical inference*. Taylor & Francis. Retrieved from <https://books.google.ch/books?id=qIKLejbVMf4C> [Referenced on page 6]
- Singhi, N., Kim, J. M., Roth, K., & Akata, Z. (2024). Improving intervention efficacy via concept realignment in concept bottleneck models. *arXiv preprint arXiv:2405.01531*. [Referenced on page 3]
- Steinmann, D., Stammer, W., Friedrich, F., & Kersting, K. (2023). *Learning to intervene on concept bottlenecks*. Retrieved from <https://doi.org/10.48550/arXiv.2308.13453> (arXiv:2308.13453) [Referenced on page 3]
- Wah, C., Branson, S., Welinder, P., Perona, P., & Belongie, S. (2011). The caltech-ucsd birds-200-2011 dataset. [Referenced on page 2, 5, 6, 14]
- Yuksekgonul, M., Wang, M., & Zou, J. (2023). Post-hoc concept bottleneck models. In *The 11th international conference on learning representations*. Retrieved from <https://openreview.net/forum?id=nA5AZ8CEyow> [Referenced on page 3]

A Dataset Details

In this section, we provide additional details on the datasets that are being used in the experiments.

A.1 Synthetic Data-Generating Mechanism

Here, we describe the data-generating mechanism of the synthetic dataset in more detail. Let N , p , and C denote the number of independent data points $\{(\mathbf{x}_n, \mathbf{c}_n, y_n)\}_{n=1}^N$, covariates, and concepts, respectively. We set $N = 50,000$, $p = 1,500$, and $C = 100$, with a 60%-20%-20% train-validation-test split. The generative process is as follows:

1. Randomly sample $\mathbf{W} \in \mathbb{R}^{C \times 10}$ s.t. $w_{i,j} \sim \mathcal{N}(0, 1)$ for $1 \leq i \leq C$ and $1 \leq j \leq 10$.
2. Generate a positive definite matrix $\mathbf{\Sigma} \in \mathbb{R}^{C \times C}$ s.t. $\mathbf{\Sigma} = \mathbf{W}\mathbf{W}^T + \mathbf{D}$. Let $\mathbf{D} \in \mathbb{R}^{C \times C}$ s.t. $\mathbf{D} = \delta \mathbf{I}$, where $\delta_i \sim \mathcal{U}_{[0,1]}$ for $1 \leq i \leq C$.
3. Randomly sample logits $\mathbf{H} \in \mathbb{R}^{N \times C}$ s.t. $\eta_n \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ for $1 \leq n \leq N$.
4. Let $c_{n,i} = \mathbb{1}_{\{\eta_{n,i} \geq 0\}}$ for $1 \leq n \leq N$ and $1 \leq i \leq C$.
5. Let $h : \mathbb{R}^C \rightarrow \mathbb{R}^p$ be a randomly initialised multilayer perceptron with ReLU nonlinearities.
6. Let $\mathbf{x}_n = h(\boldsymbol{\eta}_n) + \boldsymbol{\epsilon}_n$ s.t. $\boldsymbol{\epsilon}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ for $1 \leq n \leq N$.
7. Let $g : \mathbb{R}^C \rightarrow \mathbb{R}$ be a randomly initialized linear perceptron.
8. Let $y_n = \mathbb{1}_{\{g(\mathbf{c}_n) \geq y_{med}\}}$ for $1 \leq n \leq N$, where y_{med} denotes the median of $g(\mathbf{c}_n)$.

A.2 Natural Image Datasets

Caltech-UCSD Birds-200-2011 We evaluate on the Caltech-UCSD Birds-200-2011 (CUB)² dataset (Wah et al., 2011). It comprises 11,788 photographs from 200 distinct bird species annotated with 312 concepts, such as belly color and pattern. In this manuscript, we follow the original train-test split and revised the proposed dataset in the initial CBM work (Koh et al., 2020). Here, only the 112 most widespread binary attributes are included in the final dataset, and concepts are shared across samples in identical classes. The images were resized to a resolution of 224×224 pixels. Finally, following the original proposed augmentations, we applied random horizontal flips, modified the brightness and saturation, and applied normalization during training.

CIFAR-10 CIFAR-10³ (Krizhevsky et al., 2009) is a natural image benchmark with 60,000 32×32 colour images and 10 classes. We kept the original train-test split, with 50,000 samples in the train set and a balanced total of 6,000 images per class. We generated 143 concept labels as described in Section 4 using large language and vision models. At training time, as for CUB, we applied augmentations including modifications to brightness and saturation, random horizontal flips and normalisation. Images were rescaled to a size of 224×224 pixels.

B Implementation Details

This section provides further implementation details of SCBM and the evaluated baselines. All methods were implemented using PyTorch (v 2.1.1) (Ansel et al., 2024). All models are trained for 150 epochs for the synthetic and 300 epochs for the natural image datasets with the Adam optimizer (Kingma & Ba, 2015) with a learning rate of 10^{-4} and a batch size of 64. For the independently trained autoregressive model, we split the training epochs into 2/3 for the concept predictor and 1/3 for the target predictor. For the methods requiring sampling, the number of Monte Carlo samples is set to $M = 100$. We provide an ablation for $M = 10$ in Appendix C.2. Note that since the predictor head is very simple, the MC sampling of SCBMs is extremely fast and does not influence computational complexity by more than 0.1%. For the synthetic tabular data, we use a fully connected neural network as backbone, with 3 non-linear layers, batch normalization, and dropout. For the CUB dataset, we use a pretrained ResNet-18 (He et al., 2016), and for the lower-resolution

²https://www.vision.caltech.edu/datasets/cub_200_2011/, no license available

³<https://www.cs.toronto.edu/~kriz/cifar.html>, no license available

CIFAR-10 a simple convolutional neural network with 2 convolutional layers followed by ReLU, Dropout, and a fully connected layer. For fairness in the comparisons, all baselines have the same model architecture choices and all experiments are performed over 10 random seeds.

Resource Usage For the experiments of the main paper, we used a cluster of mostly GeForce RTX 2080s with 2 CPU workers. Over all methods, we estimate an average runtime of 8h per experiment, each running on a single GPU. This amounts to 5 methods \times 3 datasets \times 10 seeds \times 8 hours = 1200 hours. Adding to that, the Ablation Figures required another 40 runs, amounting to a full total of 1520 hours of compute. Please note that we only report the numbers to generate the final results but not the development time, which we roughly estimate to be around 10 times bigger.

C Further Experiments

In this section, we show additional experiments to provide a more in-depth understanding of SCBM’s effectiveness. We ablate multiple hyperparameters to provide an understanding of how they influence the model performance, as well as show the performance of our model in other settings.

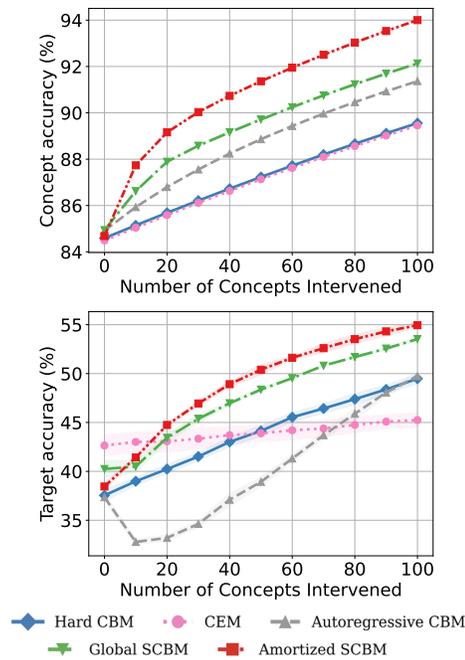


Figure 4: Performance after intervening on concepts in the order of highest predicted uncertainty in CIFAR-100 with 892 concepts. Concept and target accuracy (%) are shown in the first and second rows, respectively. Results are reported as averages and standard deviations of model performance across 3 seeds.

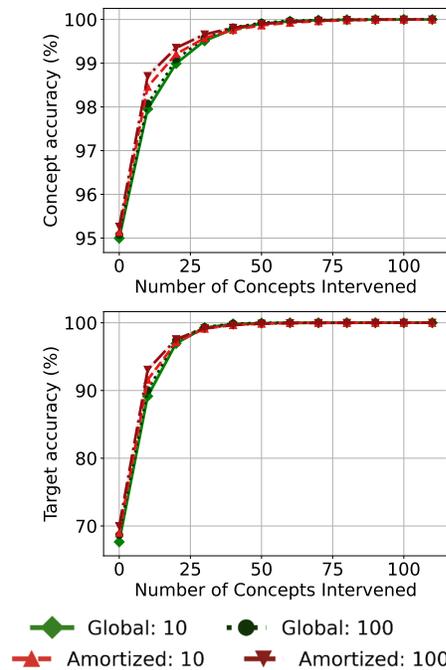


Figure 5: Intervention performance in the order of highest predicted uncertainty in CUB. Concept and target accuracy (%) are shown in the first and second rows, respectively. Results are reported as averages and standard deviations of model performance across 3 seeds.

C.1 Intervention Performance on CIFAR-100

We present the result on the CIFAR-100 dataset with 892 concepts obtained from Oikarinen et al. (2023) in Figure 4 to showcase the scalability of SCBMs. The results underline the efficiency of our method. Notably, the Autoregressive baseline has a negative dip, which is likely due to the independently trained target predictor not being aligned with the concept predictors in this noisy CLIP-annotated scenario. Note that they need to train independently to avoid the sequential MC sampling during training, which would otherwise increase training time significantly. Our jointly

trained SCBMs do not have this issue and surpass the baselines. We use the same configuration as for CIFAR-10, with the exception that we set $M = 10$ to reduce the memory requirement.

C.2 Number of Monte Carlo Samples

To showcase that SCBMs do not rely on a huge number of Monte Carlo samples, we provide an ablation of M in Figure 5. It shows that even for $M = 10$, SCBMs thrive. Note, however, that since M is not a driving factor of SCBMs computational cost, one can leave it at a high number.

C.3 Random Intervention Policy

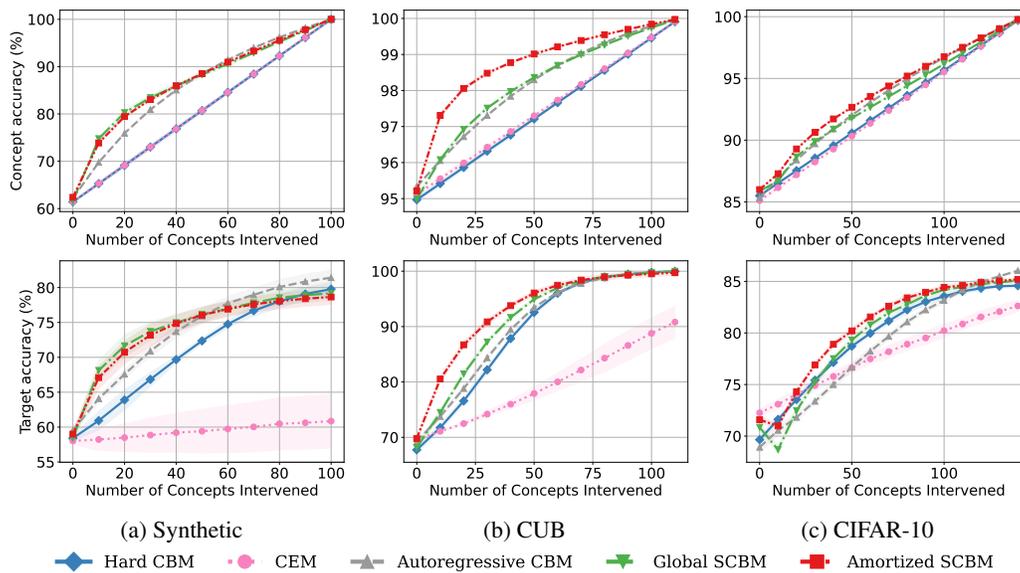


Figure 6: Performance after intervening on concepts in random order. Concept and target accuracy (%) are shown in the first and second rows, respectively. Results are reported as averages and standard deviations of model performance across ten seeds.

In Figure 6, we present the intervention performance of SCBM and baseline methods. Compared to the uncertainty-based intervention policy of Figure 2, the intervention curves of all methods are less steep, confirming the usefulness of Shin et al. (2023)’s proposed policy. Following the previous statements, SCBMs still outperform baseline methods with the amortized beating the global variant for real-world datasets. We observe that in CIFAR-10 for the first interventions, an improvement in concept accuracy is not directly reflected in improved target prediction for SCBMs, which is likely due to the low signal-to-noise ratio of the CLIP-inferred concepts.

C.4 Regularization Strength

In Figure 7, we analyze the impact of the strength of λ_2 from Equation 6. Due to environmental considerations, we conducted experiments using only 5 seeds and limited the number of interventions to 20. Our findings indicate that SCBMs are not sensitive to the choice of λ_2 , except that the unregularized amortized variant exhibits slight patterns of overfitting.

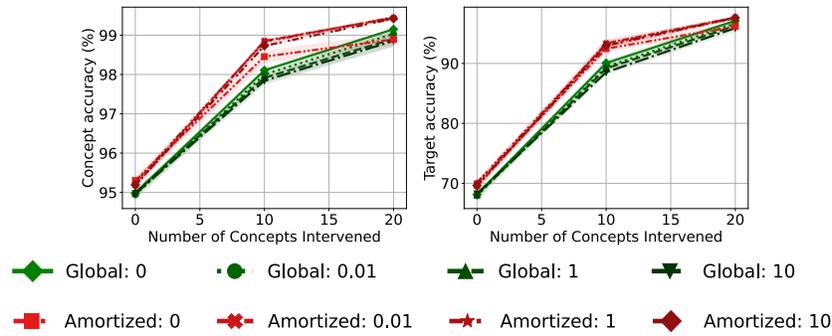


Figure 7: Performance on CUB after intervening on concepts in the order of highest predicted uncertainty with differing regularization strengths. Concept and target accuracy (%) are shown in the first and second columns, respectively. Results are reported as averages and standard deviations of model performance across five seeds. For each SCBM variant, we choose a darker color, the higher the regularization strength of λ_2 .

C.5 Intervention Strategy

In Figure 8, we analyze the effect of the intervention strategy. Our findings indicate that while SCBMs are still effective with the proposed strategy from Koh et al. (2020), that sets the logits to the 5th (if $c_i = 0$) or 95th (if $c_i = 1$) percentile of the training distribution, our proposed strategy based on the confidence region results in stronger intervenability.

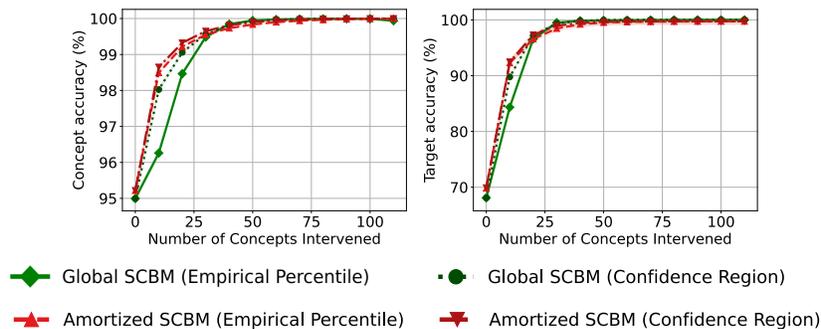


Figure 8: Performance on CUB after intervening on concepts in the order of highest predicted uncertainty, comparing the proposed intervention strategy to Koh et al. (2020)’s intervention of setting the logits to the 5th or 95th empirical percentile of the training distribution. Concept and target accuracy (%) are shown in the first and second columns, respectively. Results are reported as averages and standard deviations of model performance across five seeds.

C.6 Confidence Region Level

In Figure 9, we analyze the effect of the level $1 - \alpha$ of the likelihood-based confidence region. Our findings indicate that the SCBMs are not sensitive to the choice of $1 - \alpha$, with higher levels being slightly better in performance.

C.7 Jaccard Index

Panousis et al. (2024) propose to interpret the interpretation capacity of concepts with the Jaccard Index (Jaccard, 1901). As such, in Table 4, we extend Table 1 with this metric. It is evident that the interpretation does not change, indicating that the performance is robust to the choice of evaluation metric.

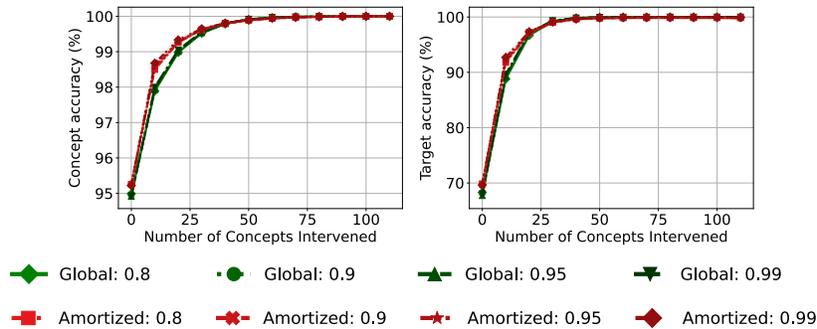


Figure 9: Performance on CUB after intervening on concepts in the order of highest predicted uncertainty with differing levels $1 - \alpha$ of the confidence region. Concept and target accuracy (%) are shown in the first and second columns, respectively. Results are reported as averages and standard deviations of model performance across three seeds.

Table 4: Test-set performance before interventions. Results are averaged across ten seeds.

Dataset	Method	Concept Accuracy	Concept Jaccard	Target Accuracy
Synthetic	Hard CBM	61.42 ± 0.07	43.80 ± 1.32	58.38 ± 0.39
	CEM	61.42 ± 0.12	44.84 ± 1.36	58.01 ± 0.49
	Autoregressive CBM	<u>62.17</u> ± 0.11	<u>45.30</u> ± 1.29	59.60 ± 0.62
	Global SCBM	61.57 ± 0.05	44.53 ± 1.02	58.39 ± 0.53
	Amortized SCBM	62.41 ± 0.20	45.85 ± 1.45	<u>58.96</u> ± 0.38
CUB	Hard CBM	94.97 ± 0.07	77.22 ± 0.33	67.72 ± 0.57
	CEM	95.12 ± 0.07	78.20 ± 0.28	<u>69.60</u> ± 0.30
	Autoregressive CBM	95.33 ± 0.07	79.21 ± 0.21	69.24 ± 0.44
	Global SCBM	94.99 ± 0.09	76.83 ± 0.47	68.19 ± 0.63
	Amortized SCBM	<u>95.22</u> ± 0.09	<u>78.29</u> ± 0.28	69.87 ± 0.56
CIFAR-10	Hard CBM	85.51 ± 0.04	81.54 ± 0.08	69.73 ± 0.29
	CEM	85.12 ± 0.14	81.06 ± 0.21	72.24 ± 0.33
	Autoregressive CBM	85.31 ± 0.06	81.31 ± 0.10	68.88 ± 0.47
	Global SCBM	<u>85.86</u> ± 0.04	<u>81.81</u> ± 0.19	70.74 ± 0.29
	Amortized SCBM	86.00 ± 0.03	81.97 ± 0.20	<u>71.66</u> ± 0.25

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: Claims are supported by evidence in the Results section and Appendix.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Yes, we have a Limitations & Future Work paragraph at the end of the conclusion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide derivations of the method's theoretical foundations (detailed up to an acceptable degree of expected math knowledge) in the Method section.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We disclose hyperparameters in the main text and Appendix. We also offer the code for reproducibility in case any information is missing.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have released an anonymized version of the repository.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.

- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See Question 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide error bars in all experiments as we believe this to be of utmost importance to reproducible research. For the Appendix, we have reduced the number of seeds and/or experiment size to save computational resources for the environment's sake.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: See Appendix

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The code of ethics was followed.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Given the more foundational work of this paper, there is not a direct negative influence that the authors can think of that might arise from this work specifically.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: To the best of our knowledge, our work does not have high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Licenses for all used datasets were clearly stated.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: In the Appendix, the data generating mechanism is clearly stated for the introduced synthetic dataset. Additionally, the new method is described in detail.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.

- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.