

MITIGATING SHORTCUT LEARNING WITH DIFFUSION COUNTERFACTUALS AND DIVERSE ENSEMBLES

Anonymous authors

Paper under double-blind review

ABSTRACT

Spurious correlations in the data, where multiple cues are predictive of the target labels, often lead to a phenomenon known as shortcut learning, where a model relies on erroneous, easy-to-learn cues while ignoring reliable ones. In this work, we propose *DiffDiv* an ensemble diversification framework exploiting Diffusion Probabilistic Models (DPMs) to mitigate this form of bias. We show that at particular training intervals, DPMs can generate images with novel feature combinations, even when trained on samples displaying correlated input features. We leverage this crucial property to generate synthetic counterfactuals to increase model diversity via ensemble disagreement. We show that DPM-guided diversification is sufficient to remove dependence on shortcut cues, without a need for additional supervised signals. We further empirically quantify its efficacy on several diversification objectives, and finally show improved generalization and diversification on par with prior work that relies on auxiliary data collection.

1 INTRODUCTION

Deep Neural Networks (DNNs) have achieved unparalleled success in countless tasks across diverse domains. However, they are not devoid of pitfalls. One such downside manifests in the form of shortcut learning, a phenomenon whereby models latch onto simple, non-essential cues that are spuriously correlated with target labels in the training data (Geirhos et al., 2020; Shah et al., 2020; Scimeca et al., 2022). Often engendered by the under-specification present in data, this simplicity bias presents easy to learn *shortcuts* that allow accurate prediction at training time, irrespective of a model’s alignment to the downstream task. For instance, previous work has found models to incorrectly rely on background signals for object recognition (Xiao et al., 2021; Beery et al., 2018), or to rely on non-clinically relevant metal tokens to predict patient conditions from X-Ray images (Zech et al., 2018). Shortcut learning has often been found to lead to significant drops in generalization performance (Agrawal et al., 2018; Torralba & Efros, 2011; Minderer et al., 2020; Li & Vasconcelos, 2019; Zech et al., 2018). Leveraging shortcut cues can also be harmful when deploying models in sensitive settings, for example, in the reinforcement of harmful biases over the use of sensitive attributes such as gender or skin color (Wang et al., 2019; Xu et al., 2020; Scimeca et al., 2022).

Addressing simplicity biases in machine learning has been a focus of extensive research, often aimed to encourage models to use a more diversified set of predictive cues. A variety of these methods have been input-centric, designed to drive models to focus on different areas of the input space (Teney et al., 2022b; Nicolicioiu et al., 2023), while others have focused on diversification strategies that rely on auxiliary data for *prediction disagreement* (Pagliardini et al., 2022; Lee et al., 2022). The latter approaches, in particular, have been instrumental in developing functionally diverse models that exhibit robustness to shortcut biases. However, they are limited by their required access to auxiliary data that is often challenging to obtain.

The primary objective of this work is to mitigate shortcut learning tendencies, particularly when they result in strong, unwarranted biases, access to *ood* data is expensive, and different features may rely on similar areas of the input space. To achieve this objective, we propose *DiffDiv*, an ensemble framework relying on unlabelled *ood* data for shortcut mitigation by ensemble diversification. To overcome the challenges of the past, we aim to synthetically generate the data for model diversification, and thus avoid the impracticality of *ood* data collection. We posit that the synthetic data should: first, lie in the manifold of the data of interest; and second, be at least partially free of the same shortcuts as the

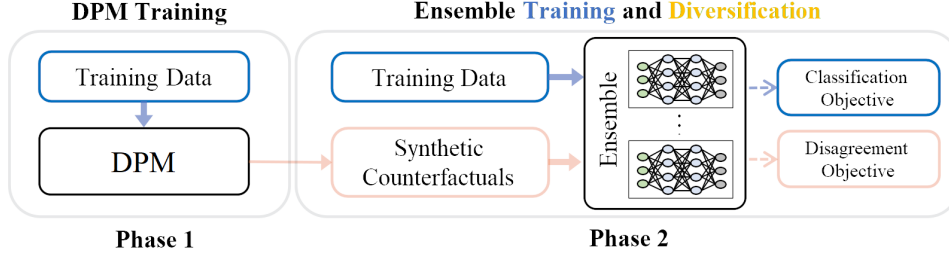


Figure 1: DiffDiv: We sample from a DPM to generate synthetic counterfactuals showcasing emergent novel feature combinations. These samples are then utilized to build a diverse model ensemble via different ensemble disagreement objectives.

original training data. We leverage Diffusion Probabilistic Models (DPMs) to generate synthetic data for ensemble disagreement, leading to strong model diversification.

We observe, and empirically show, that even in the presence of correlated features in the data, DPMs can be used to generate synthetic counterfactuals that break the shortcut signals present at training time. We further show that shortcut mitigation can be achieved via ensemble disagreement on these DPM-generated samples. Remarkably, our experiments confirm that the extent and quality of our diffusion-guided ensemble diversification is on par with existing methods that rely on additional data.

Our contributions are the following:

1. We show that DPMs can generate feature compositions beyond data exhibiting correlated input features.
2. We demonstrate that ensemble disagreement is sufficient for shortcut cue mitigation.
3. We propose *DiffDiv*, a framework to achieve bias mitigation through diversification based on diffusion counterfactuals.
4. We show appropriately trained DPM counterfactuals can lead to state-of-the-art diversification and shortcut bias mitigation, without the need for additional data collection.

Moreover, we apply for the first time the Wisconsin Card Sorting Test for Machine Learners (WCST-ML) (Scimeca et al., 2022) to the CelebA face dataset, exposing biased inference for *pale skin* features. We provide an overview of related work in Suppl. section S1

2 METHODS

2.1 DIFFDIV OVERVIEW

We apply DiffDiv in two stages. The first stage corresponds to training a DPM on the dataset of interest. The second stage corresponds to training an ensemble on the same data and jointly optimize two objectives, a standard classification objective, and a diversification objective. The standard classification objective is performed on real data, while the diversification objective is performed on synthetic counterfactuals generated by the pre-trained DPM. We only consider a fixed small set of 3k counterfactuals in all experiments, and show it to be sufficient for bias mitigation, allowing our method to scale well in complex settings. Please see Suppl. §subsection S3.5 for ablations on counterfactual set size, and their (negligible) impact on *DiffDiv*’s performance.

2.2 DPMs AND EFFICIENT SAMPLING

We utilize Diffusion Probabilistic Models (DPMs) to generate synthetic data for our experiments. DPMs operate by iteratively adding or removing noise from an initial data point x through a stochastic process governed by a predefined noise schedule. Let $\mathbf{z} = \{\mathbf{z}_t \mid t \in [0, 1]\}$ be a latent variable conditioned on t , and characterized by a noise-to-signal ratio $\lambda_t = \log [\alpha_t^2 / \sigma_t^2]$ decreasing monotonically with t . In the forward process, noise is added to x to transform it into z_t (Salimans & Ho, 2022):

$$q(\mathbf{z}_t | \mathbf{x}) = \mathcal{N}(\mathbf{z}_t; \alpha_t \mathbf{x}, \sigma_t^2 \mathbf{I}), \quad q(\mathbf{z}_t | \mathbf{z}_s) = \mathcal{N}(\mathbf{z}_t; (\alpha_t / \alpha_s) \mathbf{z}_s, \sigma_{t|s}^2 \mathbf{I}) \quad (1)$$

where and $0 \leq s < t \leq 1$, and $\sigma_{t|s}^2 = (1 - e^{\lambda_t - \lambda_s}) \sigma_t^2$. We let α_t follow a cosine schedule, thus $\alpha_t = \cos(0.5\pi t)$.

The reverse process then aims to reconstruct $x = z_0$ by iteratively denoising z_t into z_s , starting from $z_1 \sim \mathcal{N}(0, I)$. To facilitate efficient sampling, we employ Denoising Diffusion Implicit Models (DDIM) (Song et al., 2020), a first-order ODE solver for DPMs (Salimans & Ho, 2022; Lu et al., 2022), utilizing a predictor-corrector scheme to minimize the number of sampling steps:

$$\begin{aligned} \mathbf{z}_s &= \alpha_s \hat{\mathbf{x}}_\theta(\mathbf{z}_t) + \sigma_s \frac{\mathbf{z}_t - \alpha_t \hat{\mathbf{x}}_\theta(\mathbf{z}_t)}{\sigma_t} \\ &= e^{(\lambda_t - \lambda_s)/2} (\alpha_s / \alpha_t) \mathbf{z}_t + \left(1 - e^{(\lambda_t - \lambda_s)/2}\right) \alpha_s \hat{\mathbf{x}}_\theta(\mathbf{z}_t) \end{aligned} \quad (2)$$

The term $\hat{\mathbf{x}}_\theta(\mathbf{z}_t)$ is the neural network’s output, predicting the denoised data from the noisy observation at timestep t . And train the denoising model by:

$$L_\theta = \left(1 + \frac{\alpha_t^2}{\sigma^2}\right) \|\mathbf{x} - \hat{\mathbf{x}}_t\|_2^2 \quad (3)$$

We consider DPMs trained at different fidelity levels, and use the ‘number of diffusion training epochs’ as a proxy for the DPM fidelity on the modeled distribution.

2.3 ENSEMBLE TRAINING AND DIVERSIFICATION

In our setup, we wish to train a set of models within an ensemble while encouraging model diversity. Let f_i denote the i^{th} model predictions within an ensemble consisting of N_m models. Each model is trained on a joint objective comprising the conventional cross-entropy loss with the target labels, complemented by a diversification term computed on synthetic counterfactuals, represented as L_{div} .

$$\mathcal{L} = \mathcal{L}_{\text{xent}} + \gamma \mathcal{L}_{\text{div}}^{\text{obj}} \quad (4)$$

where $\mathcal{L}_{\text{xent}}$ is the cross-entropy loss, $\mathcal{L}_{\text{div}}^{\text{obj}}$ is the diversification term for a particular objective, and γ is a hyper-parameter used to modulate the importance of diversification within the optimization objective. To impart diversity to the ensemble, we investigated five diversification objectives denoted by $\text{obj} \in \{\text{div}, \text{cross}, L_1, L_2, \text{kl}\}$.

The L_1 and L_2 baseline objectives were designed to induce diversity by maximizing the distance between each model output and the moving average of the ensemble prediction, thus:

$$L_{\text{reg}}^{L_1} = -\frac{1}{N_m} \sum_{i=1}^{N_m} \left\| f_i - \frac{1}{N_m} \sum_{j=1}^{N_m} f_j \right\|_1, \quad L_{\text{reg}}^{L_2} = -\frac{1}{N_m} \sum_{i=1}^{N_m} \left\| f_i - \frac{1}{N_m} \sum_{j=1}^{N_m} f_j \right\|_2^2$$

The *cross* objective diversifies the predictions by minimizing the negative mutual cross-entropy of any two models:

$$\mathcal{L}_{\text{reg}}^{\text{cross}} = -\frac{1}{N_m(N_m - 1)} \sum_{i \neq j} \frac{\text{CE}(f_i, \text{argmax}(f_j)) + \text{CE}(f_j, \text{argmax}(f_i))}{2}$$

Finally, the *kl* objective aims at maximizing the *kl* divergence between the output distributions of any two models, while the *div* diversification objective is adapted from (Lee et al., 2022) and encourages diversity by minimizing the mutual information of any two models’ predictions:

$$\mathcal{L}_{\text{reg}}^{\text{kl}} = -\frac{1}{N_m(N_m - 1)} \sum_{i \neq j} D_{KL}(f_i || f_j), \quad \mathcal{L}_{\text{reg}}^{\text{div}} = \frac{1}{N_m(N_m - 1)} \sum_{i \neq j} D_{KL}(p(f_i, f_j) || p(f_i))$$

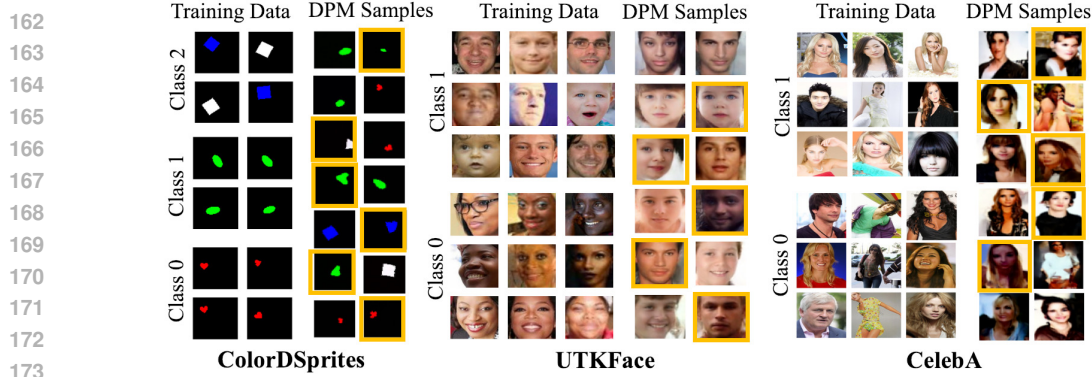


Figure 2: DPM training and counterfactual generation. While training on images showcasing a correlated set of features (left columns), DPM samples at appropriate fidelity levels can generate novel objects beyond the observed feature combinations (marked right-hand side images).

3 RESULTS

3.1 IMPLEMENTATION

To investigate the objectives of this study, we apply the WCST-ML framework (Scimeca et al., 2022) on three datasets, namely ColorDSprites, UTKFace and CelebA. Each dataset is thus divided into a training set, comprising fully correlated feature-labels groups (Figure 1), and as many test sets as there are label classes. For ColorDSprites, we consider the features of *color, orientation, scale, shape*, in UTKFace we consider the features *ethnicity, gender, age*, while in CelebA we consider the features *lightskin, oval face*. We provide more implementation details in Suppl section S2.

For the first stage we train three DPMs on the datasets with fully correlated features, each including respectively 34998, 1634 and 2914 feature-correlated samples. We generate $\approx 100k$ samples from DPMs trained at varying number of epochs between 1 to 1.2K, to be used for ensemble diversification and analysis. We perform no post-processing or pruning of the generated samples, and instead wish to observe the innate ability of DPMs to generate samples beyond the training distribution.

In the second stage, we separately train ensembles for each combination of dataset and diversification objective, and for each, we perform ablation studies by applying the diversification objective to the data generated by the DPMs at different fidelity levels. For comparison, we also consider ensemble diversification with real *ood* data, as well as a standard ensemble baseline with no diversification.

3.2 DIFFUSION COUNTERFACTUAL SAMPLING

DPMs Exhibit Generalization Capabilities Under Correlation: We show in Figure 2 the training samples for both datasets (left halves), as well as samples from the DPMs (right halves) trained for 25, 800, and 1000 epochs, respectively for each dataset. We observe how sampling from the trained DPMs generates previously unseen feature combinations, despite the correlated coupling of features during training (e.g. $\langle \text{green/white, heart} \rangle$, $\langle \text{child, female} \rangle$ or $\langle \neg \text{light skin, oval face} \rangle$), suggesting a potential for DPMs to transcend surface-level statistics of the data, extending previous findings to the special case of fully correlated input features in the target distribution.

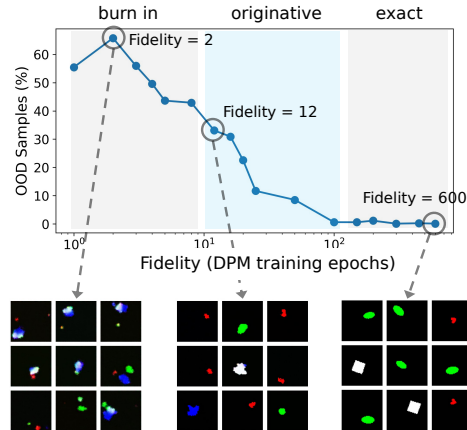


Figure 3: *ood* sample frequency for DPMs trained at different fidelities on ColodDSprites.

Early Stopping to Capture Diffusion *ood* Sampling Capabilities: To understand under which conditions sampling from DPMs can lead to the generation of samples displaying unseen feature combinations, we examine the fraction of *ood* samples generated by DPMs at different fidelities (§3.1). We use the ColorDSprites given the disentangled nature of its features, and the presence of all feature combinations in the data. To measure the fraction of *ood* samples, we train a near-perfect oracle on the full ColorDSprites dataset, and classify samples based on the *diagonal* and *off-diagonal* sets in WCST-ML (Scimeca et al., 2022).

Figure 3 shows the fraction of *ood* samples generated at varying levels of training fidelities. We identify at least three qualitative different intervals. An initial *burn-in* interval, characterized by a high frequency of *ood* generated samples, but which fails to capture the manifold of the data; an *originative* interval, where we observe a reduced number of *ood* samples, in favor for a generative distribution more aligned with the data to represent; and an *exact* interval, where the DPM’s ability to almost perfectly represent the data comes at the cost of novel emergent feature mixtures. In the context of ensemble diversification and bias mitigation, the *originative* interval is posed to provide the necessary information to break the simplicity shortcuts, while providing effective disagreement signals with respect to the visual cues available. We refer the reader to Suppl. §S3 for further insights into DPM-early stopping with ensemble diversification criteria. Interestingly, we will later show that low-fidelity samples can induce appropriate diversification within *DiffDiv*.

3.3 DIFFUSION-GUIDED ENSEMBLE DIVERSITY

We test whether diffusion counterfactuals can lead to ensemble diversification. We train ensembles comprising 100 ResNet-18 models on ColorDSprites, UTKFace, and CelebA. We train the ensembles with a cross-entropy loss on the correlated diagonal training data, as well as an additional diversification loss, computed on a separate diversification set. In our experiments, we will refer to ‘*ood*’ to the case where the data for disagreement is comprised of left-out, feature-uncorrelated, samples from the original dataset; to ‘diffusion’ when the samples are generated by a DPM; and to ‘baseline’ when no diversification objective is included. We tune γ by grid-search to provide comparable experiments under different objectives (see Suppl. §S3.1)

DPM Fidelity Significantly Impacts Diversification: We consider the case where the diversification objective is computed on samples generated by a DPM at different fidelity levels. We show in Figure 4 the diversity achieved with the *div* diversification objective when using diversification generated data at different diffusion fidelities. We observe a significant variation in the diversity of the ensemble predictions, from the low-diversity baseline (BS), trained without a diversification objective, to the high diversity achieved with *ood* samples. For both datasets, extremely low fidelities provide little use for diversification, leading to low-diversity predictions by the ensembles. The DPM trained on ColorDSprites, however, quickly fits the synthetic data distribution and provides counterfactual samples achieving similar diversification performance to the original off-diagonal samples at early fidelity levels. Ultimately, excessive diffusion training leads to limited *ood* sample generation and lower diversification performance. We report analogous results for the UTKFace and CelebA dataset in suppl Figure S3. Importantly, we find that appropriate DPM training and early stopping procedures are key to achieve diversity. We expand on early stopping signals for DPM training in Suppl. S3.3.




Figure 4: Output diversity using samples from diffusion models at different fidelities, as compared to using real *ood* samples (*ood*), and a baseline without diversification (BS) in ColorDSprites. We report all results in Suppl. Figure S3.

Diffusion-guided Diversity Leads to Comparably Diverse Ensembles: Given the difficulty and cost of collecting auxiliary *ood* data for diversification, we wish to compare the level of diversification achieved with DPM counterfactual, as opposed to using real *ood* data. Based on our previous results in Figure 4 (also Suppl. Figure S3), we choose samples drawn from DPMs trained for 100, 800 and 1200 epochs for ColorDSprites, UTKFace and CelebA respectively, providing samples especially useful for diversification. In Figure 5 we report the objective-wise normalized diversity achieved in each scenario by the ensemble. We find that diffusion counterfactuals can lead to comparable

Table 1: Comparison between diversification with *ood* samples and DiffDiv with model disagreement on five objectives. The feature columns report the fraction of models (not accuracy) attending to the respective features. The final column reports the average validation accuracy on the original training data. The shortcut feature for each dataset is highlighted in bold.

Dataset →	Obj. ↓ →	ColorDSprites					UTKFace				CelebA		
		Color (↓)	Orient.	Scale	Shape	Acc. (↑)	Age	Ethnicity (↓)	Gender	Acc. (↑)	Oval Face	Pale Skin (↓)	Acc. (↑)
Baseline	-	1.00	0.00	0.00	0.00	1.000±0.00	0.00	1.00	0.00	0.920±0.02	0.00	1.00	0.857±0.01
OOD	Cross	0.99	0.00	0.01	0.00	0.865±0.17	0.01	0.71	0.28	0.865±0.04	0.01	0.99	0.751±0.07
	Div	0.86	0.01	0.11	0.02	0.818±0.22	0.00	0.94	0.06	0.859±0.03	0.00	1.00	0.843±0.03
	KL	0.91	0.02	0.07	0.00	0.822±0.21	0.05	0.76	0.19	0.818±0.06	0.00	1.00	0.812±0.04
	L1	0.91	0.00	0.08	0.01	0.813±0.20	0.02	0.62	0.36	0.847±0.07	0.14	0.86	0.724±0.11
	L2	0.84	0.02	0.13	0.01	0.729±0.23	0.03	0.63	0.34	0.798±0.11	0.12	0.88	0.651±0.10
DiffDiv (ours)	Cross	0.96	0.00	0.04	0.00	0.856±0.16	0.00	0.94	0.06	0.836±0.05	0.01	0.99	0.745±0.10
	Div	0.94	0.00	0.06	0.00	0.916±0.13	0.00	0.98	0.02	0.826±0.05	0.00	1.00	0.857±0.01
	KL	0.89	0.01	0.10	0.00	0.786±0.20	0.00	0.94	0.06	0.837±0.06	0.04	0.96	0.672±0.07
	L1	0.89	0.00	0.09	0.02	0.784±0.20	0.00	0.77	0.23	0.816±0.11	0.08	0.92	0.659±0.12
	L2	0.93	0.02	0.03	0.02	0.762±0.22	0.01	0.83	0.16	0.757±0.12	0.09	0.91	0.650±0.11




Figure 5: Average normalized ensemble diversity comparisons across metrics (higher is better).

diversification performance with respect to real *ood* samples. In the figure, the diffusion-led diversity is almost always within 5% from the metrics achieved when using pure *ood* samples, both typically over 50% higher than the baseline.

Ensemble Diversification Breaks Simplicity Biases: By WCST-ML, we can test each model’s bias to a cue by testing the model’s output on purposefully designed test sets (Scimeca et al., 2022). We report in Table 1 the fraction of models attending to specific cues, when trained on real *ood* data and diffusion-generated samples (*DiffDiv*). A model was deemed to be attending to a cue if its validation accuracy on the cue-specific *ood* WCST-ML dataset was highest relative to all other features. Firstly, training with no diversification on fully correlated data yields especially biased ensembles, solely attending to the cues of *colore*, *ethnicity*, and *pale skin*. Notably, upon introducing the diversification objectives, we observed a perceptible shift in the models’ behavior, averting their attention from primary, easy-to-learn, cues, to other latent cues present within the data. Similarly to Figure 5, we find the diversification induced by DiffDiv in Table 1 is largely comparable with *ood* data on ColorDSprites, with respectively up to 11% and 14% of the models averting their attention from the main ‘color’ cue in ColorDSprites; up to 23% and 38% of the models averting their attention to the ethnicity cue in UTKFace; and up to 9% and 14% of the models averting their attention to the ethnicity cue in CelebA. Furthermore, we report the change in ensemble accuracy over the averted cues following diversification in Table 2, confirming a positive and significant, increase for all experiments.

Table 2: Average change in ensemble accuracy over the averted cues following diversification (reported in %).

Obj. ↓ Dataset →	Color	Face	CelebA
Cross	5.17	3.37	2.00
Div	8.16	3.61	0.45
KL	2.87	4.54	2.02
L1	3.31	3.61	1.30
L2	4.67	5.04	1.34

4 DISCUSSION AND CONCLUSION

Shortcut learning is a phenomenon undermining the performance and utility of deep learning models, where easy-to-learn cues are preferentially learned for prediction. We propose *DiffDiv*, a framework to achieve shortcut learning mitigation via ensemble diversification on low-fidelity DPM counterfactuals. We show that DPMs can generate novel feature combinations even when trained on images displaying correlated cues. We identify a DPM training stage (*originative*) that is especially useful for ensemble diversification, and show that diffusion-guided diversification can lead to comparable ensemble diversity without the need for expensive *ood* data collection.

REFERENCES

- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J Fleet. Synthetic data from diffusion models improves imagenet classification. *arXiv preprint arXiv:2304.08466*, 2023.
- Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- Yida Chen, Fernanda Viégas, and Martin Wattenberg. Beyond surface statistics: Scene representations in a latent diffusion model. *arXiv preprint arXiv:2306.05720*, 2023.
- Lisa Dunlap, Alyssa Umuno, Han Zhang, Jiezhi Yang, Joseph E Gonzalez, and Trevor Darrell. Diversify your vision datasets with automatic diffusion-based augmentation. *Advances in neural information processing systems*, 36:79024–79034, 2023.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, Nov 2020. ISSN 2522-5839. doi: 10.1038/s42256-020-00257-z. URL <https://doi.org/10.1038/s42256-020-00257-z>.
- Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Phillip Howard, Avinash Madasu, Tiej Le, Gustavo Lujan Moreno, and Vasudev Lal. Probing intersectional biases in vision-language models with counterfactual examples. *arXiv preprint arXiv:2310.02988*, 2023.
- Yeonsung Jung, Hajin Shim, June Yong Yang, and Eunho Yang. Fighting fire with fire: contrastive debiasing without bias-free data via generative bias-transformation. In *International Conference on Machine Learning*, pp. 15435–15450. PMLR, 2023.
- Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representation. *arXiv preprint arXiv:2310.02557*, 2023.
- B. Kim, H. Kim, K. Kim, S. Kim, and J. Kim. Learning not to learn: Training deep neural networks with biased data. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9004–9012, Los Alamitos, CA, USA, jun 2019. IEEE Computer Society. doi: 10.1109/CVPR.2019.00922. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2019.00922>.
- Eungyeup Kim, Jihyeon Lee, and Jaegul Choo. Biaswap: Removing dataset bias with bias-tailored swapping augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14992–15001, 2021.
- Polina Kirichenko, Pavel Izmailov, and Andrew Gordon Wilson. Last layer re-training is sufficient for robustness to spurious correlations. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Zb6c8A-Fghk>.
- Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*, 2022.
- Jungsoo Lee, Eungyeup Kim, Juyoung Lee, Jihyeon Lee, and Jaegul Choo. Learning debiased representation via disentangled feature augmentation. *Advances in Neural Information Processing Systems*, 34:25123–25133, 2021.

- Yoonho Lee, Huaxiu Yao, and Chelsea Finn. Diversify and disambiguate: Learning from underspecified data. *arXiv preprint arXiv:2202.03418*, 2022.
- Y. Li and N. Vasconcelos. Repair: Removing representation bias by dataset resampling. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9564–9573, Los Alamitos, CA, USA, jun 2019. IEEE Computer Society. doi: 10.1109/CVPR.2019.00980. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2019.00980>.
- Jongin Lim, Youngdong Kim, Byungjai Kim, Chanho Ahn, Jinwoo Shin, Eunho Yang, and Seungju Han. Biasadv: Bias-adversarial augmentation for model debiasing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3832–3841, 2023.
- Yong Lin, Lu Tan, Yifan Hao, Honam Wong, Hanze Dong, Weizhong Zhang, Yujiu Yang, and Tong Zhang. Spurious feature diversification improves out-of-distribution generalization. *arXiv preprint arXiv:2309.17230*, 2023.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- Loic Matthey, Irina Higgins, Demis Hassabis, and Alexander Lerchner. dsprites: Disentanglement testing sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Matthias Minderer, Olivier Bachem, Neil Houlsby, and Michael Tschannen. Automatic shortcut removal for self-supervised representation learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML’20*. JMLR.org, 2020.
- Arnab Kumar Mondal, Lakshya Singhal, Piyush Tiwary, Parag Singla, and AP Prathosh. Minority oversampling for imbalanced data via class-preserving regularized auto-encoders. In *International Conference on Artificial Intelligence and Statistics*, pp. 3440–3465. PMLR, 2023.
- Armand Mihai Nicolicioiu, Andrei Liviu Nicolicioiu, Bogdan Alexe, and Damien Teney. Learning diverse features in vision transformers for improved generalization. *arXiv preprint arXiv:2308.16274*, 2023.
- Maya Okawa, Ekdeep Singh Lubana, Robert P Dick, and Hidenori Tanaka. Compositional abilities emerge multiplicatively: Exploring diffusion models on a synthetic task. *arXiv preprint arXiv:2310.09336*, 2023.
- Matteo Pagliardini, Martin Jaggi, François Fleuret, and Sai Praneeth Karimireddy. Agree to disagree: Diversity through disagreement for better transferability. *arXiv preprint arXiv:2202.04414*, 2022.
- Andrew Ross, Weiwei Pan, Leo Celi, and Finale Doshi-Velez. Ensembles of locally independent prediction models. In *AAAI*, 2020.
- Shiori Sagawa*, Pang Wei Koh*, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryxGuJrFvS>.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- Mert Bulent Sariyildiz, Karteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones. In *CVPR 2023—IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- Luca Scimeca, Seong Joon Oh, Sanghyuk Chun, Michael Poli, and Sangdoo Yun. Which shortcut cues will dnns choose? a study from the parameter-space perspective. In *International Conference on Learning Representations*, 2022.

- Luca Scimeca, Alexander Rubinstein, Armand Nicolicioiu, Damien Teney, and Yoshua Bengio. Leveraging diffusion disentangled representations to mitigate shortcuts in underspecified visual tasks. In *NeurIPS 2023 Workshop on Diffusion Models*, 2023. URL <https://openreview.net/forum?id=AvUAVYRA70>.
- Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. On diffusion models for amortized inference: Benchmarking and improving stochastic control and sampling. *arXiv preprint arXiv:2402.05098*, 2024.
- Harshay Shah, Kaustav Tamuly, Aditi Raghunathan, Prateek Jain, and Praneeth Netrapalli. The pitfalls of simplicity bias in neural networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS’20, Red Hook, NY, USA, 2020. Curran Associates Inc. ISBN 9781713829546.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- Damien Teney, Ehsan Abbasnejad, Simon Lucey, and Anton van den Hengel. Evading the simplicity bias: Training a diverse set of models discovers solutions with superior OOD generalization. In *CVPR*, 2022a.
- Damien Teney, Maxime Peyrard, and Ehsan Abbasnejad. Predicting is not understanding: Recognizing and addressing underspecification in machine learning. *arXiv preprint arXiv:2207.02598*, 2022b.
- Antonio Torralba and Alexei A. Efros. Unbiased look at dataset bias. In *CVPR 2011*, pp. 1521–1528, 2011. doi: 10.1109/CVPR.2011.5995347.
- Siddarth Venkatraman, Moksh Jain, Luca Scimeca, Minsu Kim, Marcin Sendera, Mohsin Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, et al. Amortizing intractable inference in diffusion models for vision, language, and control. *arXiv preprint arXiv:2405.20971*, 2024.
- T. Wang, J. Zhao, M. Yatskar, K. Chang, and V. Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5309–5318, Los Alamitos, CA, USA, nov 2019. IEEE Computer Society. doi: 10.1109/ICCV.2019.00541. URL <https://doi.ieeecomputersociety.org/10.1109/ICCV.2019.00541>.
- Xinyue Wang, Yilin Lyu, and Liping Jing. Deep generative model for robust imbalance classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14124–14133, 2020.
- Zihao Wang, Lin Gui, Jeffrey Negrea, and Victor Veitch. Concept algebra for (score-based) text-controlled generative models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Qiucheng Wu, Yujian Liu, Handong Zhao, Ajinkya Kale, Trung Bui, Tong Yu, Zhe Lin, Yang Zhang, and Shiyu Chang. Uncovering the disentanglement capability in text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1900–1910, 2023.
- Kai Yuanqing Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=gl3D-xY7wLq>.
- Tian Xu, Jennifer White, Sinan Kalkan, and Hatice Gunes. Investigating bias and fairness in facial expression recognition. In *Computer Vision – ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part VI*, pp. 506–523, Berlin, Heidelberg, 2020. Springer-Verlag. ISBN 978-3-030-65413-9. doi: 10.1007/978-3-030-65414-6_35. URL https://doi.org/10.1007/978-3-030-65414-6_35.

Jianhao Yuan, Francesco Pinto, Adam Davies, Aarushi Gupta, and Philip Torr. Not just pretty pictures: Text-to-image generators enable interpretable interventions for robust representations. *arXiv preprint arXiv:2212.11237*, 2022.

John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl Oermann. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS Med.*, 15(11):e1002683, November 2018.

Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5810–5818, 2017.

S1 RELATED WORK

S1.1 OVERCOMING THE SIMPLICITY BIAS

Deliberate de-bias: To overcome the simplicity bias, copious literature has explored methodologies to avoid or mitigate shortcut cue learning when labels for the shortcut cues were present in the training data (Li & Vasconcelos, 2019; Kirichenko et al., 2023; Wang et al., 2019; Kim et al., 2019; Sagawa* et al., 2020; Lee et al., 2021). The access to shortcut signals is, however, a critical limitation, as these are generally hard or impossible to obtain.

Data augmentation: Data augmentation methodologies have proven useful in the generation of *bias-conflicting* or bias-free samples (Kim et al., 2021; Lim et al., 2023; Jung et al., 2023), or to augment underrepresented subgroups therein (Wang et al., 2020; Mondal et al., 2023), leading to less biased predictions. Although an important research direction, assumptions on the nature of the biases are here still necessary, and the rejection of selected cues for prediction may not always lead to improved generalization, notably when the downstream task is aligned with said cues.

Diversification: A different approach to this problem has been to enforce the use of a diverse set of signals for prediction. The use of ensembles has been one such method, where models’ diversity would lead to bias mitigation and improved generalization. Different weight initialization and architectures have previously been shown to be ineffective in the presence of strong shortcut biases (Scimeca et al., 2022). Methods ensuring mutual orthogonality of the input gradients have proven more effective, driving models to attend to different locations of the input space for prediction (Ross et al., 2020; Teney et al., 2022a;b; Nicolicioiu et al., 2023). These input-centric methods, however, may be at a disadvantage in cases where different features must attend to the same area of the input space. Instead, a different approach has been via the diversification of direct ensemble model predictions. This approach hinges on the availability of –unlabelled– out-of-distribution (*ood*) auxiliary samples that are, at least in part, free of the same shortcuts as the original training data. Through a diversification objective, the models are then made to disagree on these *ood* samples, while maintaining performance on the original data, effectively fitting functions with different extrapolation behaviors (Lee et al., 2022; Pagliardini et al., 2022; Scimeca et al., 2023; Lin et al., 2023). Even in this case, the auxiliary *ood* data dependency poses limitations, as this is often not readily accessible, and can be costly to procure.

S1.2 DPM MODELLING AND GENERALIZATION

In recent years, Diffusion Probabilistic Models have emerged as a transformative generative tool, spearheading progress in the generative domain across various applications, including vision (Ho et al., 2020; Song & Ermon, 2019), efficient sampling methods (Sendera et al., 2024), and text (Gong et al., 2022; Venkatraman et al., 2024) and even control tasks (Venkatraman et al., 2024). Numerous studies have underscored their prowess in generating synthetic images, which can then be harnessed to enrich datasets and bolster classification performance (Sariyildiz et al., 2023; Azizi et al., 2023; Yuan et al., 2022; Dunlap et al., 2023; Howard et al., 2023).

In several cases, DPMs have been shown to transcend the surface-level statistics of the data, making them invaluable in understanding data distributions and features (Chen et al., 2023; Yuan et al., 2022; Wu et al., 2023). Moreover, recent studies have indicated the ability of DPMs to achieve feature disentanglement via denoising reconstructions (Kwon et al., 2022; Wu et al., 2023; Okawa et al., 2023). In particular, work in (Wang et al., 2023) has shown how text-guided generative models can represent disentangled concepts, and how, through algebraic manipulation of their latent representations, it is possible to compositionally generalize to novel and unlikely combination of image features.

Additional work has also shown strong inductive Diffusion biases during learning, which may explain some of these phenomena (Kadkhodaie et al., 2023). In our work, we test the edge case of DPMs trained with data exhibiting correlated input features. We find that when suitably trained, DPMs can still generate samples with novel feature combinations and that these can be leveraged for ensemble diversity.

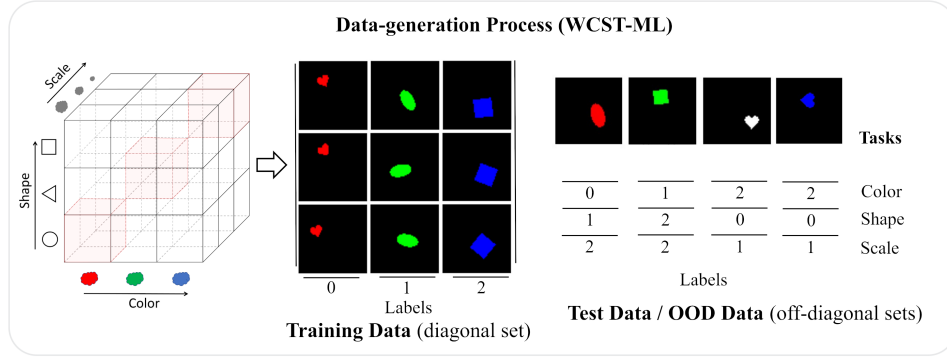


Figure S1: We consider the WCST-ML setting (Scimeca et al., 2022) as the experimental ground in our experiments. Each dataset is partitioned into training data where the task labels are perfectly correlated with the image input features (e.g. {color, shape, and scale}) and Test data, where samples are used to test a model’s tendency to a feature over another. We use a random subset of Test Data for *ood* experiments.

S2 SUPPLEMENTARY METHODS

S2.1 DATASETS

In this work, we leverage three representative datasets: a color-augmented version of DSprites (Matthey et al., 2017), UTKFace (Zhang et al., 2017), and CelebA (Liu et al., 2015).

DSprites: DSprites includes a comprehensive set of symbolic objects generated with variations in five latent variables: shape, scale, orientation, and X-Y position. We augment this dataset with a color dimension and remove redundant samples (e.g. due to rotations), resulting in 2,862,824 distinct images, which we refer to as ColorDSprites. ColorDSprites permits the examination of shortcut biases in a highly controlled setup.

UTKFace: UTKFace provides a dataset of 23,708 facial images annotated with attributes like age, gender, and ethnicity. Unlike DSprites, UTKFace presents a real-world, less controlled setup to study bias. Its inherent complexity and diversity make it an ideal candidate for understanding the model’s cue preferences when societal and ethical concerns are at stake.

CelebA: CelebA is a large-scale dataset comprising 202,599 celebrity facial images annotated with 40 different attributes, including gender, age, and facial features. It provides a more extensive and diverse real-world dataset compared to UTKFace, offering rich variations in pose, background, and lighting. CelebA is commonly used for studying bias and fairness in models due to its attribute diversity and challenging conditions.

S2.2 WISCONSIN CARD SORTING TEST FOR MACHINE LEARNERS (WCST-ML)

To isolate and investigate shortcut biases, we employ the Wisconsin Card Sorting Test for Machine Learners (WCST-ML), a method devised to dissect the shortcut learning behaviors of deep neural networks (Scimeca et al., 2022). We use the splits from WCST-ML to both train and evaluate the ensembles in this work, as it provides a systematic approach to creating datasets with multiple cues, designed to correlate with the target labels. Specifically, given K cues i_1, i_2, \dots, i_K , the method produces a *diagonal* dataset $\mathcal{D}_{\text{diag}}$ where each cue i_k is equally useful for predicting the labels Y , with the total number of classes $L = |Y|$. This level playing field is instrumental in removing the influence of feature dominance and spurious correlations, thereby allowing us to observe a model’s preference for certain cues under controlled conditions. To rigorously test these preferences, WCST-ML employs the notion of *off-diagonal* samples. These are samples where the cues are not in a one-to-one correspondence with the labels, but instead align with only one of the features under

inspection. By evaluating a model’s performance on off-diagonal samples, according to each feature, we can test and achieve an estimate of a model’s reliance on the same.

S2.3 OPERATIONALIZING WCST-ML ACROSS DATASETS

We follow the set-up in (Scimeca et al., 2022) and construct a balanced dataset $\mathcal{D}_{\text{diag}}$, which includes a balanced distribution of cues, coupled with their corresponding off-diagonal test sets (one for each feature) Figure S1. For both datasets, we define a balanced number of classes L for each feature under investigation. Where the number of feature values exceeds L , we dynamically choose ranges to maintain sample balance with respect to each new feature class. For instance, for the continuous feature ‘age’ in UTKFace, we dynamically select age intervals to ensure the same L number of categories as other classes, as well as sample balance within each category. We consider sets of features previously found to lead to strong simplicity biases. For ColorDSprites, we consider $K_{DS} = 4$, features $\{color, orientation, scale, shape\}$, and $L = 3$ as constrained by the number of shapes in the dataset. Within UTKFace we consider $K_{UTK} = 3$, features $\{ethnicity, gender, age\}$, and $L = 2$ as constrained by the binary classification on *gender*. For CelebA we consider $K_{CL} = 2$, features $\{lightskin, ovalface\}$, and $L = 2$ as enforced by the binary labels on all features. For each dataset we create one *diagonal* subset of fully correlated features and labels, available at training time, and K_{DS} , K_{UTK} and K_{CL} feature-specific *off-diagonal* datasets to serve for testing the models’ shortcut bias tendencies.

S2.4 DPM TRAINING AND SYNTHETIC COUNTERFACTUAL GENERATION

We utilize Diffusion Probabilistic Models (DPMs) to generate synthetic data for our experiments. DPMs operate by iteratively adding or removing noise from an initial data point x through a stochastic process governed by a predefined noise schedule. We base our training regime on (Ho et al., 2020). As denoiser, we train a classic U-Net architecture with 4 down-sampling blocks and 4 up-sampling blocks with $\approx 9mil$ parameters for all datasets. We train the model with the objective in Equation 3 by iterating through the relevant dataset over a maximum of 1200 epochs. We use a vanilla Adam optimizer in all experiments. All DPM schemes use a time discretization of 1000 steps. To facilitate efficient sampling, we employ Denoising Diffusion Implicit Models (DDIM) (Song et al., 2020), a first-order ODE solver for DPMs (Salimans & Ho, 2022; Lu et al., 2022), utilizing a predictor-corrector scheme to minimize the number of sampling steps, lowering the final number to 250 during sampling in framework.

For each of the experiments, we use the trained DPM to generate a fixed dataset of 3000 synthetic counterfactuals, which is independently shuffled and batched-sampled for the diversification objective during ensemble training. We perform ablation studies considering a larger batch of synthetic counterfactuals (equal to the number of data points in each dataset) in §S3.5, with only marginal performance gains compared to the smaller set.

S3 SUPPLEMENTARY RESULTS

S3.1 γ SELECTION

We perform a hyper-parameter search to find the values of γ for each diversification objective. We perform the search via the same methodology used in the *ood* main experiments, i.e. by training an ensemble of 100 ResNet-18 models on the fully correlated *diagonal* datasets, while diversifying a small subset of 30% of the original training data, randomly sampled from the de-correlated left-out set. We consider γ values ranging from $1e-3$ to $1e1$, and monitor both the validation ensemble accuracy as well as the predictive diversity on a separate de-correlated set of validation data. Figure S2 shows the performance of each metric for different values of γ . We select the values reported in Table S1 to be at the intersection of the accuracy and diversification trends for each model.

S3.2 MODEL SELECTION TO BOOST DIVERSE ENSEMBLE PERFORMANCE

The mitigation of Shortcut leaning through diversification is generally known in the literature to suffer from a decrease in ensemble ID performance, as we also observe and study in our experiments. A

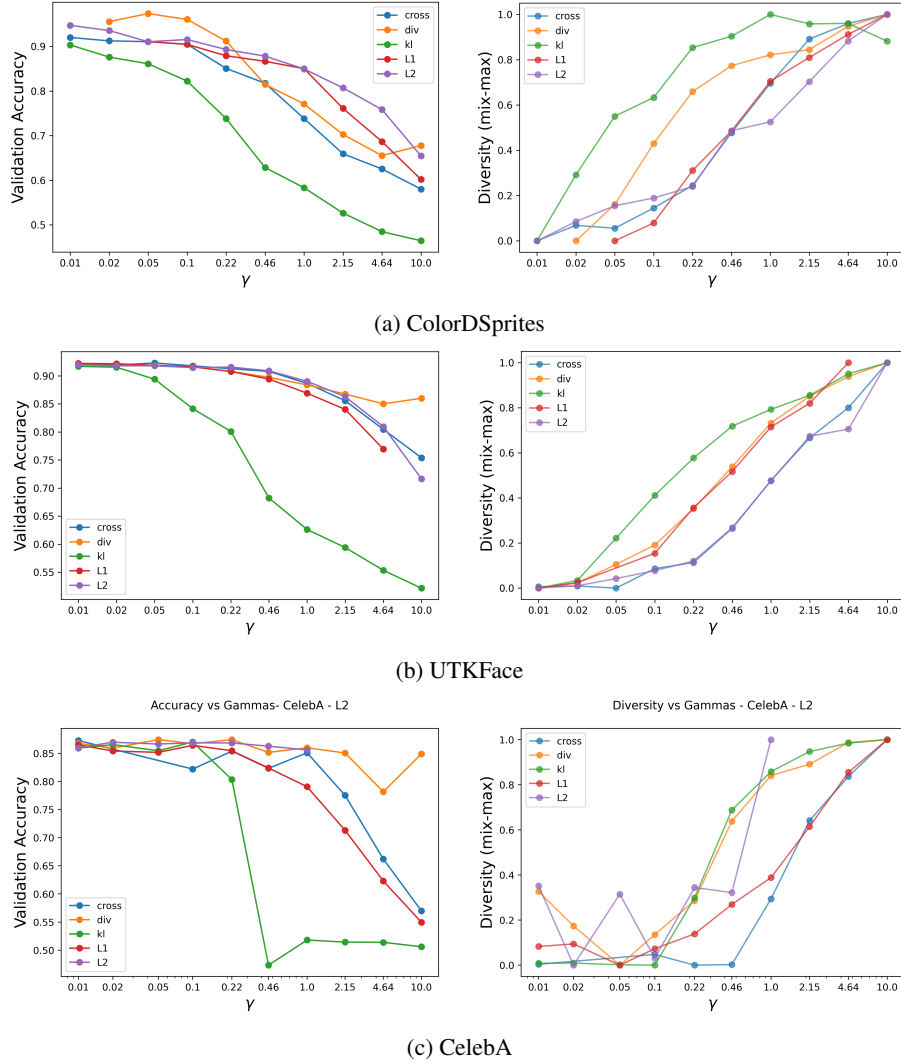


Figure S2: Hyper-parameter search on the disagreement intensity (γ) for each diversification objective. For the main experiments, we select the values of γ at the intersection of the accuracy-diversity trends by each objective (Table S1).

prominent methodology to mitigate this phenomenon in the literature is ensemble model selection, where a subset of models is selected for final ensemble inference. We perform additional experiments to assess the degree by which model selection can aid in ensemble performance within the DiffDiv framework. To ensure diversity in the final selection, we include in the selected subset any model showing shortcut-cue aversion from Table 1. Furthermore, we select additional models to reach a dynamic range between 15% and 99% of the original ensemble. We compare the performance of the ensemble before and after model selection in Table S2.

S3.3 ON THE INFLUENCE OF DPM FIDELITY TO DIVERSIFICATION

We perform experiments whereby an ensemble of 100 ResNet-18 models is trained separately with respect to all diversification objectives considered. Figure S3 shows the diversification results as a function of the fidelity of the DPM used to sample the diversification set. Although we observe that the diversification level obtained is dependent on the diffusion fidelity level, suggesting the need for appropriate early stopping procedures to achieve increased ensemble prediction diversity, we find DiffDiv to not be overly sensitive to this choice. In fact, we observe broad areas with similar diversity

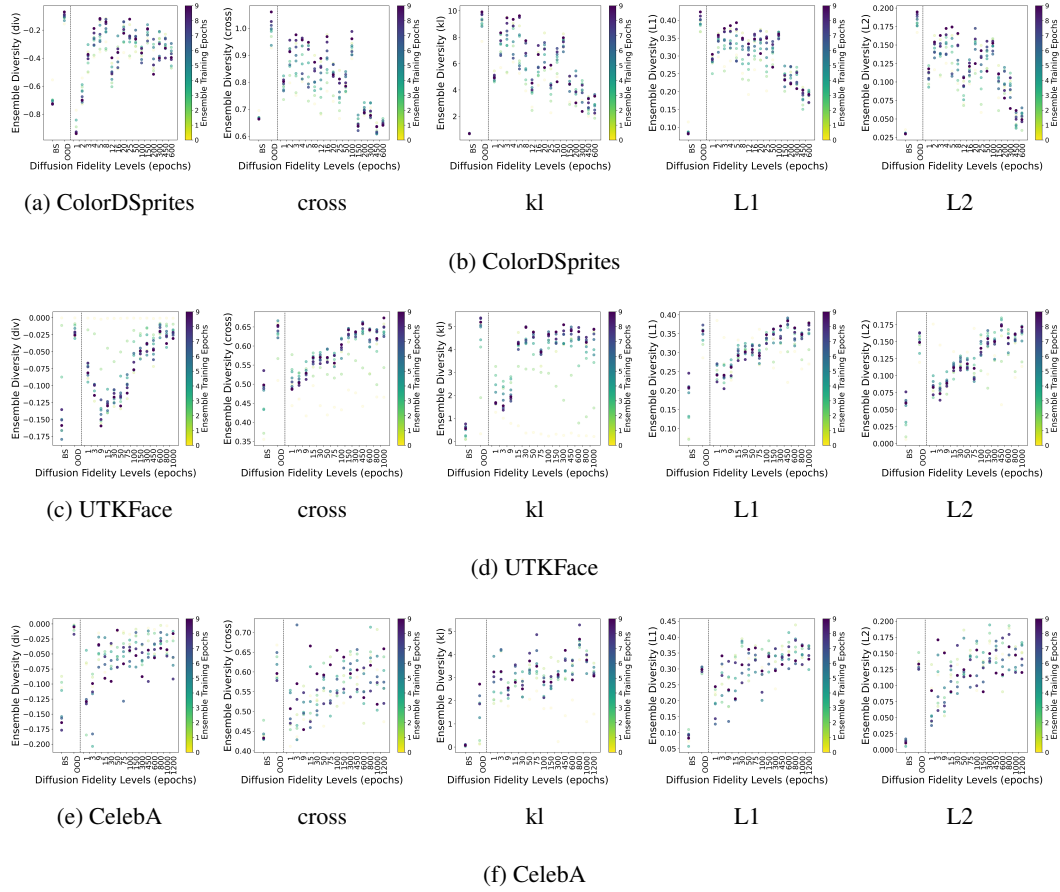


Figure S3: Ensemble diversity as enforced via samples from diffusion models trained a different fidelity levels (i.e. diffusion training epochs).

levels across several DPM fidelities. While the trends vary mildly across different disagreement objectives, we find the diversification maximized in ColorDSprites and UTKFace coherently with the analysis in Section 3.2, where improved diversification is achieved around the *originative* interval. For ColorDSprites, this is realized in around 20 DPM training epochs, for UTKFace, it is realized in around 800 DPM training epochs, while for CelebA it is around 1000 DPM training epochs.

Interestingly, our results suggest how ensemble diversification metrics can be a viable proxy for appropriate *originative* DPM training. In Figure S4, we observe the min-maxed change in accuracy and change in diversity by the ensembles with respect to baseline training on ColorDSprites. In the figure, we observe similar trends across all diversification methods, whereby the *originative* interval (highlighted in gray) is primarily identified by the highest changes in diversity, while also considering the least drop in accuracy. These results confirm our previous supervised findings (Figure 3). Importantly, we find that only when jointly looking at both diversity and validation performance we can best identify the relevant areas around the generative stage. Under this light, ensemble performance/diversity validation metrics can be directly leveraged for DPM early stopping logic. Our results align with our previous observations, where the highest number of ood-generated samples to lie within the DPM training intervals achieve the highest change in diversity while maintaining good classification performance (Figure S6.)

S3.4 DIVERSIFICATION LEADS TO ENSEMBLE MODELS ATTENDING TO DIFFERENT CUES

Figure S5 illustrates a feature-centric description of 10 ensemble models trained with a diversification objective on *ood* data (a) and Diffusion generated counterfactuals (b). The variation across models is evident: several models substantially reduce their dependency on the leading cues of the respective

Table S3: Diversification results on ColorDSprites, UTKFace, and CelebA when using the same number of *ood* samples as the training dataset. The feature columns report the fraction of models (in each row) biased towards the feature. The final column reports the average validation accuracy for the ensemble when tested on a left-out feature-correlated *diagonal* set, of the same distribution as the original training data

Dataset →	ColorDSprites					UTKFace				CelebA		
Obj. ↓	Color (↓)	Orient.	Scale	Shape	Acc. (↑)	Age	Ethnicity (↓)	Gender	Acc. (↑)	Oval Face	Pale Skin (↓)	Acc. (↑)
Baseline	1.00	0.00	0.00	0.00	1.000±0.00	0.00	1.00	0.00	0.920±0.02	0.00	1.00	0.857±0.01
Cross	0.92	0.00	0.08	0.00	0.849±0.16	0.00	0.73	0.27	0.858±0.04	0.01	0.99	0.751±0.07
Div	0.82	0.00	0.18	0.00	0.820±0.19	0.00	0.76	0.24	0.844±0.04	0.00	1.00	0.843±0.03
KL	0.90	0.03	0.05	0.02	0.801±0.20	0.02	0.68	0.30	0.820±0.07	0.00	1.00	0.812±0.04
L1	0.90	0.01	0.07	0.02	0.799±0.21	0.00	0.66	0.34	0.832±0.09	0.14	0.86	0.724±0.11
L2	0.86	0.04	0.08	0.02	0.745±0.22	0.08	0.58	0.34	0.761±0.15	0.12	0.88	0.651±0.10

S3.5 ON THE INFLUENCE OF AN INCREASED NUMBER OF *ood* SAMPLES FOR ENSEMBLE DISAGREEMENT

As per the original objective, with DPM sampling we aim to circumvent the diversification dependency on Out-Of-Distribution data, which is often not readily accessible and can be costly to procure. We test this dependency further and assess the quality of the diversification results when matching the number of *ood* data used for diversification to the original training data for the ensemble. We report in Table S3 our findings. We observe the quality of the disagreement on ColorDSprites to only marginally benefit from additional disagreement samples, with approximately 1% to 7% more of the models to avert their attention from the shortcut cue *color* as compared to the original experiments. On the other hand, we observe a strong improvement in the diversification for UTKFace, mainly registered via the *div* objective, where 24% of the models averted their attention from the *ethnicity* shortcut, as opposed to the original 6% in our previous experiments, while maintaining high predictive performance on the validation set. We observe marginal improvements on the other objectives, with approximately 4% to 8% additional models achieving cue aversion. We speculate this gain to be due to the higher complexity of the features within the data, which may require additional specimens for appropriate diversification.

Increased Ensemble Disagreement Negatively Correlates with Ensemble *id* Performance:

Achieving ensemble diversity by disagreement on unlabelled *ood* samples has previously been shown to negatively impact *id* ensemble performance (Pagliardini et al., 2022) (suppl. Figure S2), while improving *ood* performance on downstream tasks associated with non-shortcut cues. Generally, ensemble model selection has been an effective method to prune models that misalign with the original classification objective (Lee et al., 2022). In this section, we wish to understand this relationship under the scope of Diffusion-guided diversification. In Figure S6 we observe the change in ensemble average accuracy Δ_{acc} as a function of the change in diversity Δ_{ds} , spanning all the diversification objectives. We highlight a few important observations. First, increased ensemble diversity via disagreement negatively correlates with average ensemble accuracy, without additional model selection mechanisms. Second, real *ood* data achieves the highest diversity gain, comparable to the diffusion samples at the *originative* stage (§3.2), $\approx [10, 100]$ for ColorDSprites, $\approx [450, 800]$ for UTKFace, and $\approx [1000, 1200]$ for CelebA. Lastly, the diversification objectives show comparable trends, with the *div* objective displaying marginally better diversification/accuracy performance than the others on both tasks.

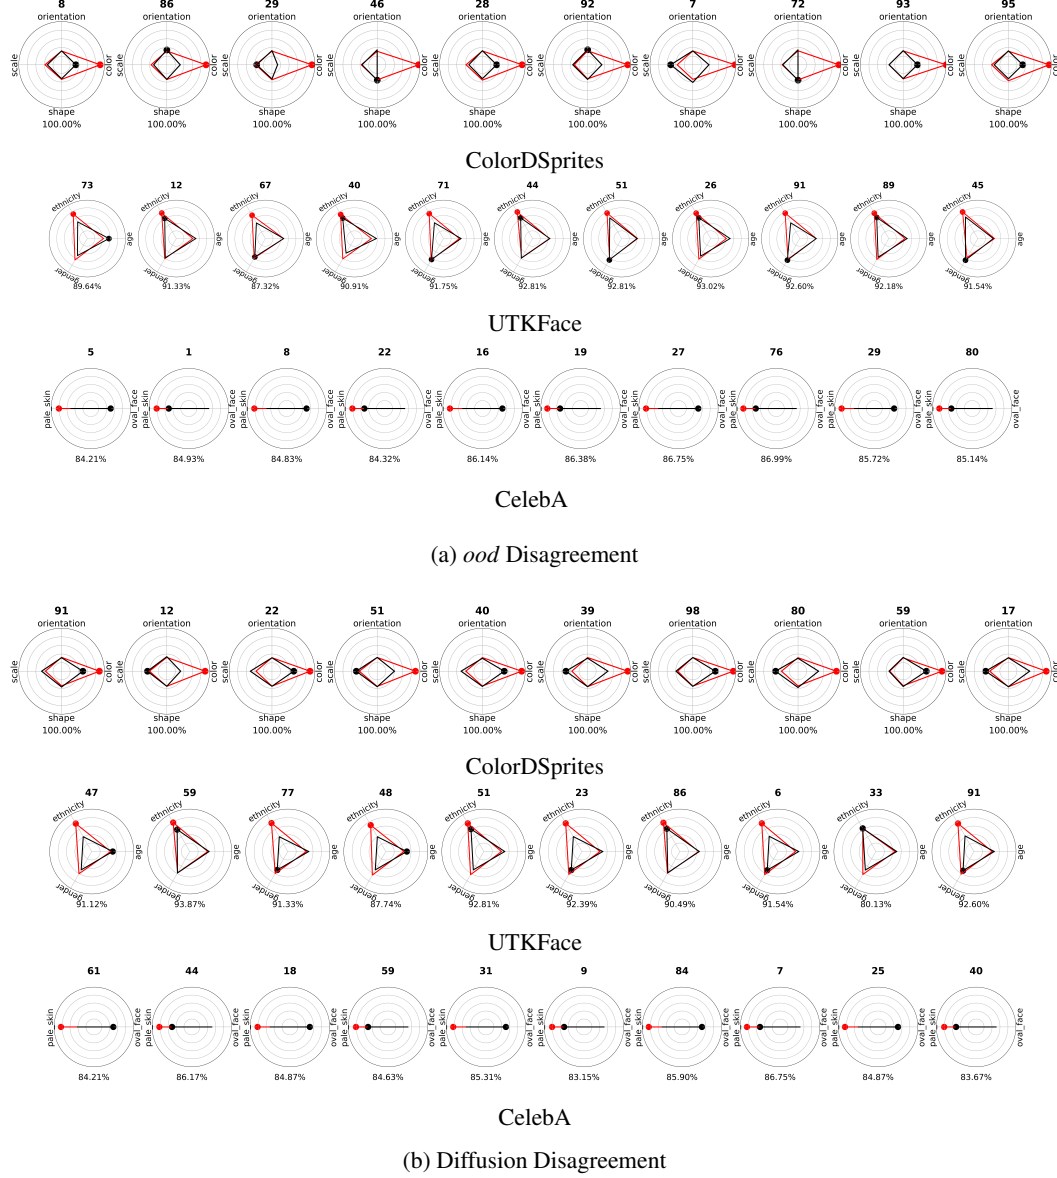


Figure S5: Comparison of 10 diversified models when training the ensemble while using (a) feature-uncorrelated *ood* data and (b) Diffusion Samples.

