
Combining Counterfactuals With Shapley Values To Explain Image Models

Aditya Lahiri¹ Kamran Alipour¹ Ehsan Adeli²
Babak Salimi¹

Abstract

With the widespread use of sophisticated machine learning models in sensitive applications, understanding their decision-making has become an essential task. Models trained on tabular data have witnessed significant progress in explanations of their underlying decision making processes by virtue of having a small number of discrete features. However, applying these methods to high-dimensional inputs such as images is not a trivial task. Images are composed of pixels at an atomic level and do not carry any interpretability by themselves. In this work, we seek to use annotated high-level interpretable features of images to provide explanations. We leverage the Shapley value framework from Game Theory, which has garnered wide acceptance in general XAI problems. By developing a pipeline to generate counterfactuals and subsequently using it to estimate Shapley values, we obtain contrastive and interpretable explanations with strong axiomatic guarantees.

1. Introduction

Understanding the decision-making rationale behind complicated black-box models has emerged as one of the most critical tasks in the greater overarching goal of making AI systems transparent and trustworthy. Amidst several proposed methods, the well-studied concept of Shapley values (Winter, 2002) from Game Theory literature has emerged as a principled framework to obtain feature attributions as explanations. By virtue of their strong axiomatic guarantees, they lend themselves favorably to the task of distributing model outputs fairly among the different input features. However, one of the major issues in using Shapley values for model explanation is the computation time (Van den Broeck et al., 2021). It grows exponentially with the number of features

¹University Of California, San Diego ²Stanford University.
Correspondence to: Aditya Lahiri <adlahiri@ucsd.edu>.

involved. This problem is particularly exacerbated in end-to-end models that operate directly on images since, usually, the features in images are defined on a pixel level. This results in thousands of features. In addition, the individual pixel values are not interpretable to humans, and therefore any attribution attached to individual pixels does not help with high-level model understanding. It has also been shown that humans understand and explain things by comparison and contrast (De Graaf & Malle, 2017). On their own, feature attribution numbers obtained by Shapley values do not allow for this contrastive notion and lack a reference or visual aid to compare to. This has been an essential barrier for generating explanations based on Shapley values (Kumar et al., 2020).

To close the gap, we propose an approach that incorporates high-level interpretable features and employs generative models to produce counterfactual images corresponding to specific changes in the interpretable features. These counterfactuals are then used to compute Shapley values which explain the difference in prediction scores between the original and counterfactual image. This process enables us to understand model behavior using the contrastive explanations (w.r.t. Shapley values) provided for arbitrary input images.

2. Related Work

Shapley values have been widely used in the context of Explainable AI (XAI) to provide feature attributions as explanations (Lundberg & Lee, 2017; Datta et al., 2016). However, since computing Shapley values is intractable (Van den Broeck et al., 2021), various approximations have been used (Sundararajan & Najmi, 2020). For models that work on images as inputs, popular techniques include aggregating neighboring pixels to form sub-pixels (Ribeiro et al., 2016) to be more interpretable. Other works require gradients to obtain pixel-level explanations and try to compute them efficiently (Li et al., 2020). We propose a model-agnostic approach that generates explanations in terms of a limited number of high-level human interpretable features.

Significant work has gone in to producing contrastive explanations (Kommiya Mothilal et al., 2021; Galhotra et al., 2021). These have primarily been shown to work on structured tabular data. Counterfactual image generation has also been an active space of research (Nemirovsky et al., 2020; Samangouei et al., 2018). Often they require re-training a generative model (Dash et al., 2022), which is expensive



Figure 1. Examples of original images and their corresponding counterfactual images. The values below the images represent the model’s prediction for Attractiveness. We generate the counterfactual images according to known specifications. For instance, C_1 was obtained from O_1 by increasing baldness and maleness while decreasing youngness, heavy makeup, and blond hair. Similarly, C_4 was generated from O_4 by increasing youngness, blond hair, heavy makeup, and decreasing maleness and baldness. We use Shapley values to explain the difference between the original images and their counterfactual images in terms of the modified interpretable attributes in Table 1.

or are specific to the choice of generative model (Abdal et al., 2021). We provide a minimal and scalable training process by utilizing only pretrained GANs for generating counterfactual images.

Recent work by (de Mijolla et al., 2020) comes closest to our work and seeks to use Shapley values to explain models trained on high dimensional data. They use classical computer vision techniques to generate interpretable features. In contrast, we use generative models to produce counterfactual images to compute Shapley values efficiently while also providing contrastive explanations.

3. Background

3.1. Contrastive Counterfactuals & Generative Models

Contrastive counterfactuals have been the building blocks of ideas in philosophy and cognition that guide people’s understanding and dictate how we explain things to one another (De Graaf & Malle, 2017) and have been argued to be central to explainable AI (Miller, 2019). We are specifically interested in the implementation of this framework in the

image classification problem in order to allow us to generate counterfactual images at will such that they increase or decrease the presence of a set of interpretable feature attributes. Generative models are vastly popular in different fields of AI, and their recent advancements in creating realistic images have made them a viable approach to produce a latent representation of an image dataset. In our experiments, we utilize StyleGAN2 (Karras et al., 2020) as a state-of-the-art generative model which can be used to generate high resolution and realistic images in different domains. StyleGAN feeds the latent variable into a mapping network that transforms it into an intermediate latent variable. Aside from its ability to produce styles, this transformation also provides the intermediary latent space as a more regulated domain to learn and traverse through interpretable attributes. We use the manipulation of this latent space to provide us with realistic counterfactual images according to our specifications.

3.2. Shapley Values

Shapley value is a concept from Game Theory that provides a unique solution to fairly allocating the total payout from a

game to its individual players. A coalition is a set of players playing the game. A grand coalition contains all the players, while an empty coalition contains none. A value function $v(S)$ provides the scores obtained by the coalition S on playing the game. The score obtained by playing the game is called the payoff.

In the context of XAI, Shapley values have been used as a means to obtain feature contribution. The input features are the players in this game of obtaining predictions from the model. The prediction is then fairly distributed among the input features. For a coalition set S , the following formula provides the Shapley value ϕ for feature i :

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S|)!}{n!} (v(S \cup \{i\}) - v(S)) \quad (1)$$

where $n = |N|$ is the total number of input features. This is the weighted sum of marginal contribution of feature i across all possible coalitions that do not contain the feature i . Shapley values satisfy desirable axioms that make them make a good choice for generating feature contributions. We describe 3 of those axioms now-

Null : The null axioms states that if a feature that does not change the output when added to any coalition, it will get a Shapley value of zero.

Efficiency : The efficiency property states that the sum of Shapley values is equal to the difference between the prediction obtained by the grand coalition and empty coalition.

Symmetry: If two features behave in the same way across all possible coalitions, then their Shapley values will end up being the same.

The grand coalition is usually the entire instance to be explained, while the null coalition is an average instance or an instance composed entirely of default replacements for missing values for each feature (since they are all out of the coalition in the null coalition). The Shapley values are then used to decompose the difference between an instance's prediction and the average prediction of the model to the individual features of the instance through the efficiency axiom. Since we iterate over all possible subsets of features that do not contain the i th feature, computing Shapley values is exponential in the number of features. This is particularly problematic in case of images, where we have large numbers of pixels as input features. For partial coalitions, we need to compute the model output given only the members in the coalition, which results in partially formed instances. This is usually done by either marginalizing (Datta et al., 2016) or conditioning (Lundberg & Lee, 2017). However, both of these techniques are known to have issues (Kumar et al., 2020). Further, this vanilla Shapley value-based decomposition explains away the difference between the prediction of the input instance and the model prediction on an "average" instance. This average model prediction does not always correspond to a sensible input (Sundararajan & Najmi, 2020).

It has also been shown that standard feature contributions in isolation do not help humans reason as well as contrastive explanations do (De Graaf & Malle, 2017). We aim to overcome these issues through our modified contrastive formulation of Shapley values using high-level interpretable input features as players.

4. Methodology

4.1. Generating Counterfactual Images

We use a Shift Predictor model to obtain counterfactual images according to our needs. A shift predictor model is an MLP model that takes latent variable of an image from a generative model G and generates the latent variable for its counterfactual based on the attributes produced by a classifier. For a generative model $G : R^d \rightarrow R^n$ having a latent space with dimension d and a classifier $C : R^n \rightarrow R^m$ that predicts m attributes, we define our shift predictor as $M(\mathbf{z}, \hat{\mathbf{y}}) : R^d \times R^m \rightarrow R^d$, where $\mathbf{z} \in R^d$ is the latent variable for the input image and $\hat{\mathbf{y}}$ denotes the attributes for the intended counterfactuals. During training, shift predictor learns the directions in the latent space of G that correspond to changes in the attributes predicted by the classifier. Without need for any manual labeling, the training procedure only requires the latent variables of images from G to input the shift predictor and supervise it with the labels generated by the classifier.

During training, shift predictor learns to produce a counterfactual latent variable that satisfies any combination of attributes defined by $\hat{\mathbf{y}}$. In other words, if the classifier predicts a set of attributes $A = \{A_1, A_2, \dots, A_m\}$, shift predictor provides a counterfactual latent variable compatible with any selected subset of attributes A :

$$\hat{\mathbf{z}} = M(\mathbf{z}, \{A_i \in A\}) \quad (2)$$

Under proper training, the shift predictor is an approximation of latent variable distribution conditioned by the subset of attributes A :

$$\hat{\mathbf{z}} \sim P(\mathbf{z} | \{A_i \in A\}) \quad (3)$$

The loss function in the training process pursues two objectives: 1) minimizing the error in prediction of attributes A for the counterfactual image, 2) assuring a level of faithfulness of the generated counterfactual to the original image. The attribute loss L_a is defined as a cross entropy between the conditioned attributes and the attributes predicted by the classifier. During training, the conditioned attributes A are distinguished from unset attributes so the loss will be only calculated for them. On the other hand, the faithfulness loss L_f is calculated as the normal distance between the original latent variables and their counterfactuals. The overall loss in the training process is defined as a combination of these

two losses with a faithfulness factor γ which defines a balance between attribute accuracy of counterfactuals and their faithfulness to the original input:

$$L = L_a + \gamma L_f = \sum_{A_i \in \mathcal{A}} \hat{y}_i \log(y_i) + \gamma \sum_{j \in \mathcal{A}} z_{jj} \quad (4)$$

4.2. Contrastive Explanations Using Shapley Values

The efficiency property of Shapley values allows us to decompose the score difference between the game played by the grand coalition and the empty coalition among all the players. We define players in terms of interpretable attributes of the images. In images of human faces, these attributes could be attributes like hair, makeup, etc. We work in this interpretable space instead of pixel-level features. We define the empty coalition as the set of all zeros for these attributes. It corresponds to their natural or default values that they take on in the image. However, we can define the grand coalition by either increasing the attributes or decreasing them. We use the shift predictor described above to obtain images corresponding to this increase or decrease in the specific attribute of the image. For instance, we can define the grand coalition to be +1 for all interpretable features, which would correspond to an image that has been forced to increase the presence of all attributes in itself. This provides us with 2 images to compare - the original image corresponding to the empty coalition and the counterfactual image with a set of increased or reduced attributes (+1, -1). We use Shapley values to decompose the difference in prediction obtained for these two images into their interpretable features. Our shift predictor can take any arbitrary vector specifying the direction of change of attributes and return a counterfactual image. This overcomes issues found in other traditional Shapley value-based frameworks of missing data for members out of coalition. For partial coalitions, we set the values for members out of the coalition to be 0. This instructs the shift predictor to not make any changes to that attribute. Hence, we directly compute value functions for any coalition using this pipeline. In this setup, we gain a contrastive notion and no longer compare to an "average" instance. Instead, we have two specific images- the original and its counterfactual, and we explain the difference in their predictions Shapley values.

5. Experiments

We run experiments to explain a classifier that is trained on face images. We have annotations for the set of interpretable attributes A that we will use to explain the model's behavior. We train a multi-task classifier built on top of a pretrained VGG cite backbone to predict the set of interpretable attributes A for new unseen images. The CelebA dataset (Liu et al., 2015) is used as the training set and provides 39 binary attributes, including attractiveness which

Table 1. Shapley Value based contributions explaining the difference in predictions between each pair of original and counterfactual images from Figure 1. Each element is of the form (\uparrow / \downarrow , attribution) where the \uparrow / \downarrow indicates whether that specific feature in the original image was increased or decreased to produce the counterfactual.

IMAGE	YOUNG	HEAVY MAKEUP	BLOND HAIR	BALD	MALE
1	$\downarrow -0.28$	$\downarrow -0.02$	$\downarrow -0.03$	$\uparrow -0.34$	$\uparrow 0.07$
2	$\downarrow -0.20$	$\downarrow -0.07$	$\downarrow -0.03$	$\uparrow -0.23$	$\downarrow -0.04$
3	$\uparrow 0.23$	$\uparrow 0.37$	$\uparrow 0.15$	$\downarrow 0.10$	$\uparrow -0.05$
4	$\uparrow 0.16$	$\uparrow 0.18$	$\uparrow 0.13$	$\downarrow 0.23$	$\downarrow 0.09$

we use as the target output Y for any black-box classifier of choice. As the set of interpretable explanatory attributes (A), we choose five other labels: blonde hair, heavy makeup, baldness, youngness, and maleness. We model attractiveness as the positive class ($y = 1.0$) and unattractiveness as the negative class ($y = 0.0$) for the classifier to predict. The black-box classifier is built on a pretrained VGG backbone. We train our shift predictor model, as described earlier, with a faithfulness γ value of 0.09. We pass the images through the multi-task classifier to obtain the attribute values and through the black-box classifier to obtain the target label. We sample images randomly from the StyleGan2 Generative Model and pass them through the black-box model. For every instance, we generate a counterfactual image through the shift predictor by defining our grand coalition in terms of increase or decrease of feature attributes. We obtain Shapley values-based contributions that explain away the difference in prediction between the original image(empty coalition) and the counterfactual image(grand coalition). Both the original and contrastive images along with the Shapley values of the interpretable features are given as explanations. We report the Shapley value-based attributions for images in Figure 1 in Table 1. The original images have their default attributes, while counterfactual images are generated through modifying the interpretable images as shown by " and # signs for the increase and decrease of attributes, respectively, as listed in Table 1. We can observe that for the original image O_1 , we generate the counterfactual image C_i by increasing maleness and baldness while decreasing youngness, heavy makeup, and blond hair. When this is done, the attractiveness score drops from 0.73 to 0.12. This drop is mainly due to increasing baldness which accounts for 0.34 of the 0.61 difference in attractiveness, while decreasing youngness(increasing age) contributed to 0.28 of the total drop. Similarly, we can look at other pairs of original images and counterfactuals, in conjunction with the assigned Shapley attributions to understand model behaviour.

6. Future Work

As next steps, we want to make the shift predictor causal by incorporating causal graphs and also work towards a computationally efficient algorithm to compute these explanations.

References

- Abdal, R., Zhu, P., Mitra, N. J., and Wonka, P. Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows. *ACM Transactions on Graphics (TOG)*, 40(3):1–21, 2021.
- Dash, S., Balasubramanian, V. N., and Sharma, A. Evaluating and mitigating bias in image classifiers: A causal perspective using counterfactuals. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 915–924, 2022.
- Datta, A., Sen, S., and Zick, Y. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE symposium on security and privacy (SP)*, pp. 598–617. IEEE, 2016.
- De Graaf, M. M. and Malle, B. F. How people explain action (and autonomous intelligent systems should too). In *2017 AAAI Fall Symposium Series*, 2017.
- de Mijolla, D., Frye, C., Kunesch, M., Mansir, J., and Feige, I. Human-interpretable model explainability on high-dimensional data. *arXiv preprint arXiv:2010.07384*, 2020.
- Galhotra, S., Pradhan, R., and Salimi, B. Explaining black-box algorithms using probabilistic contrastive counterfactuals. In *Proceedings of the 2021 International Conference on Management of Data*, pp. 577–590, 2021.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8110–8119, 2020.
- Kommiya Mothilal, R., Mahajan, D., Tan, C., and Sharma, A. Towards unifying feature attribution and counterfactual explanations: Different means to the same end. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 652–663, 2021.
- Kumar, I. E., Venkatasubramanian, S., Scheidegger, C., and Friedler, S. Problems with shapley-value-based explanations as feature importance measures. In *International Conference on Machine Learning*, pp. 5491–5500. PMLR, 2020.
- Li, X., Zhou, Y., Dvornek, N. C., Gu, Y., Ventola, P., and Duncan, J. S. Efficient shapley explanation for features importance estimation under uncertainty. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 792–801. Springer, 2020.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Lundberg, S. M. and Lee, S.-I. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.
- Miller, T. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267:1–38, 2019.
- Nemirovsky, D., Thiebaut, N., Xu, Y., and Gupta, A. CounterGAN: Generating realistic counterfactuals with residual generative adversarial nets. *arXiv preprint arXiv:2009.05199*, 2020.
- Ribeiro, M. T., Singh, S., and Guestrin, C. ” why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.
- Samangouei, P., Saeedi, A., Nakagawa, L., and Silberman, N. ExplainGAN: Model explanation via decision boundary crossing transformations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 666–681, 2018.
- Sundararajan, M. and Najmi, A. The many shapley values for model explanation. In *International conference on machine learning*, pp. 9269–9278. PMLR, 2020.
- Van den Broeck, G., Lykov, A., Schleich, M., and Suciu, D. On the tractability of shap explanations. In *Proceedings of the 35th Conference on Artificial Intelligence (AAAI)*, 2021.
- Winter, E. The shapley value. *Handbook of game theory with economic applications*, 3:2025–2054, 2002.