
DiffusionPID: Interpreting Diffusion via Partial Information Decomposition

Shaurya Dewan*

Rushikesh Zawar*

Prakanshul Saxena*

Yingshan Chang

Andrew Luo

Yonatan Bisk

Carnegie Mellon University
(srdewan, rzawar, prakanss, yingshac, afluo, ybisk}@andrew.cmu.edu
(* denotes equal contribution)

Abstract

Text-to-image diffusion models have made significant progress in generating naturalistic images from textual inputs, and demonstrate the capacity to learn and represent complex visual-semantic relationships. While these diffusion models have achieved remarkable success, the underlying mechanisms driving their performance are not yet fully accounted for, with many unanswered questions surrounding what they learn, how they represent visual-semantic relationships, and why they sometimes fail to generalize. Our work presents Diffusion **P**artial **I**nformation **D**ecomposition (DiffusionPID), a novel technique that applies information-theoretic principles to decompose the input text prompt into its elementary components, enabling a detailed examination of how individual tokens and their interactions shape the generated image. We introduce a formal approach to analyze the uniqueness, redundancy, and synergy terms by applying PID to the denoising model at both the image and pixel level. This approach enables us to characterize how individual tokens and their interactions affect the model output. We first present a fine-grained analysis of characteristics utilized by the model to uniquely localize specific concepts, we then apply our approach in bias analysis and show it can recover gender and ethnicity biases. Finally, we use our method to visually characterize word ambiguity and similarity from the model’s perspective and illustrate the efficacy of our method for prompt intervention. Our results show that PID is a potent tool for evaluating and diagnosing text-to-image diffusion models. Link to project page: <https://rbz-99.github.io/Diffusion-PID/>.

1 Introduction




Figure 1: **Concept Figure. Left:** Our "baseball" uniqueness map specifically highlights the seam region of the tennis ball as it is visually very similar to that of a baseball, **Center:** We see a high synergy for "bat" with "baseball" and "overhead" respectively which shows that it uses these contextual cues to generate the images in the right settings, **Right:** Our redundancy map between "queen" and "crown" correctly focuses on the crown and facial region.

The role of individual inputs and the pretraining data in diffusion-based image generation remains opaque, with little insight into how they translate often ambiguous text-based prompts into images. If we ask you to imagine a lush forest at the edge of a sandy beach, you likely fill in missing details, picturing golden sand, splashing waves, and sunlight filtering through dense green foliage. Our prior experience allows us to easily visualize complex semantic associations and imagine novel scenes.

In the previous example, you are able to both generate and justify your choices for how to complete the scene. Recent advances in diffusion-based generative models have achieved remarkable success in high-quality text-conditioned image generation, mimicking this human ability to a certain extent. However, the lack of transparency in these models raises important questions about their interpretability and controllability, making it challenging to identify biases, errors, or inconsistencies in the generated images. Text-to-image diffusion models can be brittle and prone to failure when faced with out-of-distribution inputs, ambiguous text descriptions, or nuanced contextual dependencies. This results in unrealistic, inconsistent, or even misleading generations, leading to unnatural (and linguistically non-sensical) prompt engineering [1, 2] to achieve desired results by identifying and exploiting spurious correlations captured by the model [3]. The lack of interpretability also hinders the ability to correct or refine the models when they produce undesirable outputs, limiting the potential applications of these generative models in real-world scenarios. Frequently, in multi-concept sentence-based conditioning of generative image models, concepts that are redundant to humans may be critical to the model. For instance, in the phrase "he swung a baseball bat", the word "baseball" may be redundant to humans, but it could be a crucial factor in the model's decision to generate a baseball bat instead of the animal as can be seen in Fig 1. This work posits that a crucial step towards more interpretable and robust models is to develop tools that can help explain higher-order contributions and interactions of individual words to the final image output. This would enable users to understand how the model is using the input text to generate the image, identify potential biases or errors, and refine the model's performance.

To that end, we introduce DiffusionPID, a novel approach that leverages partial information decomposition (**PID**), from information theory [4], to interpret the diffusion process and uncover the contributions of individual words and their interactions to the generated image. Compared to prior work which typically examined the contribution of tokens individually, our approach builds upon mutual information (**MI**) to examine the redundancy (**R**), synergy (**S**), and uniqueness (**U**) of the conditioning tokens. Specifically, our approach enables us to capture the interactions between concepts, revealing how individual words and phrases combine to influence the generation of specific image features. We further introduce an extension of PID conditioned on the rest of the prompt as context, Conditional PID (CPID). By quantifying these interactions, DiffusionPID provides a more comprehensive understanding of the diffusion process, enabling the development of more interpretable and controllable text-to-image models.

This work contributes: (1) a method to compute the image and pixel-level PID of diffusion-generated images with respect to concepts from the input prompt; (2) an extension of PID, namely CPID, to better account for the context provided by the rest of the prompt during the generative process; (3) an evaluation of the DiffusionPID framework for elucidating model dynamics (via **R**, **S**, **U**) on complex linguistic concepts such as synonyms and homonyms and exposing ethnic/gender bias in the model; and finally, (4) a method of prompt intervention via PID to remove redundant words and retain unique and synergistic words.

2 Related Work

Text-to-Image (T2I) Diffusion Models. Diffusion-based text-to-image models have shown impressive results in image synthesis [5, 6, 7, 8]. These models are capable of generating images that are both realistic and semantically consistent with the input text. This has led to the development of many large-scale and diverse datasets [9, 10, 11]. Recent datasets of challenging prompts and benchmarks based on visual question answering [12] and scene graphs have enabled researchers to assess and refine their models [13, 14, 15, 16]. Inspired by the observation that existing models struggle to accurately bind attributes to objects [17], some approaches propose to modify the cross-attention layers by explicitly amplifying the token-wise attention used for text conditioning [18, 19, 20, 21, 22], or by composing multiple outputs to generate semantically compositional images [23]. These methods aim to improve the model's ability to capture fine-grained semantic relationships between text and image. These diverse approaches demonstrate the ongoing effort to develop more effective and flexible text-to-image synthesis models and highlight the need for continued research in this area.

Interpretability in Deep Learning: Interpretability is key to understanding and trusting the decisions made by complex models. LIME [24] presents a technique for interpreting individual predictions of any machine learning (ML) model via local linear models, while SHAP [25] assigns each feature an importance value based on Shapley values from cooperative game theory. CAM [26] and Grad-CAM [27] localize image regions in a CNN using a weighted sum of intermediate layers’ activations. Network Dissection [28] operates at the neuron level, tracking image regions corresponding to an excited neuron across all the dataset images, which GAN Dissection [29] extends for GANs. MILAN [30] and CLIP-Dissect [31] identify semantic concepts for an individual neuron. Benchmarks such as ARO [32], Winoground [33], EQBEN [34], and WHOOPS [35] introduce datasets to specifically probe the compositional and semantic understanding of and show that most SOTA models struggle with relational concepts. Some works have also explored the application of PID in different settings [36]. [37] uses PID to measure the contribution of each neuron towards a target concept and [38] applies PID on generic multimodal datasets and models. Finally, MULTIVIZ [39] analyzes the behavior of multimodal models, specifically focusing on the unimodal, cross-modal, and multi-modal contributions. Overall, a lot of progress has been made in understanding and interpreting ML models. Given the recent explosion of work on diffusion models, we primarily focus on interpreting these models.

Interpretability of Diffusion Models: Interpretability of diffusion models remains an open challenge. Data attribution [40, 41, 42, 43] is a commonly used approach to identify training samples most responsible for the appearance of a generated image. Similarly, [44, 45] decompose the images associated with a particular concept into the implicit semantic concepts the generative model considers to be most similar. Others [46, 47, 48, 49, 50] have discovered a semantic latent space for diffusion models such that modifications in this space edit attributes of the original generated images. [51, 52] are similar but operate in the text embedding space. Recently, there have also been works that try to gain a neuron-level and network-level understanding of these models. For instance, [53] divides the latent vector into multiple groups of elements and designs different noise schedules for each group so that each group controls only certain elements of data, explicitly giving interpretable meaning. [54] identifies clusters of neurons corresponding to a given concept with gradient statistics. [52, 55] show that diffusion models struggle on prompts containing homonyms (words with multiple meanings) and [56, 57] suggested the presence of ethnic, gender, and semantic biases in diffusion models – we investigate both of these conditions in sections 4.3 and [4.1, 4.2], respectively. Recently, DAAM [58] proposed to analyze the cross-attention layers of diffusion models to produce attribution maps for tokens from the input prompt. However, these maps are generally very noisy and offer limited information if the model fails to generate a given concept. [59] introduced a method to compute the mutual information (MI) and conditional mutual information (CMI) between the input text prompt and each generated pixel. While their work provides stronger and more meaningful insights, their approach only provides coarse information on the model. Our work enables fine-grained model understanding by decomposing the mutual information between the text prompt and image output by using partial information decomposition (PID).

3 Method

Image diffusion models learn to model an image distribution, X , through a progressive noise addition and removal process. The forward diffusion process, defined by $x_\alpha = \sqrt{\sigma(\alpha)}x + \sqrt{\sigma(-\alpha)}\epsilon$, gradually adds Gaussian noise, $\epsilon \sim N(0, I)$, to the image, $x \sim p(X)$, over a sequence of timesteps, where α denotes the log SNR/noise schedule and σ is the sigmoid function. A model learns to reverse this process, denoising the data to gradually transform the noise back into samples from the original data distribution. This denoising process has been applied to text-to-image generation with incredible success by conditioning the process on text prompts y . We employ the denoising process to measure the influence of individual tokens (from the input prompt) and their interactions on the resulting generation. We consider the diffusion model to be an optimal denoiser that can predict the added ground-truth noise, ϵ , at noise level α as:

$$\hat{\epsilon}_\alpha(x) = \arg \min_{\hat{\epsilon}(\cdot)} E_{p(x), p(\epsilon)} [\|\epsilon - \hat{\epsilon}_\alpha(x_\alpha)\|^2] \quad (1)$$

We follow [59]’s definition of mutual information (MI) and conditional mutual information (CMI) computed using a diffusion model. Extending Eq 1 to the conditional case where we assume $\hat{\epsilon}_\alpha(x_\alpha|y)$ is the optimal denoiser for the conditional distribution $p(x|y)$, the Log Likelihood Ratio (LLR), which is also the mutual information (MI) $i(y; x)$, between a pixel x and an input text prompt y is defined as:

$$\log p(x | y) - \log p(x) = \frac{1}{2} \int E_{p(\epsilon)} [|\epsilon - \hat{\epsilon}_\alpha(x_\alpha)|^2 - |\epsilon - \hat{\epsilon}_\alpha(x_\alpha | y)|^2] d\alpha \quad (2)$$

$$\log p(x | y) - \log p(x) = \frac{1}{2} \int E_{p(\epsilon)} [|\hat{\epsilon}_\alpha(x_\alpha) - \hat{\epsilon}_\alpha(x_\alpha | y)|^2] d\alpha \quad (3)$$

Eq 3 is derived from Eq 2 using the orthogonality principle and the derivation can be found in the supplemental. We also provide graphs comparing the MMSE estimates obtained from Eq 2 versus the simplified form in Eq 3 for varying levels of noise/SNR and discuss the uncertainty in our estimator in the supplemental.

Essentially, MI quantifies the information an input variable y individually provides about an output variable x . This definition of pixel-wise MI derived from [60] can be intuitive and easily understood as the quantity by which the text prompt y increases the probability of observing x relative to the prior probability without text conditioning. The same work provides an alternate definition for MI as:

$$i(y; x) = -\log p(y) + \log p(y|x) = \log p(x|y) - \log p(x) \quad (4)$$

Similarly for the CMI of phrase y_1 given the context of the phrase y_2 in the input text prompt y :

$$i(y_1; x|y_2) = -\log p(y_1|y_2) + \log p(y_1|x, y_2) = \log p(x|y_1, y_2) - \log p(x|y_2) \quad (5)$$

From Eq 5, CMI can be understood as the information that y_1 contains about x beyond what is already contained in y_2 . Following [60], we provide the mathematical definitions for the various terms in the partial information decomposition (PID) at the pixel-level. We start with the MI of two input events (phrases) with an output variable (pixel) defined as:

$$i(y_1, y_2; x) = r(y_1, y_2; x) + u(y_1 \setminus y_2; x) + u(y_2 \setminus y_1; x) + s(y_1, y_2; x) \quad (6)$$

Here, $y_1 \setminus y_2$ means y_1 excluding y_2 . $r(y_1, y_2; x)$ is the redundant or overlapping information between y_1 and y_2 (redundancy), $u(y_1 \setminus y_2; x)$ is the unique information contributed by y_1 , $u(y_2 \setminus y_1; x)$ is the unique information contributed by y_2 , and $s(y_1, y_2; x)$ is the new information contributed by the combination of y_1 and y_2 (synergy) that neither of them could contribute on their own. It is important to note that the uniqueness defined here is with respect to the other input variable, i.e., it is the unique information the given phrase contributes that the other phrase does not.

A natural measure of redundancy is the expected value of the minimum information that any source provides about each outcome of x . This captures the idea that redundancy is the information common to all sources (the minimum information that any source provides) while taking into account that sources may provide information about different outcomes of x . However, the redundant information must capture when two predictor variables are carrying the same information about the target, not merely the same amount of information. This means that the sign/direction of information matters. To account for this and to ensure the definition complies with all the required axioms on redundancy, the redundancy is broken down into a positive $r^+(y_1, y_2; x)$ and negative $r^-(y_1, y_2; x)$ component associated with the informative $i^+(y_i; x)$ and misinformative $i^-(y_i; x)$ MI terms respectively. Thus, the equation for computing the redundancy can be derived as:

$$r^+(y_1, y_2; x) = \min_{y_i} i^+(y_i; x) = \min_{y_i} [-\log p(y_i)], \quad y_i \in \{y_1, y_2\} \quad (7)$$

$$r^-(y_1, y_2; x) = \min_{y_i} i^-(y_i; x) = \min_{y_i} [-\log p(y_i|x)] \quad (8)$$

$$r(y_1, y_2; x) = r^+ - r^- = \min_{y_i} [-\log p(y_i)] - \min_{y_i} [-\log p(y_i|x)] \quad (9)$$

$$r(y_1, y_2; x) = \min_{y_i} [-\log p(y_i)] - \min_{y_i} [-\log p(x|y_i) + \log p(x) - \log p(y_i)] \quad (10)$$

To compute the probability of an input event (phrase) $p(y)$ in the above equations, we make use of BERT [61]. One of the tasks BERT was trained on was to predict a masked token by predicting the probabilities of all possible tokens from a fixed dictionary and taking the maximum. We use this to obtain the probability of a given phrase. In the unconditional case, we only require the objective probability of the phrase with no other tokens except the special [MASK] token present in the string fed to BERT.

Given the definitions for MI and redundancy, the uniqueness of an input event/phrase and the synergy between phrases can be derived as:

$$u(y_1 \setminus y_2; x) = i(y_1; x) - r(y_1, y_2; x) \quad (11)$$

$$s(y_1, y_2; x) = i(y_1, y_2; x) - r(y_1, y_2; x) - u(y_1 \setminus y_2; x) - u(y_2 \setminus y_1; x) \quad (12)$$

We also introduce CPID (Conditional PID) as an extension of PID to take the context contributed by the rest of the prompt into consideration. It represents the PID for the conditional case where all the PID components and probability terms are now conditioned on the rest of the prompt. This is similar to the CMI extension of MI in [59]. We rewrite equations 6, 10, 11 and 12 with the required changes below for easy reference (y signifies the rest of the prompt with the terms y_1 and y_2 removed):

$$i(y_1, y_2; x|y) = r(y_1, y_2; x|y) + u(y_1 \setminus y_2; x|y) + u(y_2 \setminus y_1; x|y) + s(y_1, y_2; x|y) \quad (13)$$

$$r(y_1, y_2; x|y) = \min_{y_i} [-\log p(y_i|y)] - \min_{y_i} [-\log p(x|y_i, y) + \log p(x|y) - \log p(y_i|y)] \quad (14)$$

$$u(y_1 \setminus y_2; x|y) = i(y_1; x|y) - r(y_1, y_2; x|y) \quad (15)$$

$$s(y_1, y_2; x|y) = i(y_1, y_2; x|y) - r(y_1, y_2; x|y) - u(y_1 \setminus y_2; x|y) - u(y_2 \setminus y_1; x|y) \quad (16)$$

It is important to note that all our definitions for all the information terms require the evaluation of an integral over an infinite range of SNRs. In practice, we make use of truncated logistic-based importance sampling to evaluate this integral, similar to [59].

4 Experiments

We use our method to conduct a detailed analysis of diffusion models in various situations. We introduce several tasks along with corresponding datasets to study these models’ behavior in depth. For our experiments, we primarily focus on the pre-trained Stable Diffusion 2.1 model from Hugging Face. Latent diffusion models encode images into a lower dimensional latent space before the denoising process. Thus, all our PID computations occur in this latent space, i.e., we consider the image in this lower dimensional space as x . During visualization, the heatmaps are bilinearly interpolated from this latent space to the original image resolution. Finally, we make use of 50 samples for evaluating the integral over SNRs using importance sampling. A single A6000 GPU was used to generate the PID maps for each data sample.

4.1 Gender Bias

Setup: Gender is a social construct and a complex study in its own right. In this work, we search for the perpetuation of traditional gender roles and associations. While there is a rich literature on the implications of such biases and how models can exacerbate them [62, 63, 64, 65, 66, 67, 68, 69, 70], we simply present PID as another tool in that investigative process. In this study, we evaluate whether the diffusion model exhibits gender bias in generating images of people in various occupations. Our objective is to determine if the model consistently associates specific genders with certain occupations. For this task, we take a set of 188 common occupations such as professor, valet, receptionist, janitor, etc., and create 376 prompts (188 x 2) by joining the occupation with each gender, male or female, one at a time. More details on the dataset creation process can be found in the supplemental. To check for bias, we compare the image-level redundancy values of each gender with the occupations. The image-level value is simply computed by using the expectation of the MSE term over all pixels instead of the pixel-level MSE values in our equations. The final redundancy values are normalized across the entire dataset to the range [0, 1]. We hypothesize that whenever the redundancy of one gender is much higher than that of the other with the occupation, the model is biased towards the higher redundancy gender for that occupation.

Results: We can see in Table 1 that vocations that are typically stereotyped to be performed by males such as plumber, carpenter, and police officer, have high redundancy values for the male gender as compared to the female gender. Similarly, we see higher redundancy for the female gender for female stereotyped jobs such as babysitter and teacher. We also observe that the average redundancy for females across all occupations is very low, indicating significant model bias against generating females for any occupation. Thus, it is clear that the model has learnt gender biases.

Occupation	Male	Female
Surgeon	0.539	0.055
Soldier	0.250	0.136
Judge	0.304	0.286
Doctor	0.871	0.090
Plumber	0.605	0.038
Carpenter	0.365	0.093
Police Officer	0.390	0.091
Babysitter	0.240	0.531
Teacher	0.098	0.419
Average	0.286	0.194

Table 1: **Gender Bias:** Redundancy between gender and occupation

Occupation	Black	Asian	Caucasian	Hispanic
Athlete	0.321	0.132	0.167	0.156
Artist	0.106	0.069	0.062	0.045
Engineer	0.126	0.156	0.080	0.097
Physicist	0.109	0.209	0.162	0.064
Butcher	0.110	0.179	0.474	0.396
Coach	0.118	0.107	0.433	0.128
Nurse	0.106	0.106	0.127	0.219
Agriculturist	0.046	0.117	0.337	0.450
Average	0.133	0.233	0.255	0.236

Table 2: **Ethnicity Bias:** Redundancy between ethnicity and occupation

4.2 Ethnic Bias

Setup: Similar to gender, race and ethnicity are not hard and fast classes, but indicate how individuals self-identify. Previous work [71, 69, 72, 73, 74, 75] has shown that deep learning models are prone to learning racial and ethnic biases, especially when the underlying training data is imbalanced. With the recent shift towards diffusion-generated synthetic datasets [9, 10, 11], it has become crucial to test for the presence of such biases in diffusion. Similar to the experiment above, we examine the diffusion model for ethnic bias to see if it associates specific ethnicities with specific occupations. We use the same 188 occupations used in the gender bias experiment but now we combine them with each ethnicity (Caucasian, Black, Asian, and Hispanic) instead of gender. We end up creating a dataset of 752 (188 x 4) image-prompt pairs. Refer to the supplemental for more details on the dataset creation process. The same approach is adopted as the previous experiment where we compare the normalized image-level redundancy values of the various ethnicities with the occupations.

Results: We can see from Table 2 that jobs that are commonly stereotyped for certain ethnicities such as athlete with Black, engineer with Asian, and so on, have higher redundancies with those respective ethnicities. We also compare the redundancies of each ethnicity averaged over all occupations and observe a very low value for the Black ethnic group which means that the model is heavily biased against this group and is less likely to generate people from this group for any occupation on average. Thus, we clearly see that the model has learnt biases on the basis of ethnicity.

4.3 Homonyms

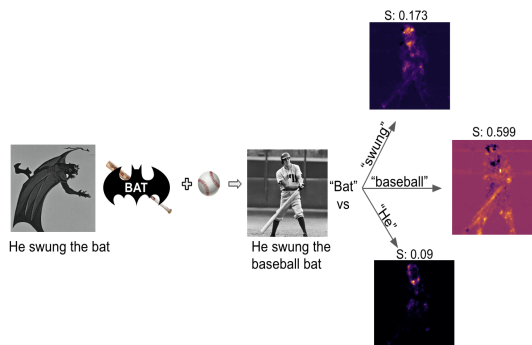


Figure 2: **Homonyms.** We see that the word "baseball" provides the required synergistic context with the homonym "bat" to pick the sports setting over the animal. This effect can be confirmed in the synergy maps and image-level synergy values (S) as well where we observe a high synergy for "bat" with "baseball" compared to other words such as "He" and "swung".

A diffusion model can behave in different manners when faced with prompts containing homonyms in different contexts. A word is a homonym when it can take on different meanings when used with

different contextual words which we term as *modifiers*. For example, the phrases "football **match**" and "he lit a **match**" should generate two completely different visuals. We examine how diffusion models handle such prompts in this experiment.

Setup: For this task, we use our approach to analyze the novel information contributed towards the image generation process by different *modifiers*. We create a dataset of 242 prompts containing homonyms in different contexts. We study the synergy map between the homonyms and the *modifiers* to see if the model is able to extract new information from their combination to generate an image of the homonym in the right context.

Results: In the left figure of Fig 3, we can see our synergy maps correctly highlight the area pertaining to the new information contributed by the combination of the homonym ("bowl") and modifier ("ceramic" and "game" respectively). However, in the right figure of Fig 3, the model fails to generate the homonym ("mole") in the right context even when given the appropriate modifiers ("coworker" and "searching"). This failure can be understood by the low synergy maps which reveal that the model fails to learn the synergistic relationship between the homonym and the modifiers. Thus, the synergy maps help us decode these kinds of failures of the model.

We also provide synergy map visuals obtained from our CPID implementation in the supplemental.

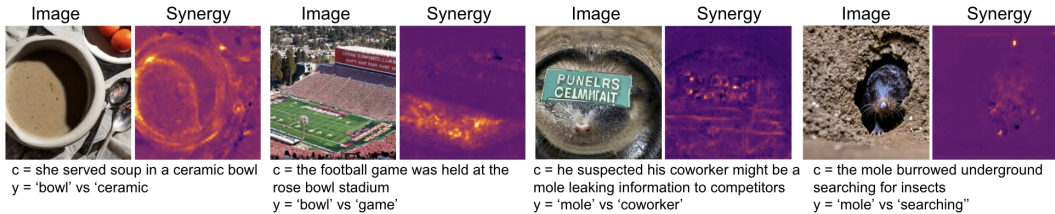


Figure 3: **Homonyms. Left:** Successful generation of homonym "bowl" in different contexts due to high synergy with modifiers "bowl" and "game". **Right:** Failure case where the model generates the homonym "mole" with the same semantic meaning, the animal, due to its failure to use contextual information from words like "coworker" as can be seen in the synergy map.

4.4 Synonyms

Setup: Dealing with synonyms, i.e., different words with the same meaning, can be challenging for diffusion models. Here, we probe the diffusion model with prompts containing known synonym pairs to see how it handles them and if it correctly identifies the given words as synonyms. To that end, we generated a dataset of images from 132 text prompts containing synonym pairs. More details on the dataset creation process can be found in the supplementary. Our hypothesis is that the redundancy will be high between words that the model considers as synonyms.

Results: In Fig 4, we can see how our approach very clearly outputs high redundancy in the region of the object that both synonyms, "cube" and "cuboid", refer to. Similarly, the redundancy map for the words "bed" and "mattress" correctly highlights the bed region. To have a fair comparison with MI, CMI, and DAAM, we take the intersection between the maps of the individual synonyms produced by each of these methods (details in supplementary). It is visible from Fig 4 that these methods fail to find the region of interest and produce very sparse maps. Hence, our method is better suited for interpreting what words/phrases the diffusion model considers semantically similar.




Figure 4: **Synonyms.** Our redundancy map is able to highlight that the model considers the pairs "bed" and "mattress" (**left**) and "cube" and "cuboid" (**right**) as semantically similar.

4.5 Co-Hyponyms

Previous work has shown that diffusion models sometimes fail to generate certain objects mentioned in the prompt. One such set of cases the model struggles with is prompts containing co-hyponyms,

i.e., concepts that are semantically very similar but not identical. The model tends to generate only one of the objects rather than both of them or a fused version of the objects. We hypothesize that this phenomenon occurs because the model confuses the two concepts to carry the same semantic meaning even if they don't for us as human beings. This means that the redundancy between them is expected to be very high which is also what we can see in our experiments. We introduce 2 prompt datasets containing co-hyponyms and are discussed below.

4.5.1 Co-Hyponym COCO

Setup: To construct this dataset, we used the COCO dataset's [76] super-category hierarchy to extract co-hyponym pairs. The supplemental contains more details on the dataset creation process.

Results: In Fig 5, the redundancy is visibly high in the region where the diffusion model has fused features from the words "cat" and "elephant". Similarly, the redundancy is highly activated in the region of the only generated object due to the high semantic similarity of the words "pizza" and "sandwich". Thus, it is apparent that for cases where the diffusion model fails to learn the difference in the semantic meaning of two co-hyponyms, it confuses the two to be referring to the same concept, signified by the high redundancy, and fails to generate one of the objects. Similar to the synonym experiment above, we compare our redundancy map with the intersection maps from MI, CMI, and DAAM and observe that they provide little information for interpreting the reason for the model's failure. Thus, the redundancy map proves to be a useful tool to understand why the model fails here.




Figure 5: **COCO Co-Hyponyms.** The redundancy map proves to be very useful in finding out the reason behind the model's failures in these figures. It confuses the co-hyponym pairs ("sandwich", "pizza")(left) and ("elephant", "cat")(right) to have the same meaning for the co-hyponyms as seen from the redundancy maps, which results in erroneous generations.

4.5.2 Co-Hyponym WORDNET

Setup: We make use of Wordnet's [77] hyponym and hypernym relations between the synsets of the words to obtain 798 co-hyponym pairs for the prompts. Refer to the supplementary for more details on the dataset creation process.

Results: In the left figure of Fig 6, we can see the fusion of features from both objects mentioned in the prompt, while in the right figure, we see only one of the objects has been generated. The redundancy maps correctly highlight the image regions pertaining to both the co-hyponyms. Again it is observed that MI, CMI, and DAAM only very sparsely highlight the correct region of overlap. This reinforces our findings from the previous COCO-based experiment that the model mixes up the co-hyponyms and our redundancy map from PID helps pin down the reason behind this phenomenon.




Figure 6: **Wordnet Co-Hyponyms.** Similar to the COCO Co-Hyponyms dataset, the redundancy maps tell us that the model confuses the co-hyponym pairs ("barrier", "railing")(left) and ("chair", "sofa")(right) to have the same meaning for the co-hyponyms, resulting in erroneous generations.

4.6 Prompt Intervention

We use our method to identify redundant words in the prompts, remove them, and verify that this intervention results in little change to the image. We test this task on all the datasets involving

redundant terms: co-hyponym wordnet, co-hyponym coco, synonyms, and occupation bias (ethnic and gender). For more details on how we edit the image, refer to the supplementary.




Figure 7: **Prompt Intervention.** The redundancy is highly activated in the face region of the giraffe. On a closer look, we see that the face is that of a cow meaning that the word "cow" is redundant. This is confirmed as the image changes very little on omitting it from the prompt.




Figure 8: **Prompt Intervention.** The redundancy is highly activated in the mouse region meaning that the word "mouse" is redundant. The uniqueness of "mouse" is also very low and spread out. This redundancy is confirmed as the image changes very little on omitting it from the prompt.

Both figures 7 and 8 clearly show that whenever the redundancy is high in a region corresponding to a particular object and the uniqueness for that object's term is low/spread out, then removing that object has minimal effect on the image. On the other hand, we see that the MI, CMI, and DAAM maps are highly activated for these redundant terms, indicating that these metrics are not ideal for identifying redundant terms in the input prompt. For instance, in the first figure, the redundancy map is highly activated in the mouse region and the uniqueness of the term "mouse" is spread out. Removing the word "mouse" results in an image very similar to the original whereas removing the word "laptop" drastically changes the image. The MI, CMI, and DAAM maps are notably high for the "mouse" term, incorrectly suggesting that removing this term should have significantly altered the image. A similar effect can be observed for the term "cow" in the second figure.

4.7 Most Representative Features

The most representative features of a concept are those characteristics that uniquely define it from the model's perspective. Given a set of classes, the uniqueness map obtained from PID can be used to localize these features. As we can see in Fig 9, when the input prompt is "hair dryer and toothbrush", even though the image is a mix of a hair dryer and a toothbrush, the uniqueness map highlights the region corresponding to the toothbrush bristles. Similarly, we can see the bear's facial features being highlighted correctly in the other example. Thus, our uniqueness maps prove useful for this task of identifying the unique characteristics of an object.




Figure 9: **Most Representative Features.** **Left:** The "toothbrush" uniqueness map correctly captures the toothbrush bristles, their most distinct feature. **Right:** The "bear" uniqueness map correctly captures the bear region, specifically the face.



Figure 10: Results of PID and CPID on complex prompts

4.8 Complex Prompts

In this experiment, we test the efficacy of our method on more complex prompts, similar to those in the diffusion training distribution. These prompts usually mention several objects and their corresponding attributes. We find that our method remains effective and informative even in these challenging examples. We visualize the information maps between objects and attribute-defining terms.

In both visuals in Fig 10, we observe a high synergy because the attribute modifies the object’s visual properties in some form. We also see that CPID provides slightly better localized results than its PID counterpart in practice. This is expected as CPID accounts for the contextual contribution of the rest of the prompt in the image generation process and better captures the specific contribution of the two terms under consideration.

5 Discussion

Limitations and Future Work: Although we have exhibited the efficacy and benefits of DiffusionPID, there remains some scope for improvement that future work can explore. In all our experiments, we compute the PID terms only for two phrases from the prompt but PID can be extended to more than two input variables. Another interesting research direction could be to compute PID on other, non-diffusion-based models. Also, our approach requires access to the diffusion model, making it difficult to apply it to closed-source models, which means that alternative methods of computing PID need to be explored. There are also some other applications that could be tried with PID such as using the uniqueness information to localize distinctive features to differentiate closely related classes in image classification. Recently there have been works such as [78] that integrate other forms of conditioning such as masks to enable better control of diffusion models. PID-based analysis of these multimodal forms of conditioning is another worthwhile research direction for the future.

Conclusion: The fine-grained breakdown of MI afforded by our approach, DiffusionPID, allows us to understand the decision-making process of diffusion models, identify their shortcomings, and pin down the reasons for their failures. This understanding is critical given the misalignment between the model’s world understanding and ours. We can use the insights furnished from our approach to make diffusion models better aligned with humans’ conceptual understanding and counter its limitations and biases. Our work can serve as a strong fundamental basis for further research on using information theoretic concepts such as PID to dissect, study, and enhance deep learning models.

References

- [1] Vivian Liu and Lydia B Chilton. Design guidelines for prompt engineering text-to-image generative models. CHI '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3501825. URL <https://doi.org/10.1145/3491102.3501825>.
- [2] Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Hard prompts made easy: Gradient-based discrete optimization for prompt tuning and discovery. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=V0stHxDdsN>.
- [3] Junhyun Nam, Sangwoo Mo, Jaeho Lee, and Jinwoo Shin. Breaking the spurious causality of conditional generation via fairness intervention with corrective sampling. *arXiv preprint arXiv:2212.02090*, 2022.
- [4] Paul L. Williams and Randall D. Beer. Nonnegative decomposition of multivariate information, 2010.
- [5] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- [6] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [7] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [8] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023.
- [9] Zijie J. Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models, 2023. URL <https://arxiv.org/abs/2210.14896>.
- [10] Keqiang Sun, Junting Pan, Yuying Ge, Hao Li, Haodong Duan, Xiaoshi Wu, Renrui Zhang, Aojun Zhou, Zipeng Qin, Yi Wang, Jifeng Dai, Yu Qiao, Limin Wang, and Hongsheng Li. Journeydb: A benchmark for generative image understanding, 2023. URL <https://arxiv.org/abs/2307.00716>.
- [11] Rushikesh Zavar, Shaurya Dewan, Andrew F. Luo, Margaret M. Henderson, Michael J. Tarr, and Leila Wehbe. Stablesemantics: A synthetic language-vision dataset of semantic representations in naturalistic images, 2024. URL <https://arxiv.org/abs/2406.13735>.
- [12] Aviral Agrawal, Carlos Mateo Samudio Lezcano, Iqui Balam Heredia-Marin, and Prabhdeep Singh Sethi. Listen then see: Video alignment with speaker attention, 2024. URL <https://arxiv.org/abs/2404.13530>.
- [13] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, Seyedeh Sara Mahdavi, Raphael Gontijo Lopes, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. *ArXiv*, abs/2205.11487, 2022.
- [14] Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A. Smith. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20349–20360, 2023.

- [15] Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. Evaluating text-to-visual generation with image-to-text generation. *ArXiv*, abs/2404.01291, 2024.
- [16] Jaemin Cho, Yushi Hu, Jason Baldridge, Roopal Garg, Peter Anderson, Ranjay Krishna, Mohit Bansal, Jordi Pont-Tuset, and Su Wang. Davidsonian scene graph: Improving reliability in fine-grained evaluation for text-to-image generation. In *ICLR*, 2024.
- [17] Colin Conwell and Tomer David Ullman. Testing relational understanding in text-guided image generation. *ArXiv*, abs/2208.00005, 2022.
- [18] Weixi Feng, Xuehai He, Tsu-Jui Fu, Varun Jampani, Arjun Reddy Akula, Pradyumna Narayana, Sugato Basu, Xin Eric Wang, and William Yang Wang. Training-free structured diffusion guidance for compositional text-to-image synthesis. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PUIqjT4rzq7>.
- [19] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42:1 – 10, 2023.
- [20] Royi Rassin, Eran Hirsch, Daniel Glickman, Shauli Ravfogel, Yoav Goldberg, and Gal Chechik. Linguistic binding in diffusion models: Enhancing attribute correspondence through attention map alignment. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=AOKU4nRw1W>.
- [21] Ruichen Wang, Zekang Chen, Chen Chen, Jiancang Ma, H. Lu, and Xiaodong Lin. Compositional text-to-image synthesis with attention map control of diffusion models. *ArXiv*, abs/2305.13921, 2023.
- [22] Quynh Phung, Songwei Ge, and Jia-Bin Huang. Grounded text-to-image synthesis with attention refocusing. *ArXiv*, abs/2306.05427, 2023.
- [23] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 423–439. Springer, 2022.
- [24] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “why should i trust you?”: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [25] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Neural Information Processing Systems*, 2017.
- [26] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.
- [27] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [28] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6541–6549, 2017.
- [29] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B Tenenbaum, William T Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. *arXiv preprint arXiv:1811.10597*, 2018.

- [30] Evan Hernandez, Sarah Schwettmann, David Bau, Teona Bagashvili, Antonio Torralba, and Jacob Andreas. Natural language descriptions of deep visual features. In *International Conference on Learning Representations*, 2021.
- [31] Tuomas P. Oikarinen and Tsui-Wei Weng. Clip-dissect: Automatic description of neuron representations in deep vision networks. *ArXiv*, abs/2204.10965, 2022.
- [32] Mert Yuksekgonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. When and why vision-language models behave like bags-of-words, and what to do about it? In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=KRLUvxh8uaX>.
- [33] Anuj Diwan, Layne Berry, Eunsol Choi, David Harwath, and Kyle Mahowald. Why is winoground hard? investigating failures in visuolinguistic compositionality. *arXiv preprint arXiv:2211.00768*, 2022.
- [34] Tan Wang, Kevin Lin, Linjie Li, Chung-Ching Lin, Zhengyuan Yang, Hanwang Zhang, Zicheng Liu, and Lijuan Wang. Equivariant similarity for vision-language foundation models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11964–11974, 2023.
- [35] Nitzan Bitton-Guetta, Yonatan Bitton, Jack Hessel, Ludwig Schmidt, Yuval Elovici, Gabriel Stanovsky, and Roy Schwartz. Breaking common sense: Whoops! a vision-and-language benchmark of synthetic and compositional images. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2616–2627, 2023.
- [36] Faisal Hamman and Sanghamitra Dutta. Demystifying local and global fairness trade-offs in federated learning using partial information decomposition, 2024. URL <https://arxiv.org/abs/2307.11333>.
- [37] David Alexander Ehrlich, Andreas Christian Schneider, Viola Priesemann, Michael Wibral, and Abdullah Makkeh. A measure of the complexity of neural representations based on partial information decomposition. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=R8TU3pfzFr>.
- [38] Paul Pu Liang, Yun Cheng, Xiang Fan, Chun Kai Ling, Suzanne Nie, Richard J. Chen, Zihao Deng, Nicholas Allen, Randy P. Auerbach, Faisal Mahmood, Ruslan Salakhutdinov, and Louis-Philippe Morency. Quantifying & modeling multimodal interactions: An information decomposition framework. In *Neural Information Processing Systems*, 2023.
- [39] Paul Pu Liang, Yiwei Lyu, Gunjan Chhablani, Nihal Jain, Zihao Deng, Xingbo Wang, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multiviz: Towards visualizing and understanding multimodal models. In *International Conference on Learning Representations*, 2022.
- [40] Xiaosen Zheng, Tianyu Pang, Chao Du, Jing Jiang, and Min Lin. Intriguing properties of data attribution on diffusion models. In *International Conference on Learning Representations (ICLR)*, 2024.
- [41] Kristian Georgiev, Joshua Vendrow, Hadi Salman, Sung Min Park, and Aleksander Madry. The journey, not the destination: How data guides diffusion models. In *Arxiv preprint arXiv:2312.06205*, 2023.
- [42] Zheng Dai and David Kenneth Gifford. Training data attribution for diffusion models. *ArXiv*, abs/2306.02174, 2023.
- [43] Sheng-Yu Wang, Alexei A. Efros, Jun-Yan Zhu, and Richard Zhang. Evaluating data attribution for text-to-image models. In *ICCV*, 2023.
- [44] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *ArXiv*, abs/2208.01618, 2022.
- [45] Hila Chefer, Oran Lang, Mor Geva, Volodymyr Polosukhin, Assaf Shocher, Michal Irani, Inbar Mosseri, and Lior Wolf. The hidden language of diffusion models. *arXiv preprint arXiv:2306.00966*, 2023.

- [46] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. *abs/2210.10960*, 2022.
- [47] Hang Li, Chengzhi Shen, Philip Torr, Volker Tresp, and Jindong Gu. Self-discovering interpretable diffusion latent directions for responsible text-to-image generation. *arXiv preprint arXiv:2311.17216*, 2023.
- [48] René Haas, Inbar Huberman-Spiegelglas, Rotem Mulayoff, and Tomer Michaeli. Discovering interpretable directions in the semantic latent space of diffusion models. *ArXiv*, *abs/2303.11073*, 2023.
- [49] Yong-Hyun Park, Mingi Kwon, Jaewoong Choi, Junghyo Jo, and Youngjung Uh. Understanding the latent space of diffusion models through the lens of riemannian geometry. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 24129–24142. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/4bfcebedf7a2967c410b64670f27f904-Paper-Conference.pdf.
- [50] Yusuf Dalva and Pinar Yanardag. Noiseclr: A contrastive learning approach for unsupervised discovery of interpretable directions in diffusion models, 2023.
- [51] Qiucheng Wu, Yujian Liu, Handong Zhao, Ajinkya Kale, Trung M. Bui, Tong Yu, Zhe Lin, Yang Zhang, and Shiyu Chang. Uncovering the disentanglement capability in text-to-image diffusion models. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1900–1910, 2022.
- [52] Jennifer C. White and Ryan Cotterell. Schrodinger’s bat: Diffusion models sometimes generate polysemous words in superposition, 2022. URL <https://arxiv.org/pdf/2211.13095.pdf>.
- [53] Sangyun Lee, Gayoung Lee, Hyunsu Kim, Junho Kim, and Youngjung Uh. Diffusion models with grouped latents for interpretable latent space. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2023.
- [54] Zhiheng Liu, Ruili Feng, Kai Zhu, Yifei Zhang, Kecheng Zheng, Yu Liu, Deli Zhao, Jingren Zhou, and Yang Cao. Cones: Concept neurons in diffusion models for customized generation. In *International Conference on Machine Learning*, 2023.
- [55] Royi Rassin, Shauli Ravfogel, and Yoav Goldberg. DALLE-2 is seeing double: Flaws in word-to-concept mapping in text2image models. *CoRR*, *abs/2210.10606*, 2022. doi: 10.48550/ARXIV.2210.10606. URL <https://doi.org/10.48550/arXiv.2210.10606>.
- [56] Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. Stable bias: Evaluating societal representations in diffusion models. In *Neural Information Processing Systems*, 2023.
- [57] Hadas Orgad, Bahjat Kawar, and Yonatan Belinkov. Editing implicit assumptions in text-to-image diffusion models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7030–7038, 2023.
- [58] Raphael Tang, Linqing Liu, Akshat Pandey, Zhiying Jiang, Gefei Yang, Karun Kumar, Pontus Stenetorp, Jimmy Lin, and Ferhan Ture. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*, 2022.
- [59] Xianghao Kong, Ollie Liu, Han Li, Dani Yogatama, and Greg Ver Steeg. Interpretable diffusion via information decomposition. *arXiv preprint arXiv:2310.07972*, 2023.
- [60] Conor Finn and Joseph T. Lizier. Pointwise information decomposition using the specificity and ambiguity lattices. *ArXiv*, *abs/1801.09010*, 2017.
- [61] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019.

- [62] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5310–5319, 2019.
- [63] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Conference on Empirical Methods in Natural Language Processing*, 2017.
- [64] Tejas Srinivasan and Yonatan Bisk. Worst of both worlds: Biases compound in pre-trained vision-and-language models. *ArXiv*, abs/2104.08666, 2021.
- [65] Kaylee Burns, Lisa Anne Hendricks, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. *arXiv preprint arXiv:1803.09797*, 2018.
- [66] Mustafa Atay, Megh Poudyel, and Saul Evora. Evaluation of gender bias in masked face recognition with deep learning models. *2023 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 829–835, 2023.
- [67] Ahsan Ul Islam and Abm Rezbaul Islam. Gender bias in pre-trained deep learning models. *2023 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, pages 1–6, 2023.
- [68] Siobhan Mackenzie Hall, Fernanda Gonçalves Abrantes, Hanwen Zhu, Grace Sodunke, Aleksandar Shtedritski, and Hannah Rose Kirk. Visogender: A dataset for benchmarking gender bias in image-text pronoun resolution. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=BNwsJ4bFsc>.
- [69] Yusuke Hirota, Yuta Nakashima, and Noa García. Quantifying societal bias amplification in image captioning. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13440–13449, 2022.
- [70] Shruti Bhargava and David A. Forsyth. Exposing and correcting the gender bias in image captioning datasets and models. *ArXiv*, abs/1912.00578, 2019.
- [71] Pierre Stock and Moustapha Cisse. Convnets and imagenet beyond accuracy: Understanding mistakes and uncovering biases. In *Proceedings of the European conference on computer vision (ECCV)*, pages 498–512, 2018.
- [72] Zaid Khan and Yun Raymond Fu. One label, one billion faces: Usage and consistency of racial categories in computer vision. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021.
- [73] Kimmo Kärkkäinen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1547–1557, 2021.
- [74] Kinshuk Sengupta and Praveen Ranjan Srivastava. Causal effect of racial bias in data and machine learning algorithms on user persuasiveness & discriminatory decision making: An empirical study. *ArXiv*, abs/2202.00471, 2022.
- [75] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. The risk of racial bias in hate speech detection. In *Annual Meeting of the Association for Computational Linguistics*, 2019.
- [76] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [77] George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38:39–41, 1995.

- [78] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.
- [79] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Conference on Computer Vision and Pattern Recognition 2023*, 2023.
- [80] Linoy Tsaban and Apolinário Passos. Ledits: Real image editing with ddpm inversion and semantic guidance. 2023.
- [81] OpenAI. Chatgpt: A large language model trained by openai. <https://chat.openai.com/>, 2023. Accessed: 2024-05-20.
- [82] Meta AI — meta.ai. <https://www.meta.ai/>. [Accessed 22-05-2024].
- [83] Demis Hassabis and Google DeepMind Team. Introducing gemini: Google’s most capable ai model yet, 2024. URL <https://blog.google/technology/ai/introducing-gemini/>. Accessed: 2024-05-20.

6 Appendix

6.1 MI and PID

Here, we explain the concepts of mutual information and partial information decomposition in greater detail.

Mutual Information (MI) is a measure that estimates the amount of information obtained for an output variable from a predictor variable passed through a function. It measures to what extent an input variable reduces the uncertainty in the model's prediction. In deep learning, where understanding the dependencies and mapping between input features and target variables is vital, MI serves as a powerful tool to quantify these dependencies accurately.

Partial Information Decomposition (PID), as defined by [60], shows that two predictor variables X_1 and X_2 can contribute information about a target variable Y in four possible ways. X_1 can provide some information about Y independently/uniquely, X_2 too can provide some information about Y independently/uniquely, both X_1 & X_2 can contribute some overlapping/redundant information, i.e., some of the information contributed by both is the same, and X_1 & X_2 can contribute synergistic information which is information neither of them can contribute on their own and only in conjunction with the other. [60] also state that mutual information can be derived from first principles as fundamentally pointwise quantities where they measure the information content of individual events. The overall (average) mutual information can then be calculated by taking the expectation over all events for the relevant variables.

6.2 Simplification of Mutual Information

Here we outline the derivation of Eq 3 from Eq 2 using the orthogonality principle:

$$I(X; Y) = \mathbb{E}_{p(x,y)} [\log p(x|y) - \log p(x)] \quad (1)$$

$$I(X; Y) = \mathbb{E}_{p(x,y)} \left[\frac{1}{2} \int \mathbb{E}_{p(\epsilon)} [|\epsilon - \hat{\epsilon}_\alpha(x_\alpha)|^2 - |\epsilon - \hat{\epsilon}_\alpha(x_\alpha|y)|^2] d\alpha \right] \quad (2)$$

By expanding all the squares and re-arranging we get:

$$\begin{aligned} I(X; Y) = & \mathbb{E}_{p(x,y)} \left[\frac{1}{2} \int \mathbb{E}_{p(\epsilon)} [|\hat{\epsilon}_\alpha(x_\alpha) - \hat{\epsilon}_\alpha(x_\alpha|y)|^2] d\alpha \right] + \\ & 2\mathbb{E}_{p(y)} \left[\frac{1}{2} \int \mathbb{E}_{p(x|y), p(\epsilon)} [(\hat{\epsilon}_\alpha(x_\alpha) - \hat{\epsilon}_\alpha(x_\alpha|y)) \cdot (\hat{\epsilon}_\alpha(x_\alpha|y) - \epsilon)] d\alpha \right] \end{aligned} \quad (3)$$

Here,

$$+2\mathbb{E}_{p(y)} \left[\frac{1}{2} \int \mathbb{E}_{p(x|y), p(\epsilon)} [(\hat{\epsilon}_\alpha(x_\alpha) - \hat{\epsilon}_\alpha(x_\alpha|y)) \cdot (\hat{\epsilon}_\alpha(x_\alpha|y) - \epsilon)] d\alpha \right] \equiv \ominus \quad (4)$$

Based on the orthogonality principle [1], which states:

$$\forall f, \quad \mathbb{E}_{p(x|y)p(\epsilon)} [f(x_\alpha, y) \cdot (\hat{\epsilon}_\alpha(x_\alpha|y) - \epsilon)] = 0 \quad (5)$$

The term $(\hat{\epsilon}_\alpha(x_\alpha|y) - \epsilon)$ represents the error of the MMSE estimator, which is orthogonal to any estimator. Therefore, the second term becomes zero, leading to:

$$I(X; Y) = \mathbb{E}_{p(x,y)} \left[\frac{1}{2} \int \mathbb{E}_{p(\epsilon)} [|\hat{\epsilon}_\alpha(x_\alpha) - \hat{\epsilon}_\alpha(x_\alpha|y)|^2] d\alpha \right] \quad (6)$$

This is the same as Eq 3.

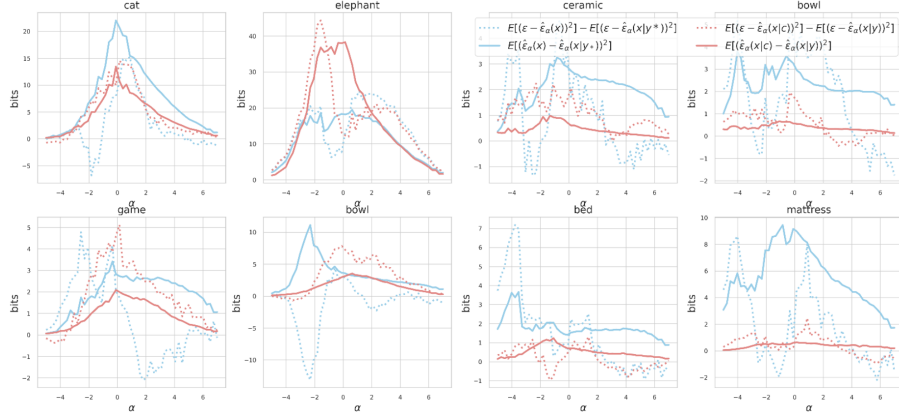


Figure 11: MMSE curves comparing the standard (Eq 2) and orthogonal (Eq 3) estimators

6.3 Standard vs Orthogonal Estimators

It can be seen that the original form (dotted line) is more unstable with many zigzag patterns. We also see the orthogonal/simplified form (continuous line) enforces better consistency between the MMSE (blue) and conditional MMSE (red). Thus, this simplification works better in practice.

6.4 Estimator Uncertainty

There is no guarantee that the estimator provides an upper or lower bound for the PID terms. It is dependent on the conditional and unconditional denoising MMSEs obtained from the diffusion model which is assumed to be an optimal denoiser for our experiments. However, in practice, this assumption need not hold because a neural network trained on a regression problem to minimize MSE need not converge to the global minima and instead may converge to a local minima. That being said, neural networks have been found to do really well on regression problems and the diffusion model, specifically, has been found to perform well on the denoising problem. Thus, we expect reasonable estimates despite the inherent uncertainty.

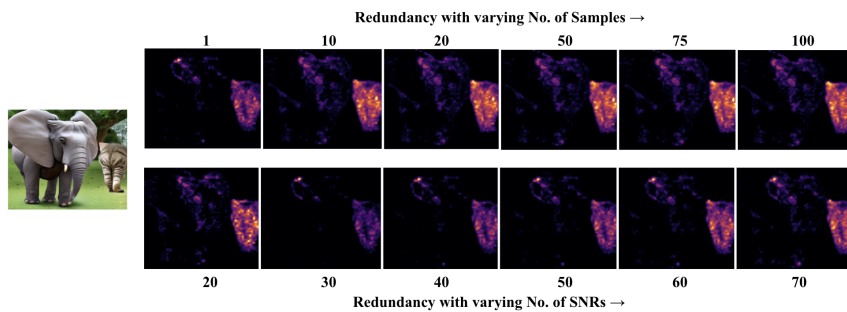


Figure 12: Redundancy maps for the varying levels of noise/SNR and number of samples

A measure of uncertainty is also introduced based on the number of samples under the same noise level, α , in Eq 3's expectation term and from the number of values sampled to evaluate the integral in the equation. Thus, we can obtain more confident information maps by using higher values for both of these hyperparameters. In Fig 12, we provide visuals of the information maps for varying values of these hyperparameters on the "cat and elephant" sample from the COCO co-hyponym experiment (Fig 5 in the main text). We observe that the maps depict the same information, i.e., they are highly activated in the same regions across variations but do become less noisy at higher values.

6.5 Baselines for Redundancy

When comparing CMI, MI, and DAAM with our redundancy maps, we take the intersection of the maps these methods produce for the two words/phrases individually. This is done in the following steps:

- The individual maps are thresholded to obtain binary masks. The threshold we employ is $\mu + 1.5\sigma$, where μ refers to the mean of the map being thresholded and σ refers to the standard deviation. We use 1.5 times the standard deviation to discard roughly 86.64 % of the map’s values as we observe that they are irrelevant and correspond to the low-activated background region.
- We multiply the binary masks of the two words’ maps to obtain the final regions of intersection/overlap between the two maps.
- Finally, the above intersection map im is multiplied with each of the two maps, m_1 and m_2 , followed by the computation of the pixel-level mean of the two masked maps to obtain the final map f as:

$$f = (m_1 * im + m_2 * im) / 2 \tag{7}$$

6.6 Image Editing via Prompt Intervention

We follow a similar process to other image editing methods [79, 80] where we progressively add Gaussian noise to the image and pass it through the diffusion model for denoising but now this reverse diffusion process is conditioned on the edited prompt instead of the original prompt. We edit the original prompt by removing a word from it to see if the edited image is similar to the original to help determine if the omitted word is redundant. For instance, we remove the word "desk" from the prompt "The desk was a sturdy table perfect for working" to produce the edited prompt "The was a sturdy table perfect for working" which is then used to condition the denoising process. The intuition behind this experiment is that for words that contribute little information on top of the rest of the prompt, the model will keep the edited image very similar to the original as it will follow a similar path as that would be followed if it were conditioned on the original prompt.

6.7 Datasets

We create datasets for each of our numerous experiments. We primarily rely on existing hierarchical datasets, Wordnet [77] and COCO [76], and free-access Large Language Models, namely Chat-GPT [81], Meta.ai [82], and Gemini [83]. We probe these models with prompts like "make sentences with pairs of synonyms of common tangible nouns", "make sentences with category and subcategory of common tangible nouns", "generate sentences with synonyms of common tangible nouns and both of those synonyms must be present in the sentence", "generate pairs of sentences containing homonyms in different contexts" and so on. A combination of the outputs produced by these models was used to obtain the final set of prompts for each task. Once we had the prompts, we fed them into the open-source diffusion model, Stable Diffusion 2.1 [7], to generate images corresponding to those prompts. During the image generation, we kept the seed constant (42). Following are the prompt-image pair datasets we introduce:

- **Gender Bias Dataset:** Here we generated a list of 188 most common occupations from the LLMs ([82],[81],[83]). These occupations are then combined with each gender one at a time to produce a total of 376 prompts. For instance, the occupation "Doctor" is used to generate two prompts, "Female Doctor" and "Male Doctor". A sample of this dataset is visible in Fig 13.
- **Ethnic Bias Dataset:** Instead of combining the previously mentioned 188 occupations with a gender, we now combine them with one of four ethnicities (*Caucasian, Asian, Black, Hispanic*) in each prompt to produce a total of 752 prompts. For example, the occupation "Doctor" is used to generate four prompts, "Caucasian Doctor", "Asian Doctor", "Black Doctor", and "Hispanic Doctor". A sample of this dataset is visible in Fig 14.
- **Homonym Dataset:** For this dataset, we generated 121 pairs of prompts, where each pair corresponds to one homonym but in different contexts. Thus, in total, we produced 242

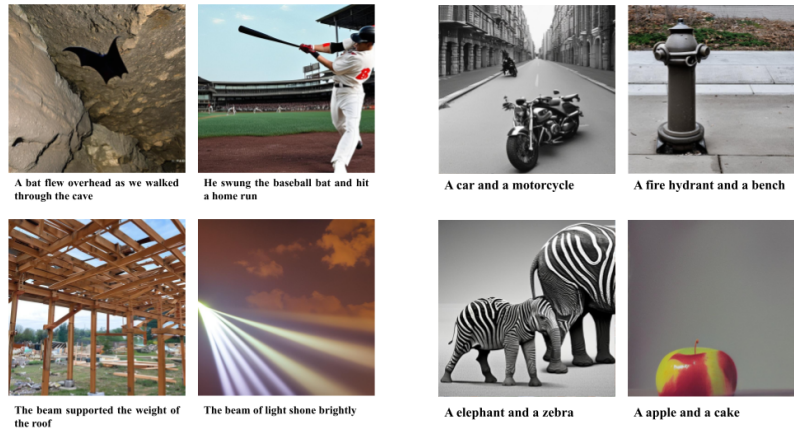


Figure 13: Gender Bias Dataset Sample



Figure 14: Ethnic Bias Dataset Sample

prompts, all of which were generated from the LLMs [82],[81],[83]. A sample of this dataset is visible in figures 15a.



(a) Homonym Dataset Sample

(b) Co-Hyponym Dataset (COCO) Sample

- **Synonym Dataset:** The entire dataset of 132 prompts containing synonym pairs was generated using the previously mentioned LLMs [82],[81],[83]. A sample of this dataset is visible in Fig 16.
- **Co-Hyponym Dataset (Wordnet):** To create this dataset, we first sampled co-hyponyms from Wordnet for the most frequently occurring tangible nouns in the COCO dataset. Next, we generated prompts with the format: "a *co-hyponym* and a *co-hyponym*". For example, one of the most commonly used nouns across the COCO dataset is "bench", so we find its co-hyponyms from Wordnet: "box", "seat", "chair", "ottoman", "stool", "sofa" and "couch". Then, these are used to produce prompts such as "a box and a bench", "a chair and a bench", "a sofa and a bench" and so on. A sample of this dataset is visible in Fig 17.
- **Co-Hyponym Dataset (COCO):** We took the super-categories of the COCO dataset and paired the objects within each individual super-category to produce a list of 197 co-hyponym




Figure 16: Synonym Dataset Sample



Figure 17: Wordnet Dataset Sample


pairs which are then converted to prompts using a similar format as that used for the Wordnet dataset above. For example, "vehicle" was a super-category, so "a car and a airplane", "a car and a boat" and "a motorcycle and a truck" are a few prompts produced from that super-category. A sample of this dataset is visible in Fig 15b.

6.8 More Results

In this section, we present more results from our experiments.

6.8.1 Homonyms

Below are some more examples of results for Homonym Experiments.



c = the calf frolicked in the field alongside its mother
y = 'calf' vs 'field'



c = she stretched her calf muscles before going for a run
 y = 'calf' vs 'run'




c = she admired the stained glass windows in the church
 y = 'glass' vs 'windows'



c = she sipped her water from a tall glass
 y = 'glass' vs 'water'



c = the gang worked together
 y = 'gang' vs 'together'



c = the gang of gears was complex
y = 'gang' vs 'gears'



c = the hook on the door was useful
y = 'hook' vs 'door'



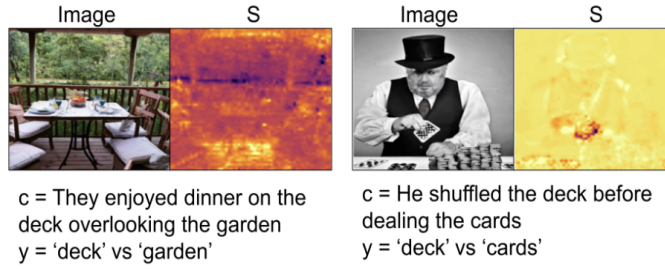
c = he baited the hook with a worm
y = 'hook' vs 'worm'

Below are some Synergy map visuals obtained from CPID on Homonym samples.






c = we saw a crane standing gracefully in the marsh
y = 'crane' vs 'gracefully'


c = the construction crew used a crane to hoist steel beams into place
y = 'crane' vs 'crew'



6.8.2 Synonyms

Below are some more examples of results for Synonyms Experiments.







c = the saw was a sharp tool cutting through wood easily
y = 'saw' vs 'tool'

6.8.3 COCO Co-Hyponyms


Below are the results of Co-Hyponyms.




c = a car and a bus
y = 'car' vs 'bus'



c = a baseball bat and a surfboard
y = 'baseball' vs 'surfboard'




c = a carrot and a pizza
y = 'carrot' vs 'pizza'




c = a cat and a dog
y = 'cat' vs 'dog'

6.8.4 Wordnet Co-Hyponyms


Below are some more results for the Co-Hyponym experiments on the images generated from prompts made through Wordnet.




c = a chamber and a classroom
y = 'chamber' vs 'classroom'



c = a door and a tower
y = 'door' vs 'tower'






c = a glass and a window
y = 'glass' vs 'window'




c = a fox and a dog
y = 'fox' vs 'dog'


6.8.5 Prompt Intervention

Below are some more results for the Prompt Intervention experiments.







c = a skis and a surfboard
y = 'skis'



c = a skis and a surfboard
y = 'surfboard'



c = a clock and a vase
y = 'clock'



c = a clock and a vase
y = 'vase'

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: Sections 3 and 4

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: Section 5 has limitations mentioned

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: Section 3 outlines various theoretical concepts with the required proofs and detailed explanations provided in the appendix in sections 6.1, 6.2, 6.3, and 6.4.

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: Sections 3 and 4 have the required information to reproduce the experiments.

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 shared our code and data in the supplementary. We will be releasing the code publicly before the main conference.

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: Sections 3 and 4 have the required details.

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: [NA]

Justification: Our experiments don't really involve such results that would require error bars or statistical significance tests as they are not distribution-dependent.

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: Yes, the GPU used was an A6000 and it took around 7 seconds to get an image generated with its maps. This information is mentioned in section 4.

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: It is as per the guidelines.

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: [Yes]

Justification: The required discussions are spread across the work. More of it can be found in section 1.

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: Our model doesn't pose such risks.

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: The assets used are properly cited and credited.

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: The shared data and code is described along with itself.

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: No crowdsourcing done.

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: No such experiments or participants involved.

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.