

Towards Geographic Inclusion in the Evaluation of Text-to-Image Models

Melissa Hall
Meta AI
United States

Samuel J. Bell
Meta AI
France

Candace Ross
Meta AI
United States

Adina Williams
Meta AI
United States

Michal Drozdal*
Meta AI
Canada

Adriana Romero Soriano*
Meta AI, Mila, McGill University,
Canada CIFAR AI Chair
Canada

ABSTRACT

Rapid progress in text-to-image generative models coupled with their deployment for visual content creation has magnified the importance of thoroughly evaluating their performance and identifying potential biases. In pursuit of models that generate images that are realistic, diverse, visually appealing, and consistent with the given prompt, researchers and practitioners often turn to automated metrics to facilitate scalable and cost-effective performance profiling. However, commonly-used metrics often fail to account for the full diversity of human preference; often even in-depth human evaluations face challenges with subjectivity, especially as interpretations of evaluation criteria vary across regions and cultures. In this work, we conduct a large, cross-cultural study to study how much annotators in Africa, Europe, and Southeast Asia vary in their perception of geographic representation, visual appeal, and consistency in real and generated images from state-of-the-art public APIs. We collect over 65,000 image annotations and 20 survey responses. We contrast human annotations with common automated metrics, finding that human preferences vary notably across geographic location and that current metrics do not fully account for this diversity. For example, annotators in different locations often disagree on whether exaggerated, stereotypical depictions of a region are considered geographically representative. In addition, the utility of automatic evaluations is dependent on assumptions about their set-up, such as the alignment of feature extractors with human perception of object similarity or the definition of “appeal” captured in reference datasets used to ground evaluations. We recommend steps for improved automatic and human evaluations. This includes collecting annotations from people located inside and outside the region of interest, instructing annotators on whether they should follow specific definitions of evaluation criteria or utilize their own interpretation, and reporting assumptions underlying automatic evaluations.

*Equal contribution in research and engineering leadership.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

FACCT '24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3658927>

CCS CONCEPTS

• **Social and professional topics**; • **Applied computing**; • **Computing methodologies** → **Computer vision**;

KEYWORDS

text-to-image generation, geography, evaluation

ACM Reference Format:

Melissa Hall, Samuel J. Bell, Candace Ross, Adina Williams, Michal Drozdal, and Adriana Romero Soriano. 2024. Towards Geographic Inclusion in the Evaluation of Text-to-Image Models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FACCT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3630106.3658927>

1 INTRODUCTION

Recent years brought unprecedented progress in generative models for visual content creation, with works achieving impressively photorealistic image generations [18, 35, 43, 46]. However, it is important to understand whether the true diversity of the real world is conveyed in generated images. In particular, generated images should be *representative* of the world, capturing the myriad variability of people, objects, and scenes across geographic regions, while also *visually appealing* or interesting, and *consistent* with the input text description. Assessing these properties is crucial for an in-depth understanding of the performance and potential biases of text-to-image systems.

Commonly used automated metrics, such as the Fréchet Inception Distance (FID) [25], Inception Score [49] and precision and recall (PR)-based metrics [29, 34, 44, 48, 52], already exist for evaluating the representativeness of generated images. Other metrics such as LPIPS [56] and Vendi Score [17] target diversity within a corpus of generated images explicitly. Metrics such as CLIPScore [23] have also been introduced to automatically evaluate the consistency of generated images with respect to the input text description. Recently, several metrics have been extended to measure the ability of text-to-image systems to create faithful depictions of objects around the globe [22].

However, these metrics have notable challenges and are undergoing debates about their robustness [32, 54]. For example, the aforementioned metrics rely on pre-trained feature extractors, such as Inceptionv3 [55], which can lead to unreliable model rankings and susceptibility to undesirable “fringe features” [28]. FID and PR-based metrics rely on reference datasets to define representativeness that may not capture the true diversity of the real world.

(Q1) Select all image(s) that show a dog.

Top Image

Bottom-Left Image

Bottom-Right Image

None

Please select at least 1 option to continue.

(Q2) Review the top image's dog. After reviewing the top image's dog, select the bottom image that shows something MOST SIMILAR to the dog in the top image.

Bottom-Left Image

Bottom-Right Image

(Q3) What is your confidence level in your answer to the previous question (Q2)?

Low - Not confident

Medium - Somewhat confident


High - Very confident

(Q4) Which of the two bottom images is more visually appealing?


Bottom-Left Image

Bottom-Right Image

Group 1:



Group 2:



Does the image show a dog?

Yes

No

Select any region where dogs in that region look similar to the dog shown in this image, based on your personal knowledge and experience. Focus only on the dog, and ignore background scenes and additional objects.

Africa

Europe

Southeast Asia

None / Cannot determine

Please select at least 1 option to continue.

Select all geographic regions that could be depicted in the image, based on your personal knowledge and experience. Focus on the background scene and context, and ignore the dog.

Africa

Europe

Southeast Asia

None / Cannot determine

Please select at least 1 option to continue.

(a) Task 1
(b) Task 2

Figure 1: We collect 65k annotations performed by people located in Africa, Europe, and Southeast Asia corresponding to evaluation criteria of text-to-image models including geographic representation, similarity, visual appeal, and object consistency in real and generated images. We develop recommendations for improved human and automatic evaluations of text-to-image models.

These metrics operationalize representations in a constrained way and fail to account for the diversity of human preferences [40].

Given the complexity and contestedness of automatic metrics, human evaluations remain the gold standard to benchmark text-to-image systems. Human evaluations often focus on side-by-side comparisons of images taken from different sources [10, 35, 43, 46] or task human auditors with identifying concepts or group information in an image based on instructional criteria [5]. However, human evaluations also face challenges. Most notably, they struggle to account for the subjectivity across regions and cultures. Recent studies suggest that human evaluations tend to be somewhat ad hoc, with task design impacting results and leading to high variance in estimates [57]. Furthermore, annotations of model performance and bias presuppose definitions that may not encapsulate the full range of human perceptions, may be mutually incompatible between different groups of people, and may not map to what the model creator is genuinely interested in.

In this paper, we present a human study to understand how annotators conceptualize desirable properties of text-to-image systems across geographic locations. We focus on images of *objects*, as notable disparities have been observed between geographic regions in both text-to-image generative systems [5, 6, 22] and discriminative models [12, 21, 45, 53].

We collect over 65,000 annotations with 20 qualitative survey responses to study how people in Africa, Europe, and Southeast Asia vary in their perception of geographic representation, visual appeal, and consistency. The study includes both real, geographically diverse images and images created by two state-of-the-art generative systems that depict six types of cross-culturally common objects. We contrast the results of the human annotations with commonly used automated metrics. Through our study, we find that:

- Perceptions of geographic representation, appeal, and consistency vary across annotator locations.
- Annotators can reinforce geographic stereotypes and contradict each other in their interpretation of appeal. Preferred qualities include realism, foregrounding, exaggerated backgrounds, bold colors, and luxurious objects.
- Annotators vary in how much they follow concrete definitions or apply their subjective interpretation when completing tasks. Some rely on personal knowledge while others utilize external reference sources.
- CLIP and DINOv2 feature extractors correlate better with human judgment of similarity than Inceptionv3.
- Although CLIPScore is highly reliable in assessing the presence of an object in an image when used with a threshold, it is generally unreliable for assessing geographic representation.

We release the subset of annotations corresponding to object consistency and regional representation in real images to enable further analysis of geographic inclusion in evaluations. They may be found here.

In addition, we recommend steps to improve evaluations of text-to-image generative models. To strengthen human evaluations, we suggest clearly instructing annotators about whether they should attempt to capture their subjective perspective or to follow specified definitions, and whether they should leverage external resources or rely only on their existing knowledge. In addition, we suggest collecting perspectives from both in-region and out-of-region annotators to capture more breadth. To strengthen automatic evaluations, we suggest selecting reference data sources with thoughtful consideration of the interpreted criteria and using more modern feature extractors that better capture image qualities relevant for the evaluation, such as shape. When reporting evaluations, avoid “majority-vote” types of aggregations that reduce uncertainty to

binary indicators of evaluation criteria, and communicate clearly the assumptions of criteria definitions used in the measurement process.

We hope that this work strengthens evaluations of geographic representation in text-to-image generative models and paves the way towards text-to-image systems that truly work for everyone.

2 EXPERIMENTAL SET-UP

We outline evaluation criteria for text-to-image models and describe our annotations, summarized in Table 1.

2.1 Evaluation Criteria for Text-to-Image Generative Models

We discuss evaluation criteria for text-to-image models and automatic metrics used for their approximation. In later analyses, we encapsulate the realism / diversity and similarity criteria into an overall focus of geographic representation. We also study image visual appeal and object consistency.

2.1.1 Realism / Diversity. Realism is the degree to which generated samples *resemble* the real world, and diversity is the degree to which samples capture the *full variability* of the real world [22, 29, 48]. Prior works [5, 22] have utilized realism and diversity in the assessment of the extent to which the real-world variability and nuance of a given region is depicted in a corpus of generated images. We call this “geographic representation.”

Automatic metrics. Image realism and diversity are often measured with precision and coverage, respectively [22, 29, 48]. Precision highlights whether generated images look similar to real images, while coverage quantifies whether generated images capture the breadth of variation in real images. Specifically, precision measures the proportion of generated images that lie within the manifold of real, reference images. Coverage measures the proportion of real images that have generated images nearby. In Section 3, we analyze the vulnerability of these metrics to predefined feature extractors and reference manifolds through the context of similarity, which we introduce next. In addition, the distances *i.e.*, cosine similarity, between the CLIP text embedding corresponding to group terms and embedding of a generated image has been used to identify demographic groups [9] and geographic representation [5]. We use the CLIPScore [23] with region information (“Region-CLIPScore”) as an indicator of geographic representation, analogously to Lee et al. [31].

2.1.2 Similarity. While similarity, *i.e.*, whether two images closely resemble each other, is not commonly evaluated in text-to-image generative models, it is the foundation of criteria like realism and diversity. For an object in an image to have “realism,” it needs to resemble instances of the object in the real world. For a dataset of images to display enough “diversity,” the objects in the images must meaningfully differ from each other.

Automatic metric. Measurements of similarity are closely related to automatic evaluations of realism and diversity, as they approximate relative distances between images in a predefined feature space. In Section 3, we study how the older, more ubiquitous Inceptionv3 [55] compares to the more recent CLIP ViT-B/32 [41]

and DINO ViT-L/14 [36] feature extractors trained on larger data sources in approximating human perceptions of similarity.

2.1.3 Visual Appeal. Visual appeal corresponds to whether images have visual attractiveness or interest and has become an increasingly common evaluation criteria [10, 38, 47]. We study how automatic metrics capture subjectivity in appeal across geographic regions.

Automatic metric. According to some previous works, appeal can be automatically measured by comparing to a manifold of real images with the precision metric [22, 29]. In Section 3, we study whether generated images that fall in the real image manifold align with human perceptions of visual appeal.

2.1.4 Object Consistency. Consistency corresponds to whether an image includes all components of its prompt. This relates to visual concreteness [7] and the relationship between a concept’s meaning and its human-perceptible form [37]. Per [22], we focus on objects included in the prompt, *e.g.*, whether an image generated with “car in Africa” depicts a car.

Automatic metric. In Section 3 we explore the CLIPScore [23], which can be used to measure the input-output consistency of the generative model [22, 31, 47]. Following [22], we use the text embedding corresponding to the object that should be present in the image and refer to the metric as “Object-CLIPScore.”

2.2 Annotations

We now describe the annotations we collected to better understand how humans in different geographic locations vary in their perception of the aforementioned evaluation criteria for text-to-image models.

2.2.1 Images. Our tasks include real images from GeoDE [42], a diverse, geographically representative dataset of images taken across multiple regions. Each image contains an object filling at least 25% of the image and occurs in everyday settings. We also include generated images from a multi-modal implementation of a generative pre-trained transformer leveraging CLIP image embeddings. This model, which we call “DM w/ CLIP”, has approximately 3.5 billion parameters [43]. We also use a latent diffusion model trained on a public dataset of approximately 5 billion images, excluding explicit material [46], which we refer to as “LDM 2.1.” We use paid versions of both models and generate images using the prompts “{object}” and “{object} in {region}” to capture representations across regions.

We focus on six objects (*bag, car, cooking pot, dog, plate of food, and storefront*) that exist across regions [42] and have been used to audit generative and discriminative models [4, 12, 22, 45]. We include the regions *Africa, Europe, and Southeast Asia*. Generated images of these objects vary in their geographic representation and consistency [22], as shown in randomly selected images in Figure 2. For example, generated images of *car, cooking pot* and *storefront* often depict geographic stereotypes not seen in the real dataset, such as rudimentary infrastructure, low-income tools, and vibrant colors for Africa and Southeast Asia. Similarly, generated images have exaggerated backgrounds, with rural and cobblestone/alley-filled scenes for Africa and Europe, respectively, that are rarely

Evaluation Criteria	Definition	Automatic Metric	Our Annotator Data
Realism / Diversity	Whether, for a geographic region, generated samples resemble the real world / capture the variability of the real world	Precision / Coverage	Geographic representation of objects: "Which regions have similar objects?"
		Region-CLIPScore	Geographic representation of backgrounds: "Which regions have similar backgrounds?"
Similarity	Whether two images closely resemble each other	Distance	Similarity: "Which image has an object more similar to the reference image?"
Visual appeal	Whether images have visual attractiveness or interest	Precision	Visual appeal: "Which image is more appealing?"
Object consistency	Whether images show elements of prompt used in generation	Object-CLIPScore	Object consistency: "Is the object shown?"

Table 1: We study geographic representation (encapsulating realism/diversity and similarity), visual appeal, and object consistency. The table summarizes their definitions, metrics used in their approximation, and relevant annotation data that we collected.

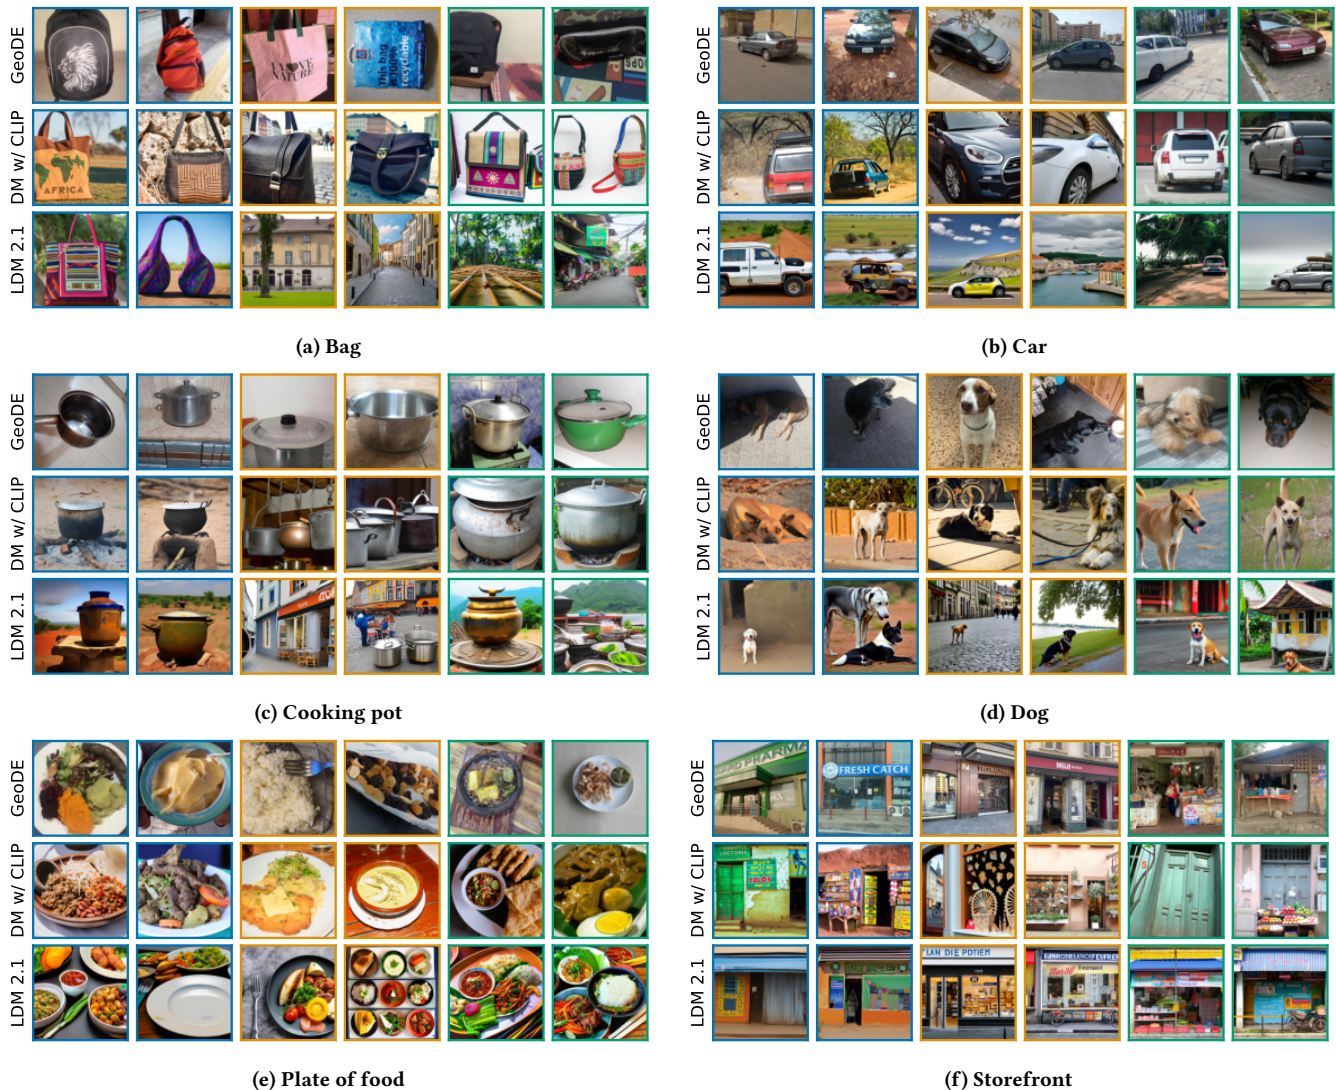


Figure 2: Random examples of real images from GeoDE in each region and generated images from DM w/ CLIP and LDM 2.1 using the prompt {object} in {region}. The first two columns correspond to Africa, the next to Europe, and the last to Southeast Asia.

depicted in GeoDE. Furthermore, LDM 2.1 shows consistency issues with intended objects like *bag*, *car*, and *cooking pot* often left out of generated images.

2.2.2 Tasks. Our study comprises two main annotation tasks that we populate with different image combinations, balanced across objects and regions. Figure 1 shows examples of the user interface for Task 1 and Task 2.

Task 1: Image Comparison & Object Consistency. The first task, in Figure 1a, focuses on understanding object consistency and human perceptions of image similarity and appeal. As with classical ABX testing [33], we display triplets of a single reference image and two images for comparison, one real and one generated. The annotator is asked to indicate whether the intended object is shown in each image, providing insight into their perception of *object consistency*. Then, they are asked to select which of the two comparison images shows an object more similar to the object in the reference image and to provide a confidence score to their answer. This allows us to measure potential variations in annotator perception of *similarity*. Finally, annotators decide which comparison image is more *visually appealing*.

For Task 1, images in a given triplet have the region in common or are generated with the object prompt. Generated images in a triplet are from the same model. Across triplets, the comparison images vary in their distance from each other and from the reference images in the Inceptionv3 feature space. We use Inceptionv3 as it is well-established when assessing performance of generative models [8, 24, 43, 46] (and we study additional feature extractors in our analysis of similarity). We define the triplet distance as $\|f(I_{ref}) - f(I_{far})\|^2 - \|f(I_{ref}) - f(I_{close})\|^2$, where f represents the feature extractor, I_{ref} is the reference image, I_{close} is the closer comparison image to the reference in the embedding space, and I_{far} is the further image. A higher triplet distance means that one comparison image is much further from the reference image than the other comparison image, suggesting it is much more “similar” to the reference image. For every model, we compile 425 triplets for each object-region combination, yielding 15,300 triplets. Appendix A.1 contains additional details about the triplets.

Task 2: Geographic Representation & Object Consistency. The second task, in Figure 1b, focuses on understanding consistency and geographic representation of objects and scenes in images. A single image is shown, and annotators are asked to indicate if the intended object is present. If so, they are then asked to select any region from a pre-specified list that they feel the object could be from, i.e., that they feel has objects from the same category (e.g. “bag”, “car”) that look similar to the object in the image. They are also asked the same question about the background scene / context depicted in the image. These questions correspond to annotator perception of geographic representation.

For Task 2, we randomly select 100 images from GeoDE for each region-object combination. In addition, for LDM 2.1 and DM w/ CLIP, we used 100 images with generated {object} in {region} for each object-region combination and 100 images generated with {object} for each object. This yields 6,600 images.

2.2.3 Annotators. Each task is annotated by three people from our pool: one person located in Africa, Europe, and Southeast Asia.

The annotator pool consists of 15-25 people in each region. After both tasks are completed, we filter out approximately 13% of Task 1 annotations and 2% of Task 2 annotations that do not meet quality checks on real images. The remaining annotations yield the corpus with which we perform our analyses. Appendix A.1 describes the annotator training process and compensation, quality criteria, and statistics about annotator background and job execution.

2.3 Annotator Survey

Once all annotations were collected, we also elicited voluntary survey responses to learn more about how annotators interpreted questions related to geographic representation. The first set of questions focuses on the depiction of objects. We ask three versions of the question “When performing these tasks, how did you determine if an object looked like similar objects in [region]? If possible, please provide specific examples,” where we consider Africa, Europe, and Southeast Asia. We also ask, “Did you ever mark that none of the regions had similar objects as the ones shown in the image, or that you could not determine? If so, what prompted you to make that decision?” For the second set of questions, we asked the same questions, but about the “background scene / context” of the image. We received 23 responses to the annotator survey, from 5 annotators in Africa, 8 in Europe, and 10 in Southeast Asia. We applied an inductive coding approach to categorize themes discussed by the annotators, resulting in 16 descriptive codes. After an initial round to validate the coding scheme, responses were individually coded by three of the authors. We binarize the presence of each code per response and evaluate inter-coder agreement using Fleiss’ Kappa [16]. We discard codes with a coefficient less than 0.6 to ensure consistent deployment by coders. Final codes were assigned via a majority vote over the three authors, and responses were grouped according to whether they correspond to background scene or object. All codes and their Fleiss’ Kappa coefficients can be found in Appendix Table 3.

3 RESULTS

For each criteria, we first present analyses corresponding to *Human Interpretations*, leveraging only annotation and survey data (where relevant). Then we discuss the *Human & Automatic Metric Interaction*, where we analyze the extent to which automatic metrics capture human interpretation of evaluation criteria. Along the way, we suggest recommendations for improved human and automatic evaluations of text-to-image models based on our findings.

3.1 Geographic Representation

3.1.1 Human Interpretations: Realism / Diversity. We first analyze how perceptions of geographic representation in objects and backgrounds depicted in images differ based on where the annotator is located. In Figure 3a, we focus on real images that were taken in a specific region and generated images created with prompts mentioning that region. We compare perceptions of geographic representation between in-region and out-of-region annotators. We find that annotators located in the same region as a *real* image tend to identify both the object and the background as representative of




Figure 3: (a) Proportion of images that in-region and out-of-region annotations consider depicting geographic representation for objects and backgrounds. (b) Relationship between Region-CLIPScore and annotator designation of geographic representation. The x-axis shows the average CLIPScore for all images within a bucket with size 0.01. The y-axis shows the proportion of images where annotators said the object was present. We include 95% confidence intervals generated via bootstrapping. Annotator perceptions of geographic representation differ according to whether the annotators are located in the region of focus or outside it. Region-CLIPScore does not always capture variations in perceived geographic representation across annotator location.

Theme	Quote	Re: Region
Built environment	R5: "The types of houses are really distinct"	Africa
Culture, art & religion	R6: "distinctive religious architecture such as Buddhist temples or mosques"	Southeast Asia
Absence of detail	R11: "simply ground or tile floor with international design"	NA
External search	R16: "Sometimes I used Google picture to search the photo"	Africa
Personal lived experience	R5: "I'm familiar with European environments so I compare it to them"	Europe
Nature & natural world	R6: "Presence of tropical vegetation such as palm trees or bamboo"	Southeast Asia
People	R7: "Black and brown people in the background"	Africa
Stereotyped representations	R19: "If place is classy or elegant"	Europe
Stereotyped representations	R7: "houses were not developed and tools surrounding the objects were rudimental"	Africa
Stereotyped representations	R6: "traditional houses like still houses"	Southeast Asia

Table 2: Selected survey responses about geographic representation for themes that satisfy our inter-annotator agreement threshold. "R[N]" indicates respondent ID. "Re: Region" indicates the region specified in the question. Annotators deploy positive and negative stereotypes and use personal experience and external resources when considering geographic information in images.

that region moreso than annotators located in other regions. However, perceptions between in-region and out-of-region annotators for *generated* images are more mixed.

To better understand why in- and out-of-region annotators disagree, we include example images in Figure 4. In-region annotators often consider objects and backgrounds with few geographically-distinguishing features as representative more than out-of-region annotators. For example, this is the case for real images depicting bags without geographically-associated patterns (Figure 4a). It is also true of generated images of cars and dogs with backgrounds dominated by agnostic outdoor streets (Figure 4b). Meanwhile, when generated images show stereotypical objects, such as simple cooking pots for Southeast Asia from LDM 2.1 and saucy pastas

and "meat-and-potato" meals for Europe from DM w/ CLIP (as in Figure 4c), out-of-region annotators often consider the images more representative than in-region annotators do. Furthermore, stereotypically rural scenes and rudimentary infrastructure are common in images of Africa and Southeast Asia that out-of-region annotators consider geographically representative and in-region annotators do not (Figure 4d).

The results of the qualitative survey (see Table 2 for a summary) also suggest an over-reliance on stereotypical image features among out-of-region annotators, e.g., the expectation that European backgrounds be "classy or elegant" (respondent R19, located in Southeast Asia), or that African "houses were not developed" (R7, located in Europe). We find that annotators report using a range of

image features in determining geographic representation in both background scene and object, with a greater proportion of Europe-based respondents relying on features including buildings, nature, and people than annotators based in other locations. In addition, when asked about the geographic representation of objects and background scenes, annotators frequently refer to common stereotypes when explaining their judgments, such as “stilt houses” for Southeast Asia. Finally, some annotators reported using external references such as “Google picture to search the photo,” which can introduce additional biases from online web data. We quantify these findings according to annotator location in the Appendix (Figure 9).

Recommendation: When annotating for geographic representation, include in-region and out-of-region annotators: In-region annotators tend to indicate realistic, visual qualities representative of a region while out-of-region annotators often indicate unique, exaggerated identifiers of a region. Contrasting these perspectives can illuminate regional stereotyping in generated images.

3.1.2 Human Interpretations: Similarity. We also study how annotators vary in their assessment of image similarity. Figures 5a and 5b show examples of image triplets for which the annotators agree and disagree, respectively, on which comparison image is more similar to the reference image. We find that factors in perceptions of object similarity include color, size, type (e.g., dog breed), and camera angle.

In addition, in Figure 5c, we stratify rates of annotator agreement in similarity by the source of the triplet images and the region where the image was taken / prompted with. We do not find that annotators agree on similarity more consistently for any image region. However, annotators agree the least when determining which DM w/ CLIP comparison images are more similar to a GeoDE reference image. But when this is swapped, and they are selecting between two GeoDE images compared to a single DM w/ CLIP image, they have a much higher rate of agreement. We hypothesize that this asymmetry may be due to the relatively low conditional diversity between the DM w/ CLIP images as compared to GeoDE (observed in Figure 2): Selecting between less diverse images can make agreeing on the more similar image challenging, whereas two very different comparison images may allow for easier consensus about which is more similar. Such variation in annotator perception recalls prior work suggesting alternatives to majority vote aggregations of annotations [11].

Recommendation: Collect multiple annotations for image similarity. Aggregations like majority-vote may mask the strength of disagreement between annotators or obfuscate factors relevant to the minority perspective.

3.1.3 Human & Automatic Metric Interaction: Region-CLIPScore. We next analyze how Region-CLIPScore approximates perceptions of geographic representation. Higher CLIPScore values indicate stronger similarity between the image embedding and corresponding text embedding. For each region, we calculate the Region-CLIPScore

for all images and compare it to the proportion of images that annotators designate as representative. Results are shown in Figure 3b.

We find that Region-CLIPScore tends to correspond to a larger proportion of real images perceived as representative of that region by in-region annotators than out-of-region annotators. Furthermore, Region-CLIPScore has a larger disparity in its approximation of geographic representation across annotator locations when used for identifying Southeast Asia representation than Africa and Europe. In addition, Region-CLIPScores can be 2 – 5 points higher for generated images of Africa and Southeast Asia, which tend to show strong regional stereotyping, than for Europe and real images. This suggests that Region-CLIPScore is susceptible to stereotypical representations of regions.

Recommendation: Avoid Region-CLIPScore or use it with caution: it does not capture variation in perception of geographic representation between annotator locations and may favor stereotypes over realistic images.

3.1.4 Human & Automatic Metric Interaction: Distance. We next study the alignment between annotator perception of similarity and relative distances between reference and comparison images using different feature extractors,¹ with results shown in Figure 6. We consider a good feature extractor one that has a high rate of alignment between the comparison image that is closer to the reference image in the feature space and the image that annotators consider more similar across all triplet distances, indicating that it correlates well with human judgment.

For all models, the chance that annotators select the closer comparison image as the more similar one is close to random for low triplet distance values. However, the probability of alignment grows as the triplet distance increases, and annotators almost always identify the image that is closer to the reference image as the more similar one at large distances. Both DINOv2 and CLIP features correlate significantly better with human judgement than InceptionV3, with DINOv2 features are slightly more associated with human judgment than CLIP features at low distance values. This is true for all image regions and annotator locations, suggesting there is little effect of geography in assessing similarity.

Recommendation: Use modern feature extractors like CLIP and DINOv2 for automatic measures of object similarity.

3.2 Visual Appeal

3.2.1 Human Interpretations. To understand annotator subjectivity in determining visual appeal, we measure annotator agreement in deciding which image looks more appealing to them (Figure 7a). We observe that the agreement is the lowest for images depicting Europe. Furthermore, annotators tend to agree on which image is more appealing more often for GeoDE and LDM 2.1 than DM w/ CLIP. In the Appendix, Figures 10a and 10b show examples where

¹We correct for the bias of using InceptionV3 distances when selecting images for annotations by subsampling to triplet distances for which there exist samples for all feature extractors. We normalize the plot by the maximum triplet distance and rescale the values to compare different feature extractors.

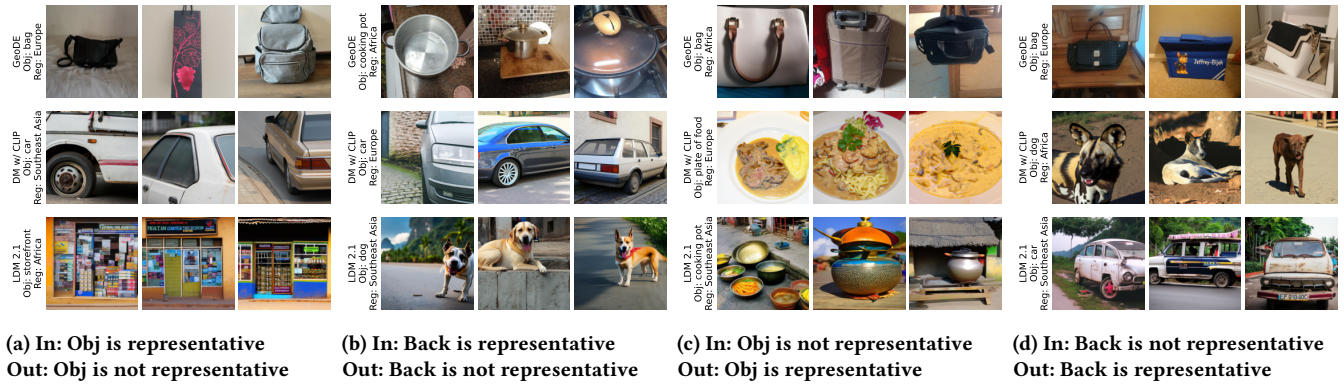


Figure 4: Examples of disagreement between in- and out-of-region annotators about geographic representation of *objects* and *backgrounds*. Out-of-region annotators tend to consider stereotypes representative more than in-region annotators.

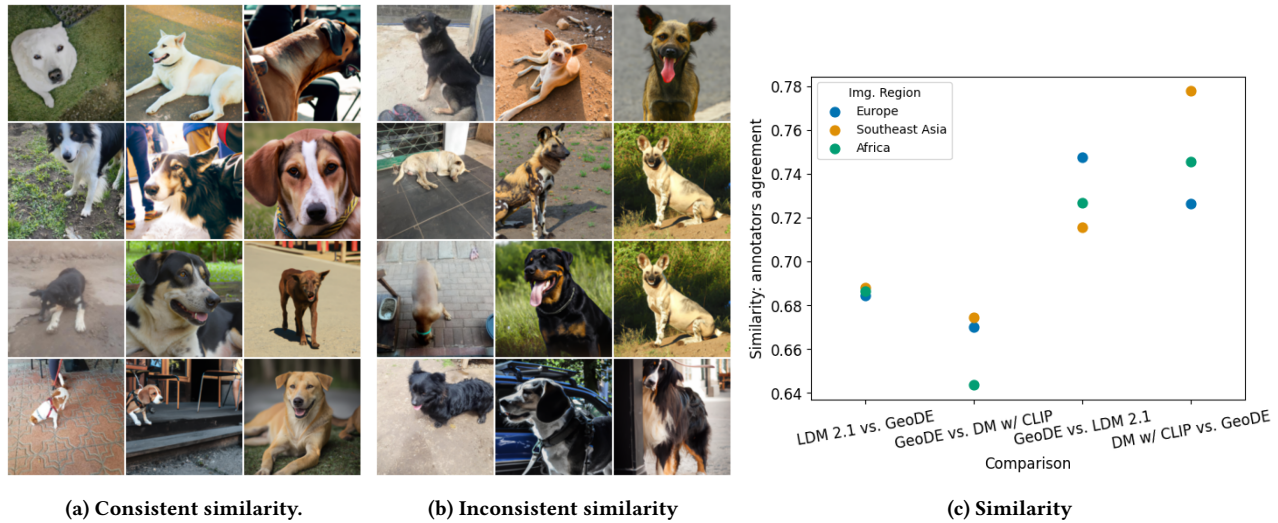


Figure 5: (a-b) Qualitative examples of consistent and inconsistent annotator perceptions of similarity, where (a) depicts (from left): reference image from GeoDE and comparison images from DM w/ CLIP designated as more and less similar by all three annotators and (b) depicts (from left): reference image from GeoDE and two comparison images from DM w/ CLIP with inconsistent similarity annotations. (c) Rates of annotator agreement in similarity. Variations in perception of object similarity can depend on the diversity of the images. Factors in perception of similarity include color, size, type (e.g., dog breed), and camera angle.

annotators agree and disagree, respectively, that one images is more appealing. We find that consistent appeal can occur when one sample has significantly lower visual quality. However, agreement is lower when annotators are shown two high quality images.

Moreover, annotators in different locations tend to agree more on similarity than appeal, suggesting perceptions of appeal are more variable across regions. This could be because similarity is grounded in a reference image, while visual appeal evokes latent interpretations of attractiveness. In addition, Figure 7b shows that annotators in Southeast Asia more often select the same image as more similar *and* appealing, suggesting that their interpretation of appeal coincides with that of similarity moreso than annotators in other regions. For annotators across all locations, appeal coincides

with similarity the least for images that correspond to their location, suggesting that annotators may leverage qualities for appeal beyond similarity more when they have greater familiarity with what is presented in the image.

Recommendation: Collect annotations of visual appeal from annotators in multiple regions. While collecting more preferences will likely not resolve disagreements in aesthetic preferences, they may suggest directions for more inclusive evaluation criteria or definitions. Furthermore, they may be useful in challenging a universal notion of appeal by pinpointing specific preferences that vary between certain groups.




Figure 6: Frequency that annotators consider the closer comparison image to the reference image as having the similar object as a function of triplet distance. CLIP and DINOv2 best reflect annotator perceptions of object similarity across all locations.




Figure 7: (a) Annotator agreement in visual appeal. (b) Co-occurrence of appeal and similarity. (c) Proportion of generated appealing images that fall in the GeoDE manifold. Annotators agree on appeal less than similarity and at different rates between depicted regions. They may leverage qualities for appeal beyond similarity more when they have greater familiarity with what is in the image. The majority of appealing annotations of generated images fall in the manifold of real images.

3.2.2 Human & Automatic Metric Interaction: Manifold Presence. Because visual appeal is approximated with precision, we construct an analogue to precision by counting the proportion of visually appealing annotations for generated images that fall within the manifold of all the real images in GeoDE. In this set-up, GeoDE is representative of the real world’s geographic diversity. Figure 7c shows that for most objects and regions, the majority of visually appealing annotations of generated images fall within the manifold of real images.

In Appendix Figure 11a, we inspect examples and find that annotators often pick images similar to those in GeoDE are more appealing, *i.e.*, realistic, with the object in at least 25% of the image, and high quality [42]. This tends to occur more for images that fall within the manifold, although some appealing images with these characteristics are outside the manifold. However, we also find that some annotators’ interpretation of “visually appealing” diverge from GeoDE and fall outside of the manifold. Examples are shown in Appendix Figure 11b. In some cases, annotators give greater weight to the aesthetics of the background than the object. This is especially true when images have greater background stereotyping,

such as for images of *cars* in LDM 2.1. Annotators can also prefer realistic objects even if they are lower quality or smaller. For example, a far-away dog is marked as more appealing than a close-up dog with unrealistic colors. The “niceness” of objects also relates to visual appeal as annotators tend to consider dirty, rusty cars, dull pots, and plain bags less appealing than newer, more pristine, and colorful counterparts.

Recommendation: When using manifold-based metrics to measure visual appeal, select a reference dataset that aligns with the definition used by annotators or needed for model deployment. When reporting on appeal, clarify whose and what definition is captured, including whether appeal is associated with exaggerations or stereotypes.

3.3 Object Consistency

3.3.1 Human Interpretations. Finally, we investigate the rate at which annotators identify the desired object in a photo and the extent to which this coincides with the geographic representation of

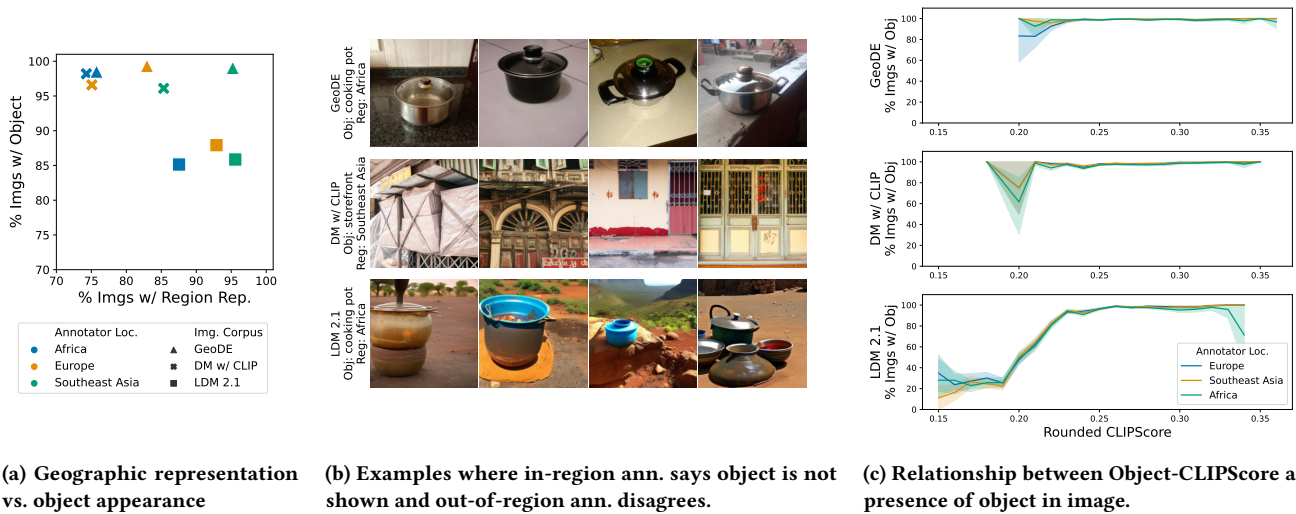


Figure 8: (a) Relationship between object consistency and geographic representation. (b) Disagreements in consistency between annotators. (c) Comparison of Object-CLIPScore and annotator perception of object consistency. Annotators in different regions vary slightly in their recognition of concrete objects. Beyond a threshold, Object-CLIPScore corresponds to consistent presence of objects for all annotator locations.

the image. In Figure 8a, we find that annotators across all locations consider LDM 2.1 as having the least object consistency while also indicating that LDM 2.1 tends to have the greatest geographic representation. This aligns with qualitative observations from [22] that LDM 2.1 often depicts strong geographic representation at the cost of object consistency.

In Figure 8b, we visualize examples in which in-region annotators say the desired object is not shown while out-of-region annotators say it is. For example, annotators located in Southeast Asia do not consider a subset of images from DM w/ CLIP generated with the prompt “storefront in Southeast Asia” as storefronts while annotators located outside of Southeast Asia do. The same pattern is true of in-region and out-of-region annotators for generated images of a “cooking pot in Africa.” This suggests that annotators located in different regions vary in what they consider to be accurate visual representations of concrete objects. Interestingly, in-region annotators can even disagree about objects in real images that were collected from people living in that region, as in the example of cooking pots from GeoDE in the Figure. This suggests possible *intra*-region variations in the definition of objects as well as *inter*-region variations.

Recommendation: When annotating for object consistency, include perspectives from in- and out-of-region. Multiple annotators per region are ideal for capturing within-region variations. If uniformity is desired, instruct annotators to focus on a specific object definition / representation, including positive and negative examples.

3.3.2 Human & Automatic Metric Interaction: Object-CLIPScore. In Figure 8c, we plot bucketed CLIPScores and average per-image

annotations of whether the object appears in the image. By examining LDM 2.1, which suffers from some consistency challenges, we find that the lower range of CLIPScores is a useful indicator of whether the object appear in the image. This is aligned with findings in [22]. However, at higher scores the object is consistently in the image. As with the Region-CLIPScore, we see that the range of CLIPScores can vary between types of images. We find empirically that a threshold of 0.21-0.25 best indicates consistent appearance of an object. This range accounts for variations between datasets, with a threshold 0.21 corresponding to nearly constant object presence for GeoDE and DM w/ CLIP and 0.25 for LDM 2.1. Unlike the Region-CLIPScore, the Object-CLIPScore seems equally informative regardless of which location the annotator is in.

Recommendation: Use a threshold for Object-CLIPScore to include multiple representations of the same concept.

4 RELATED WORK

Beyond text-to-image models, our work is situated within a wider discussion of geographic and cultural performance disparities in computer vision systems. For example, De Vries et al. [12] report that object detection models suffer marked drops in performance when evaluated on images from non-Western geographies, and particularly so for those from lower-income countries. Goyal et al. [20] find similar disparities in the performance of visual feature extractors, as does Gustafson et al. [21] for the popular image-text foundation model CLIP [41]. These failures are likely to be, at least in part, the result of poor representation in the training datasets: ImageNet [13] displays a sample that is overwhelmingly biased toward North American and Western European contexts [51]. Remarkably, Richards et al. [45] report that as performance has improved on

classic object detection benchmarks, the performance gap *between* regions has worsened.

Accounting for geographic context is a challenge not limited to vision systems: in NLP, a growing body of research investigates subjectivity among annotators [2, 11, 19] and, in particular, cross-cultural sensitivity. While cultural diversity is not exactly proxied by geographic diversity, we expect cultural differences to often co-occur with geographic shifts, and consider work on cross-cultural inclusion deeply relevant to our work. Researchers working on machine translation have long sought to improve translation both to and from so-called low-resource languages, though models continue to exhibit failures along culturally-specific axes [1]. Recent work explores the role of cultural context in automated detection of stereotypes [14, 27], detecting toxic text or hate speech [30], and values espoused by large language models [3, 15]. Others have focused on how evaluation practices suffer from a narrow cultural frame. Hutchinson et al. [26] explore how cultural factors are overlooked during NLP system evaluation, highlighting the obfuscation of annotators' "socio-demographic standpoints" [p. 1867]. Similarly, Prabhakaran et al. [39] call for "culturally situated" evaluations [p. 3], while Sambasivan et al. [50] argue that fairness evaluations must account for non-Western perspectives on fairness itself. Our work responds to such calls, exploring the extent to which automated evaluation metrics can account for geographic differences and cultural expectations.

5 CONCLUSION

We study how annotators in different regions vary in their perceptions of evaluation criteria for text-to-image generative models and how well automatic metrics capture these variations. From our analyses, we recommend that annotations of geographic representation include perspectives from people located inside *and* outside regions of focus and that instructions clarify whether subjectivity and reference to external sources are encouraged. In addition, modern feature extractors like CLIP and DINOv2 can better capture aspects of image similarity when performing automatic evaluations than the standard Inceptionv3. We find that visual appeal is often interpreted in contradictory preferences that can reinforce stereotypes and recommend improving automatic metrics like precision via careful selection of reference datasets. Perceptions of object consistency can have *within* region variations but are generally well approximated by CLIPScore across all annotator locations. Finally, contradictory responses for criteria like appeal, similarity, and consistency can reveal true ambiguity that is lost in "majority-vote" type aggregations. We also release annotations corresponding to geographic representation and consistency in real images to enable further study. We hope that this work is a meaningful step towards greater geographic inclusion in human and automatic evaluations of text-to-image generative models.

Limitations. Our study is limited to three broad regions, six objects, and three image sources. Annotators may vary in their familiarity of objects presented in the study. Due to the vast variations in human subjectivity within regions, findings based on annotator responses are limited in their extensibility to those not included in the study. Task construction may also have affected annotator

responses, and conclusions are limited without a deep focus group to collect qualitative feedback and explanation from annotators.

ACKNOWLEDGMENTS

We thank Carleigh Wood, Ida Cheng, Emerson Bacud, and Adam Hakimi for their many contributions throughout the annotation process. We also thank Laura Gustafson, Mark Ibrahim, and Pietro Astolfi for providing technical advice.

Ethical considerations. This work would not have been possible without dozens of individual human annotators. In our data collection process, we followed recommendations pertaining to responsible data collection, including (i) compensating annotators fairly for their work at an hourly rate, (ii) respecting the privacy of annotators by using anonymous ids in the data collection process, (iii) avoiding questions about personal information, (iv) providing voluntary and informed consent about the annotation task and use of data, (v) providing multiple bi-directional channels of communication with annotators, and (vi) allowing annotators to opt-out of the task at any time. In addition, we attempted to mitigate risks of exposure to harmful or uncomfortable content by focusing on images of everyday objects rather than people or divisive topics.

Researcher positionality. The authors of this work recognize the role that our personal experience and institutional resources and constraints play in the annotation collection process and subsequent analyses and recommendations. The choice of questions, objects, regions, and images included in the annotation tasks were selected with the intention of increasing the variety of perspectives about annotator perceptions of evaluation criteria within our infrastructure and financial constraints. While we attempted to mitigate bias in our interpretation of quantitative annotations by including a qualitative survey, not all annotators completed the survey. We therefore remain susceptible to our subjective interpretation of responses. Furthermore, it is likely that nuance in annotator interpretation of the evaluation criteria is lost in our quantification and summarization of analyses, even with our effort to inspect qualitative examples. In addition, the interpretation and coding of the qualitative survey data was conducted by three authors based in Europe or North America. Our inductive coding approach is likely to be biased by our own experiences. Finally, our recommendations pre-suppose a degree of resourcing that is not accessible to all, such as the potentially costly suggestion of collecting multiple annotations for a single sample to capture subjectivity. While we hope this work informs improved evaluations of text-to-image generative models, there is much room for further investigation of the metrics and annotator behaviors.

Adverse impacts. Findings about variations in human and automatic evaluations studied in this work may differ when applied to other types of images, such as depictions of people, complex scenes, or creative imagery, or across different demographic groups such as gender or age. Readers should use caution when extending findings to broader regional trends or when informing future annotator patterns. A future study would be strengthened by including multiple annotators per region for the same task, to account for intra-region variability.

REFERENCES

- [1] Idris Akinade, Jesujoba Alabi, David Adelani, Clement Odoje, and Dietrich Klakow. 2023. Epsilon kú mask: Integrating Yorùbá cultural greetings into machine translation. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*. Association for Computational Linguistics, Dubrovnik, Croatia, 1–7. <https://doi.org/10.18653/v1/2023.c3nlp-1.1>
- [2] Hala Al Kuwatly, Maximilian Wich, and Georg Groh. 2020. Identifying and Measuring Annotator Bias Based on Annotators' Demographic Characteristics. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, Seyi Akiwowo, Bertie Vidgen, Vinodkumar Prabhakaran, and Zeerak Waseem (Eds.). Association for Computational Linguistics, Online, 184–190. <https://doi.org/10.18653/v1/2020.alw-1.21>
- [3] Arnab Arora, Lucie-aimée Kaffee, and Isabelle Augenstein. 2023. Probing Pre-Trained Language Models for Cross-Cultural Differences in Values. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, Sunipa Dev, Vinodkumar Prabhakaran, David Adelani, Dirk Hovy, and Luciana Benotti (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 114–130. <https://doi.org/10.18653/v1/2023.c3nlp-1.12>
- [4] Hritik Bansal, Da Yin, Masoud Monajatipoor, and Kai-Wei Chang. 2022. How well can Text-to-Image Generative Models understand Ethical Natural Language Interventions?. In *EMNLP (Short)*.
- [5] Abhisa Basu, R Venkatesh Babu, and Danish Pruthi. 2023. Inspecting the Geographical Representativeness of Images from Text-to-Image Models. *arXiv preprint arXiv:2305.11080* (2023).
- [6] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2022. Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale. *arXiv:2211.03759* [cs.CL]
- [7] Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. Concrete ratings for 40 thousand generally known English word lemmas. *Behavior research methods* 46, 3 (2014), 904–911.
- [8] Arantxa Casanova, Marlene Careil, Jakob Verbeek, Michal Drozdal, and Adriana Romero Soriano. 2021. Instance-conditioned gan. *Advances in Neural Information Processing Systems* 34 (2021), 27517–27529.
- [9] Jaemin Cho, Abhay Zala, and Mohit Bansal. 2022. DALL-Eval: Probing the Reasoning Skills and Social Biases of Text-to-Image Generative Transformers. (2022). *arXiv:2202.04053* [cs.CV]
- [10] Xiaoliang Dai, Ji Hou, Cihui-Yao Ma, Sam Tsai, Jialiang Wang, Rui Wang, Peizhao Zhang, Simon Vandenhende, Xiaofang Wang, Abhimanyu Dubey, Matthew Yu, Abhishek Kadian, Filip Radenovic, Dhruv Mahajan, Kungpeng Li, Yue Zhao, Vladan Petrovic, Mitesh Kumar Singh, Simran Motwani, Yi Wen, Yiwen Song, Roshan Sumbaly, Vignesh Ramanathan, Zijian He, Peter Vajda, and Devi Parikh. 2023. Emu: Enhancing Image Generation Models Using Photogenic Needles in a Haystack. *arXiv:2309.15807* [cs.CV]
- [11] Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2021. Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations. *CoRR* abs/2110.05719 (2021). *arXiv:2110.05719* <https://arxiv.org/abs/2110.05719>
- [12] Terrance De Vries, Ishan Misra, Changhan Wang, and Laurens Van der Maaten. 2019. Does object recognition work for everyone?. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 52–59.
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [14] Sunipa Dev, Jaya Goyal, Dinesh Tewari, Shachi Dave, and Vinodkumar Prabhakaran. 2023. Building Socio-culturally Inclusive Stereotype Resources with Community Engagement. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=u1j1Dc8k6>
- [15] Esin Durmus, Karina Nyugen, Thomas I. Liao, Nicholas Schiefer, Amanda Askill, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish, Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli. 2023. Towards Measuring the Representation of Subjective Global Opinions in Language Models. *arXiv:arXiv:2306.16388*
- [16] J.L. Fleiss et al. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin* 76, 5 (1971), 378–382.
- [17] Dan Friedman and Adji Bousso Dieng. 2022. The Vendi Score: A Diversity Evaluation Metric for Machine Learning. *arXiv preprint arXiv:2210.02410* (2022).
- [18] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. 2022. Make-a-scene: Scene-based text-to-image generation with human priors. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XV*. Springer, 89–106.
- [19] Nitesh Goyal, Ian Kivlichan, Rachel Rosen, and Lucy Vasserman. 2022. Is Your Toxicity My Toxicity? Exploring the Impact of Rater Identity on Toxicity Annotation. *arXiv:2205.00501* [cs.HC]
- [20] Priya Goyal, Adriana Romero Soriano, Caner Hazirbas, Levent Sagun, and Nicolas Usunier. 2022. Fairness Indicators for Systematic Assessments of Visual Feature Extractors. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAcCT '22). Association for Computing Machinery, New York, NY, USA, 70–88.
- [21] Laura Gustafson, Megan Richards, Melissa Hall, Caner Hazirbas, Diane Bouchacourt, and Mark Ibrahim. 2023. Pinpointing Why Object Recognition Performance Degrades Across Income Levels and Geographies. *arXiv:2304.05391* [cs.CV]
- [22] Melissa Hall, Candace Ross, Adina Williams, Nicolas Carion, Michal Drozdal, and Adriana Romero Soriano. 2023. DIG In: Evaluating Disparities in Image Generations with Indicators for Geographic Diversity. *arXiv:2308.06198* [cs.CV]
- [23] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718* (2021).
- [24] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2018. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *arXiv:1706.08500* [cs.LG]
- [25] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Nash Equilibrium. *CoRR* abs/1706.08500 (2017). *arXiv:1706.08500* <http://arxiv.org/abs/1706.08500>
- [26] Ben Hutchinson, Negar Rostamzadeh, Christina Greer, Katherine Heller, and Vinodkumar Prabhakaran. 2022. Evaluation Gaps in Machine Learning Practice. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAcCT '22). Association for Computing Machinery, New York, NY, USA, 1859–1876. <https://doi.org/10.1145/3531146.3533233>
- [27] Akshita Jha, Aida Davani, Chandan K. Reddy, Shachi Dave, Vinodkumar Prabhakaran, and Sunipa Dev. 2023. SeeGULL: A Stereotype Benchmark with Broad Geo-Cultural Coverage Leveraging Generative Models. *arXiv:arXiv:2305.11840*
- [28] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. 2023. The Role of ImageNet Classes in Fréchet Inception Distance. *arXiv:2203.06026* [cs.CV]
- [29] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2019. Improved Precision and Recall Metric for Assessing Generative Models. *arXiv:1904.06991* [stat.ML]
- [30] Nayeon Lee, Chani Jung, and Alice Oh. 2023. Hate Speech Classifiers are Culturally Insensitive. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, Sunipa Dev, Vinodkumar Prabhakaran, David Adelani, Dirk Hovy, and Luciana Benotti (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 35–46. <https://doi.org/10.18653/v1/2023.c3nlp-1.5>
- [31] Tony Lee, Michihiro Yasunaga, Chenlin Meng, Yifan Mai, Joon Sung Park, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Benita Teufel, Marco Bellagente, Minguk Kang, Taesung Park, Jure Leskovec, Jun-Yan Zhu, Li Fei-Fei, Jiajun Wu, Stefano Ermon, and Percy Liang. 2023. Holistic Evaluation of Text-To-Image Models. *arXiv:2311.04287* [cs.CV]
- [32] Stanislav Morozov, Andrey Voynov, and Artem Babenko. 2021. On Self-Supervised Image Representations for {GAN} Evaluation. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=NeRdBeTionN>
- [33] WA Munson and Mark B Gardner. 1950. Standardizing auditory tests. *The Journal of the Acoustical Society of America* 22, 5, Supplement (1950), 675–675.
- [34] Muhammad Ferjad Naem, Seong Joon Oh, Youngjung Uh, Yunjeon Choi, and Jaegun Yoo. 2020. Reliable Fidelity and Diversity Metrics for Generative Models. *CoRR* abs/2002.09797 (2020). *arXiv:2002.09797* <https://arxiv.org/abs/2002.09797>
- [35] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2021. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. *CoRR* abs/2112.10741 (2021). *arXiv:2112.10741* <https://arxiv.org/abs/2112.10741>
- [36] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2023. DINOv2: Learning Robust Visual Features without Supervision. *arXiv:2304.07193* [cs.CV]
- [37] Ben Phillips. 2019. The Shifting Border Between Perception and Cognition. *Nous* 53 (06 2019). <https://doi.org/10.1111/nous.12218>
- [38] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. *arXiv:2307.01952* [cs.CV]
- [39] Vinodkumar Prabhakaran, Rida Qadri, and Ben Hutchinson. 2022. Cultural Incongruencies in Artificial Intelligence. *arXiv:arXiv:2211.13069*
- [40] Rida Qadri, Renee Shelby, Cynthia L. Bennett, and Emily Denton. 2023. AI's Regimes of Representation: A Community-centered Study of Text-to-Image Models in South Asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (<conf-loc>, <city>Chicago</city>, <state>IL</state>, <country>USA</country>, </conf-loc>) (FAcCT '23). Association for Computing Machinery, New York, NY, USA, 506–517. <https://doi.org/10.1145/3593013>

- 3594016
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
 - [42] Vikram V. Ramaswamy, Sing Yu Lin, Dora Zhao, Aaron B. Adcock, Laurens van der Maaten, Deepti Ghadiyaram, and Olga Russakovsky. 2023. Beyond web-scraping: Crowd-sourcing a geodiverse dataset. In *arXiv preprint*.
 - [43] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. arXiv:2204.06125 [cs.CV]
 - [44] Suman V. Ravuri and Oriol Vinyals. 2019. Classification Accuracy Score for Conditional Generative Models. *CoRR* abs/1905.10887 (2019). arXiv:1905.10887 <http://arxiv.org/abs/1905.10887>
 - [45] Megan Richards, Polina Kirichenko, Diane Bouchacourt, and Mark Ibrahim. 2023. Does Progress On Object Recognition Benchmarks Improve Real-World Generalization? arXiv:arXiv:2307.13136
 - [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. *CoRR* abs/2112.10752 (2021). arXiv:2112.10752 <https://arxiv.org/abs/2112.10752>
 - [47] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. arXiv:2205.11487 [cs.CV]
 - [48] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lučić, Olivier Bousquet, and Sylvain Gelly. 2018. Assessing Generative Models via Precision and Recall. In *Advances in Neural Information Processing Systems (NeurIPS)*.
 - [49] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training gans. *Advances in neural information processing systems* 29 (2016).
 - [50] Nithya Sambasivan, Erin Arnesen, Ben Hutchinson, Tulsee Doshi, and Vinodkumar Prabhakaran. 2021. Re-Imagining Algorithmic Fairness in India and Beyond. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (Virtual Event, Canada) (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 315–328. <https://doi.org/10.1145/3442188.3445896>
 - [51] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D Sculley. 2017. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. *arXiv preprint arXiv:1711.08536* (2017).
 - [52] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. 2018. How good is my GAN? *CoRR* abs/1807.09499 (2018). arXiv:1807.09499 <http://arxiv.org/abs/1807.09499>
 - [53] Mannat Singh, Laura Gustafson, Aaron Adcock, Vinicius de Freitas Reis, Bugra Gedik, Raj Prateek Kosaraju, Dhruv Mahajan, Ross Girshick, Piotr Dollár, and Laurens van der Maaten. 2022. Revisiting Weakly Supervised Pre-Training of Visual Perception Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 804–814.
 - [54] George Stein, Jesse C. Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Leigh Ross, Valentin Vilecroze, Zhaoyan Liu, Anthony L. Caterini, J. Eric T. Taylor, and Gabriel Loaiza-Ganem. 2023. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. arXiv:2306.04675 [cs.LG]
 - [55] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. Rethinking the Inception Architecture for Computer Vision. *CoRR* abs/1512.00567 (2015). arXiv:1512.00567 <http://arxiv.org/abs/1512.00567>
 - [56] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
 - [57] Sharon Zhou, Mitchell L. Gordon, Ranjay Krishna, Austin Narcomey, Durim Morina, and Michael S. Bernstein. 2019. HYPE: Human eYe Perceptual Evaluation of Generative Models. *CoRR* abs/1904.01121 (2019). arXiv:1904.01121 <http://arxiv.org/abs/1904.01121>

A APPENDIX

We discuss additional details regarding the set-up and results of our analysis in the following sections.

A.1 Additional Set-up Details

Tasks. For Task 1, we construct triplets with real images from GeoDE and generated images from DM w/ CLIP and LDM 2.1. In order to have representation across different types of real and generated images as well as relative distances in the embedding

space, we define different combinations of triplets each object-region as follows:

- A One real image and two generated images, for 85 real images of object O in region R .
 - Reference Image: One real image of O from R .
 - Comparison Images: One image generated with prompt “ O in R ” which is the closest to the reference image in the feature space. and one image generated with prompt “ O in R ” which is the eighth furthest from the reference image.
- B One generated image and two real images, for 85 images generated with prompt “ O in R ”.
 - Reference Image: One image generated with prompt “ O in R ”.
 - Comparison Images: One real image of O from R which is the closest to the reference image in the feature space, and one real image of O from R which is the eighth furthest from the reference image.
- C One real image and two generated images, for 85 real images of object O in region R .
 - Reference Image: One real image of O from R .
 - Comparison Images: One image generated with prompt “ O in R ” which is the closest feature to the reference image, and one image generated with the prompt “ O ” that is the 48th furthest from the reference image.
- D One real image and two generated images, for 85 real images of object O in region R .
 - Reference Image: One real image of O from R .
 - Comparison Images: One image generated with prompt “ O ” which is the sixth feature to the reference image, and one image generated with the prompt “ O in R ” that is the eighth furthest from the reference image.
- E One generated image and two real images, for 85 images generated with prompt “ O ”.
 - Reference Image: One image generated with prompt “ O ”.
 - Comparison Images: One real image of O from R which is the closest to the reference image in the feature space, and one real image of O from R which is the eighth furthest from the reference image.

To select reference images, we pick the 60 images for each object-region combination that has another image nearest to it and the 25 images that has its nearest image furthest away. This allows us to capture statistics about both images that are relatively similar to the others (the former) and images that are quite different (the latter).

For our Visual Appealing analysis, we use Combo A. This allows a comparison between real images of a given object taken in a specific region and generated images prompted with the same object and region. In addition, for the Visual Appeal analysis we select only triplets where exactly one image falls into the GeoDE manifold.

Annotators. Potential annotators were first recruited according to their geographic location. They were provided an initial task to complete to demonstrate general competency then went through a training program consisting of written instructions for each task and practice examples. Because our Tasks require an amount of subjectivity throughout them yet we still want to ensure annotators have a clear understanding of the intended task, we created a trial

task that involved annotators identifying similar images of cars. Annotators who passed with a 75% rate were then graduated to the production tasks and paid a standard rate per hour based on a country-specific living wage.

For production tasks, annotators within a given location worked through queues for Task 1 and Task 2 for each object and image region. They varied in the proportion of jobs completed, with some annotators completing less than 10 jobs and others hundreds jobs. Each unique job was annotated by one annotator from each region.

Following the completion of the production tasks, we performed an analysis of annotation completion quality. We use annotations of whether the real GeoDE images contain a given object as our “sanity” check, as we have a ground truth answer for these images. We found that while the accuracy for real GeoDE images was quite high for most annotators, a very small number of annotators had low accuracy for identifying objects in GeoDE that were included as comparison image in the task construction. In those cases, we believe the annotators did not fully understand the task instructions to select *all* images showing the object or were operationalizing the task incorrectly. Thus, we introduce a requirement that annotators have at least 90% accuracy in identifying objects in GeoDE images, to ensure some conformity to definition of object representation and task completion. For each object-region combination, we remove annotations from annotators who did not meet this bar.

At the conclusion of filtering, we had annotations from 20 annotators located in Africa (split across Egypt, Nigeria, and South Africa), 17 annotators in Europe (located in Great Britain, Italy, Romania and Spain) and 23 annotators in Southeast Asia (located in Indonesia, the Philippines, and Thailand). All annotators were fluent in English, and tasks were presented in English.

A.2 Additional annotation details

Table 3 contains descriptive codes used to categorize responses from annotators. Figure 9 shows proportion of survey respondents exhibiting key themes in their response to questions about their interpretation of the annotation task.

A.3 Additional Results

Visual appeal. The top of Figure 12 shows that when annotators pick the comparison image that is closer to the reference image as the more similar image, all annotators differentiate between visual appeal and similarity *less* as comparison images move away from each other in InceptionV3 space (solid line) and *slightly more* as images move further from the reference image (dashed line). Furthermore, annotators in Europe tend to differentiate between similar and visually appealing images *more* than annotators in Africa, and both more than annotators in Southeast Asia. This is consistent even as distances between comparison and reference images vary.

In addition, the bottom of Figure 12 shows that when annotators pick the further image as the more similar image, all annotators differentiate *more* between similarly and visually appealing than when they were picking the closer image as the similar image. However, this differentiation seems less impacted by the distance of comparison and reference images, and there is less between-region variation.




Figure 9: (a) Proportion of survey respondents exhibiting key themes in their responses to questions about identifying locations from both background context and object specifics. (b) Proportion of survey respondents suggesting they utilized common stereotypes when annotating geographic representativeness. Annotators in all locations frequently exhibit stereotypes in their responses.



Figure 10: (a-b) Examples of generated images of dogs: (b) depicts images selected as more and less appealing (respectively) by all three annotators; (c) depicts images with inconsistent appeal annotations. Annotators may leverage qualities for appeal beyond similarity more when they have greater familiarity with what is in the image.

Code	Description	Kappa
External search	Using an external reference, e.g. search engine	0.91
People	Mentions people or faces	0.80
Nature & natural world	Mentions nature, plants, animals, landscape, weather or climate	0.77
Feature-based representation	Mentions looking for feature-based representations, e.g. looking at color, size or shape	0.76
Unclassified	Unclassified response	0.73
Stereotyped representations	Utilizes common stereotypes of appearance of items within a region	0.72
Built environment	Mentions buildings, architecture, or the built environment	0.71
Absence of detail	Mentions an absence detail or noteworthy characteristics	0.70
Culture, art & religion	Mentions culture, food, art, sculpture, religion, fashion and clothing	0.63
Personal lived experience	Leverages experience living in a location	0.60
Media portrayals	Reliance on portrayals in media	0.60
Brand identification	Uses a known reference, such as recognising a brand or store	0.52
Overgeneralised representations	Considers subsets of a region indicative of a broader entire region	0.48
Objects and tools	Uses everyday objects such as tools or appliances	0.41
Erroneous representation	Mentions erroneous or unrealistic object or background representation	0.41
Uniqueness	Expectation that a representation be unique to a region	0.15

Table 3: Descriptive codes used to categorize survey responses from annotators. Codes are sorted by inter-annotator agreement Fleiss’ Kappa, and only codes with Kappa greater than 0.6 (horizontal rule) were included in our analysis.



(a) Examples where *more* appealing generated image falls in GeoDE manifold

(b) Examples where *less* appealing generated image falls in GeoDE manifold

Figure 11: (a-b) Example pairs where the *more* or *less* appealing image falls in the manifold of real images. Annotators often consider realistic images more appealing, aligning with the real manifold, but also interpret appeal in contradictory ways, e.g., prioritizing background aesthetics, stereotypes, or “niceness” of objects.



Figure 12: When annotators pick the further reference image as the more similar image, all annotators differentiate *more* between similarly and visually appealing than when they were picking the closer image as the similar image.