

ARTWHISPERER: A DATASET FOR CHARACTERIZING HUMAN-AI INTERACTIONS IN ARTISTIC CREATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

As generative AI becomes more prevalent, it is important to study how human users interact with such models. In this work, we investigate how people use text-to-image models to generate desired target images. To study this interaction, we created ArtWhisperer, an online game where users are given a target image and are tasked with iteratively finding a prompt that creates a similar-looking image as the target. Through this game, we recorded over 50,000 human-AI interactions; each interaction corresponds to one text prompt created by a user and the corresponding generated image. The majority of these are repeated interactions where a user iterates to find the best prompt for their target image, making this a unique sequential dataset for studying human-AI collaborations. In an initial analysis of this dataset, we identify several characteristics of prompt interactions and user strategies. People submit diverse prompts and are able to discover a variety of text descriptions that generate similar images. Interestingly, prompt diversity does not decrease as users find better prompts. We further propose a new metric to quantify the *steerability* of AI using our dataset. We define steerability as the expected number of interactions required to adequately complete a task. We estimate this value by fitting a Markov chain for each target task and calculating the expected time to reach an adequate score in the Markov chain. We quantify and compare AI steerability across different types of target images and two different models, finding that images of cities and natural world images are more steerable than artistic and fantasy images. These findings provide insights into human-AI interaction behavior, present a concrete method of assessing AI steerability, and demonstrate the general utility of the ArtWhisperer dataset.

1 INTRODUCTION

Direct human interaction with AI models has become widespread following a number of technical innovations improving the quality of text-to-text Brown et al. (2020); Ouyang et al. (2022); Anil et al. (2023) and text-to-image models Rombach et al. (2022a); Ramesh et al. (2022), enabling the public release of high-quality AI-based services like ChatGPT chatGPT, Bard Bard, and Midjourney Midjourney. These models have seen rapid interest and adoption largely due to the ability of the general public to interact with and steer the AI in diverse contexts including engineering, creative writing, art, education, medicine, and law Dakhel et al. (2023); Nguyen & Nadi (2022); Ippolito et al. (2022); Cetinic & She (2022); Qadir (2023); Cascella et al. (2023); Sloan (2023).

A key challenge in developing these models is aligning their output to human inputs. This is made challenging by the broad domain of use cases as well as the diverse prompting styles of different users. Many approaches can be categorized as “prompt engineering,” where specific strategies for prompting are used to steer a model Oppenlaender (2022); Liu & Chilton (2022); Zhou et al. (2022b); Wei et al. (2022); White et al. (2023). Great success has also been found by fine-tuning models using relatively small datasets to follow human instructions Ouyang et al. (2022), respond in a specific style Hu et al. (2021), or behave differently to specified prompts Zhou et al. (2022a); Gal et al. (2022).

In this work, we take interest in the fact that human interaction with these models is often an iterative process. We develop a dataset to study this interaction. The dataset is collected through an interactive game we created where players try to find an optimal prompt for a given task (see Figure 1). In particular, we focus on text-to-image models and ask the player to generate a similar image (AI

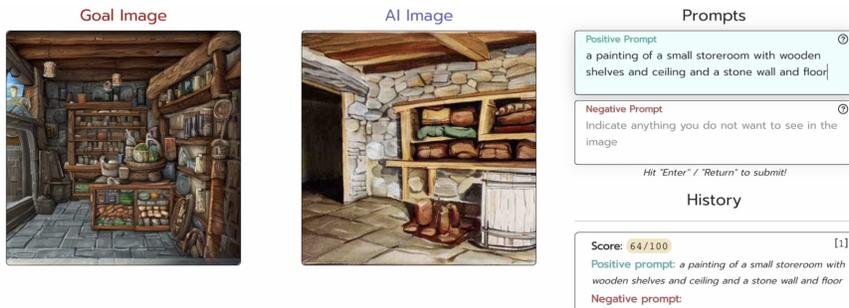


Figure 1: Interface of the ArtWhisperer game. Prompts entered on right. Target (goal) image and player-generated image on left. Previous prompts and scores are displayed in the lower right.

Image) to a given *target image*. The player is allowed to iterate on their prompt, using the previously generated image(s) as feedback to help them adjust their prompt. A score is also provided as feedback to help the user calibrate how “close” they are to a similar image.

Using this setup, we collected data on 51,026 interactions from 2,250 players across 191 unique target images. The target images were selected from a diverse set of AI-generated and natural images. We also collected a separate dataset of 4,572 interactions, 140 users, and 51 unique target images in a more controlled setting to assess the robustness of our findings.

Based on this data, we find several interesting patterns in how people interact with AI models. Players discover a diverse set of prompts that all result in images similar to the target. To discover these prompts, players typically make small, iterative updates to their prompts. Each update improves their image with a moderate success rate (40 – 60% for most target images). Based on these findings, we define and evaluate a metric for model steerability using the stopping time of an empirical Markov model. We use this metric to assess steerability across image categories and across two AI models.

Ethical considerations One of the main goals of this work is to help improve the quality of human-AI interaction. Our dataset and findings provide quantitative insights on how people interact with generative AI and can potentially be used to design AI that are easier for people to use. It does not address the broader concern that bad actors may abuse generative AI models.

Our contributions We release a public dataset on human interactions with an AI model. To our knowledge, this is the first such dataset showing repeated interactions of people with a text-to-image model to accomplish specified tasks. We also provide an initial analysis of this data and propose a simple-to-calculate metric for assessing model steerability. Our dataset and associated code is made available at [link redacted for submission].

Related Works Human-AI interaction datasets for text-to-text and text-to-image models typically focus on single interactions and generally do not provide users with a specific task. Public text-to-image interaction datasets typically contain the generated AI images and/or prompts Santana (2022); Wang et al. (2022) and optionally some form of human preference rating Pressman et al. (2022); Wu et al. (2023); Xu et al. (2023); Kirstain et al. (2023). These datasets generally rely on scraping online repositories like Lexica Lexica or Discord servers focused on AI art. Though some of these datasets include metadata that may allow for reconstruction of prompt iteration, there is no guarantee the user has the same desired output in mind over the iteration. Public text-to-text interaction datasets are much more limited as the best performing models are generally accessible only through APIs with no public user interaction datasets. While some researchers have investigated how human-AI interaction for text-to-text can be improved through various tools Wu et al. (2022a;b), the amount of data collected is limited and not publicly available. There are also repositories containing prompt strategies for various tasks Bach et al. (2022) but no human interaction component.

We seek to rectify two of the shortcomings of the existing datasets—namely, that they do not contain extended interactions as the user attempts to steer the AI, and they do not have a predefined goal. In our work, we create a controlled environment where we allow extended interactions and have a known goal for human users. As shown by our initial analysis, our dataset may enable deeper understanding of user prompting strategies and assessing model steerability.

2 INTERACTION GAME

In the game, players are shown a target image. A few example target images are provided in Figures 2, 4. Players are also given a limited interface to a text-to-image model, Stable Diffusion (SD) v2.1 model Rombach et al. (2022b). In particular, players can enter a “positive prompt” (describes the desired content) and a “negative prompt” (describes what should be omitted) to steer the AI model. All models hyperparameters are fixed. Upon inputting a prompt, the player is shown the image generated by the AI model, along with a similarity score between their generated image and the target image. The interface is shown in Figure 1.

2.1 HOW TARGET IMAGES ARE SELECTED

We randomly sample target images from two sources. The first is a collection of Wikipedia pages, and the second is a dataset of prompts AI artists have used with SD Santana (2022). In addition to sampling target images, we need to ensure the task is feasible to users. We do not allow users to adjust the seed or other parameters of the model, so we need to ensure the selected model parameters can generate reasonably similar images to the target image. We find that selecting an appropriate random seed is sufficient, and fix all other model parameters (see Appendix A.4 for details and discussion).

Wikipedia Images A collection of 35 Wikipedia pages on various topics including art, nature, cities, and various people. A full list of pages sampled from is provided in Appendix 2. From these pages, we scraped 670 figures licensed under the Creative Commons license. These figures were then filtered by which had captions, as well as which images were JPG or PNG images (i.e., not animated, and not PDF files), resulting in 557 images.

For each of the 557 images, we first resize and crop the image to size 512×512 . The Wikipedia caption is used as the ground truth “prompt”. Let the image-caption pair be denoted as (t_i, p_i^*) . We sample the model on 50 random seeds, with p_i^* as the prompt input. This generates a set of 50 images: $S_i = \{(x_i, s_i) : i = 1, \dots, 50\}$ for generated image x_i and seed s_i . Let $C(x)$ denote the CLIP image embedding Radford et al. (2021) of an image x . Then we select the seed as s_{i^*} , where

$$i^* := \min_{i=1, \dots, 50} \left\| \frac{C(x_i)}{\|C(x_i)\|_2} - \frac{C(t_i)}{\|C(t_i)\|_2} \right\|_2$$

Here, s_{i^*} is selected to minimize the distance to the target image given the target prompt.

AI-Generated Images A collection of 2,000 AI-art prompts are randomly sampled from the Stable Diffusion Prompts dataset Santana (2022). For each prompt, p_i^* , we generate two sets of images. As before, we use 50 unique random seeds to select the seed, s_{i^*} and an additional 10 random seeds to use for selecting the generated target image (so in total, we use 60 unique random seeds): the first set, $S_{i,1} = \{(x_{i,1}, s_i) : i = 1, \dots, 10\}$ and $S_{i,2} = \{(x_{i,2}, s_i) : i = 1, \dots, 50\}$. We select the target image, $t_{i_1^*}$, from $S_{i,1}$:

$$i_1^* := \min_{i=1, \dots, 10} \text{median} \left(\left\{ \left\| \frac{C(x_{i,1})}{\|C(x_{i,1})\|_2} - \frac{C(x_{j,2})}{\|C(x_{j,2})\|_2} \right\|_2 : j = 1, \dots, 50 \right\} \right)$$

We select the random seed, $s_{i_2^*}$, using $t_{i_1^*}$ and $S_{i,2}$, with

$$i_2^* := \min_{i=1, \dots, 50} \left\| \frac{C(x_{i,2})}{\|C(x_{i,2})\|_2} - \frac{C(t_{i_1^*})}{\|C(t_{i_1^*})\|_2} \right\|_2$$

Here, $t_{i_1^*}$ is chosen to be more representative of the types of images we may expect given the fixed prompt, p_i^* . This is because $t_{i_1^*}$ is selected to be close to the center of the sampled images, $S_{i,2}$. The intuition for selecting $s_{i_2^*}$ is the same as selecting s_{i^*} for the Wikipedia images.

2.2 SCORING FUNCTION

To provide feedback to players, we created a scoring function to assess the similarity of a player’s generated image and the target image. We define the scoring function as

$$\text{score}(x, t) = \max(0, \min(100, \alpha \cdot \left\| \frac{C(x)}{\|C(x)\|_2} - \frac{C(t)}{\|C(t)\|_2} \right\|_2 + \beta)),$$

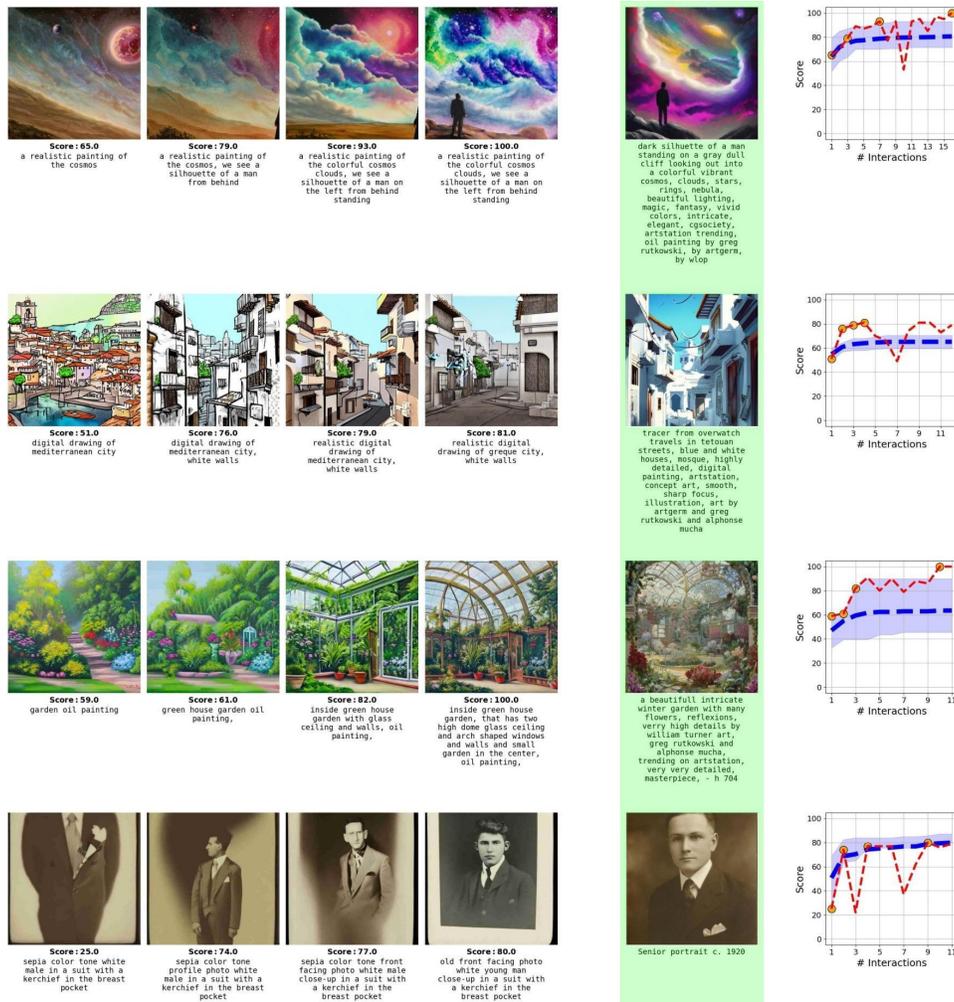


Figure 2: Example user trajectories. In each row, the first 4 images show the prompt progression for a given user. Target image shown in green column. Plot shows the average score trajectory across all users for this target image (blue) and the specific user’s full score trajectory (red). Orange circles indicate the displayed images.

for generated image x , target image t , and constants α, β . Note the range of $score(x, t)$ is integers in the interval $[0, 100]$. Details on how α, β are selected parameters are provided in Appendix A.3.

While this scoring function is often reasonable, it does not always align with the opinions of a human user. To assess how well $score(x, t)$ follows a user’s preferences, we acquire ratings from a subset of users (see *ArtWhisperer-Validation* in Section 2.3). We find $score(x, t)$ has a Pearson correlation coefficient of 0.579 indicating reasonable agreement. Further assessment is performed in Section 4.3 and discussed at length in Appendix A.13.

2.3 DATASET OVERVIEW

We collected two datasets: *ArtWhisperer* and *ArtWhisperer-Validation*. We use *ArtWhisperer* for most analysis and results; for some of the results in Sections 2.2, 4.2, and 4.3, we also use *ArtWhisperer-Validation* (when referenced). Data was collected from March-May 2023. IRB approval was obtained.

ArtWhisperer: A public version of our game was released online, with three new target images released daily. We collected data from consenting users playing the game. Users were not paid. Users were anonymous and we only collected data related to the prompts submitted to ensure privacy of

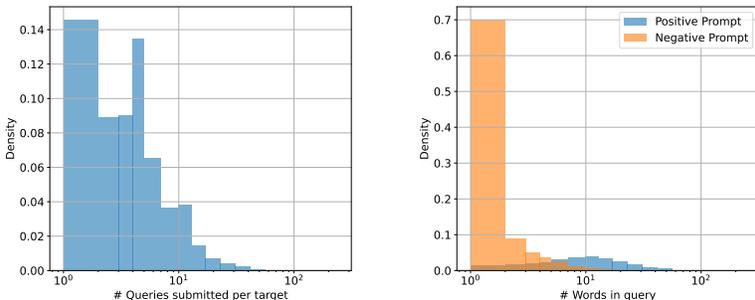


Figure 3: Left, Distribution of # of user queries per target image. The average number of queries per image is 9.18. Right, Distribution of the # of words submitted in a query. The average number of words submitted in a positive and negative prompt are 20.02 and 2.32 respectively.

users. While we expect some users played the game across multiple days, we did not track them. A summary of the *ArtWhisperer* dataset is provided in Table 1. In total, we have 2,250 (potentially non-unique) players corresponding to 51,026 interactions across 191 target images. Players interacted with the model SD v2.1. In Figure 3, we plot the number of queries submitted by players across different target images.

***ArtWhisperer-Validation*:** The game (with a near identical interface) was also released as a controlled user study to paid crowd workers on Prolific Academic. The crowd workers were compensated at a rate of \$12.00 per hour for roughly 20 minutes of their time. Workers played the game across 5 randomly selected target images from a pre-selected subset of 51 target images chosen to have diverse content. Workers were also asked to rate each of their images on a scale of 1-10 (i.e., self-scoring their generated images). In total, we collected data on 4,572 interactions, corresponding to 140 users and 51 unique target images across two different diffusion models, SD v2.1 and SD v1.5. Additional details and demographic information are provided in Appendix A.6.

Table 1: *ArtWhisperer* Dataset Overview. Each row contains summary data for a different subset of the dataset. Subsets may overlap. Similar information for *ArtWhisperer-Validation* is in Appendix A.6.

# Players	# Target Images	# Interactions	Average # Prompts	Average Score	Median Duration	Category
2250	191	51026	9.29	58.93	18 s	Total
377	25	3884	8.65	56.70	19 s	Contains famous person?
353	32	3785	8.26	61.64	21 s	Contains famous landmark?
2005	140	40290	9.24	59.83	18 s	Contains man-made content?
1177	58	18255	10.93	57.21	17 s	Contains people?
344	77	6972	8.81	62.01	20 s	Is real image?
2140	103	43524	9.42	58.37	17 s	Is AI image?
1483	82	24913	9.14	59.45	17 s	Is art?
623	29	7297	9.14	53.77	18 s	Contains nature?
160	14	1355	7.28	65.74	19 s	Contains city?
1239	39	15872	9.91	56.74	16 s	Is fantasy?
618	19	8359	10.51	57.88	17 s	Is sci-fi or space?

3 PROMPT DIVERSITY

We quantify prompt diversity by looking at the distribution of prompts in the text embedding space. In particular, we use the CLIP text embedding Radford et al. (2021), though we do find the choice of embedding is not particularly important for our results (see Appendix A.7).

3.1 DIVERSE PROMPTS USED FOR HIGH SCORES

People achieve high scores with a diverse set of prompts. It is not surprising that this is possible (i.e., that the score metric has multiple local maxima), but it is potentially surprising that users consistently find these differing local maxima. Examples are shown in the four leftmost columns of Figure 4.

We quantify this finding in Figure 5, where we plot two metrics defined as follows. Let z_0, z_n be normalized embeddings of the initial and best prompt/image found by a user. Let z^* be the normalized embedding of the target prompt/image. We define the difference in embedding distance to ground truth as $\|z_n - z^*\|_2 - \|z_0 - z^*\|_2$. In blue, we use the CLIP text embeddings of the prompts; in orange, we use the CLIP image embeddings of the generated images. We note two findings here: (1) the metric applied to the image embeddings is guaranteed to be non-positive as the embedding distance is monotonically decreasing with the score, and (2) the metric applied to the text embeddings is apparently symmetric around 0, indicating that unlike the image embedding, distance in the text embedding space does not monotonically decrease with score. Together, these findings illustrate that users tend to discover diverse prompts and *do not converge in their prompt design*.

The right two columns of Figure 4 provide another visualization at an individual level. Here, we plot a UMAP McInnes et al. (2018) projection for the CLIP image and text embeddings of the target image and a sample of the submitted prompts (and corresponding generated images). The target embedding is in orange, the first prompt embedding is in red, and the last (best scoring) prompt embedding is in green. The arrows connect a given user’s first and last prompt. We see that a decrease in distance in the image embedding space (which is inversely correlated with the user’s score) does not always correspond to a decrease distance in the text embedding space.

Additionally, we find the distribution of prompts does not converge. In the left of Figure 6, we plot the distribution of distances between the first prompt (in blue) and the last prompt (in orange) to the average prompt for the corresponding target image. Despite the average score improving from 51.9 to 70.3 (out of 100) indicating a significant improvement in score, prompt diversity does not significantly diminish. That is, users do not converge to similar prompts to achieve high scores. Similar analysis of the image embedding space suggests image diversity *decreases* (Figure 5).

3.2 PEOPLE SUBMIT SIMILAR PROMPTS THROUGHOUT THEIR INTERACTION

In the center of Figure 6, we plot the distribution of the standard deviation of prompts for users (blue) and for permuted users (orange). Permuted users are generated by sampling from all prompts for a given target image uniformly, using the same distribution of number of prompts as for real users. The gap between the two distributions shows that individuals do not randomly sample prompts each interaction, but base new prompts off of previously submitted prompts (p-value $< 10^{-10}$, t-test for independent variables). An analysis of how scores change between adjacent prompts shows that this strategy has a moderate success rate and improves the score 40 – 60% of the time, with an average rate of 48.6% (note that score changes < 1 are counted as unchanged; this occurs 10.2% of the time).

While this is not a surprising result (that users often do not make significant changes to their prompt), but it is an important result to understand how typical users interact with AI models. Moreover, it suggests that user initialization (i.e., the first prompt they submit) is critical.

3.3 PEOPLE HAVE SIMILAR PROMPT STYLES ACROSS IMAGES

We quantify user style by computing the difference (in the CLIP text embedding space) between the average prompt of a given user and the average prompt across all users for a given target image. To quantify style variation for a user, we then compute the standard deviation of the user style across the target images the user generated. In the right of Figure 6, we plot the distribution of user style variation for real users (blue) and permuted users (orange). Permuted users are generated by randomly sampling user styles. This allows us to test whether users have a consistent prompting style. We find users do indeed have specific styles of prompting (p-value $< 10^{-10}$, t-test for independent variables). However, the difference is not seemingly not large, suggesting that while user style may a component to prompting, other factors related to the target image may be more important.

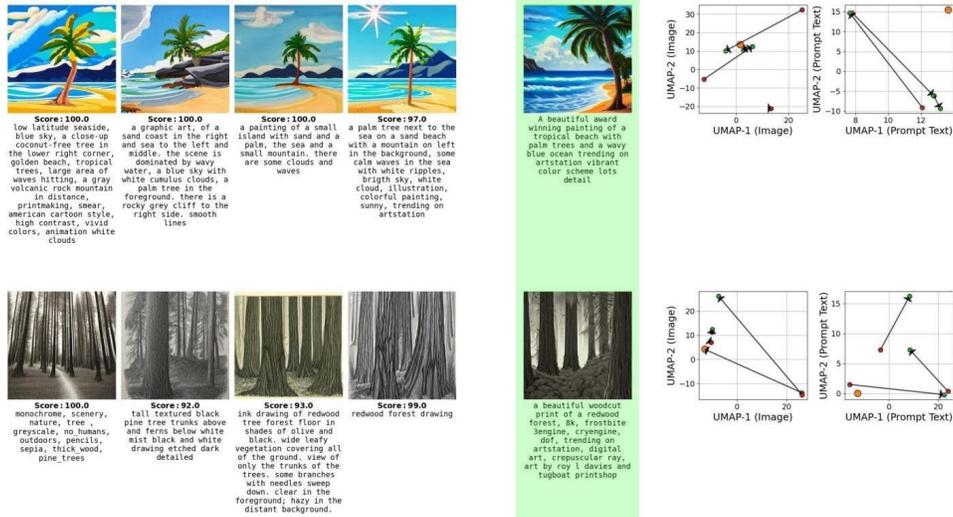


Figure 4: Left: Diverse, high-scoring prompt submissions from different player. Target image in green column. Right: Image (left) and text (right) embeddings of displayed images (target in orange; user submissions in green (best) and red (first)) using the UMAP McInnes et al. (2018) projection.

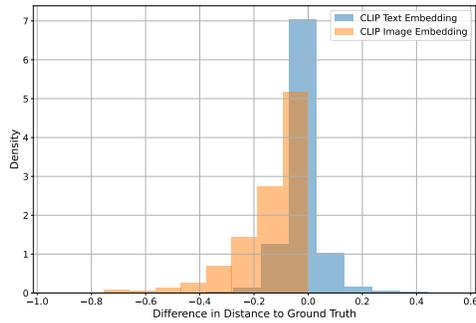


Figure 5: Difference of distance from the first prompt to ground truth and distance from the last (best) prompt to ground truth for CLIP text (blue) and CLIP image embeddings (orange).

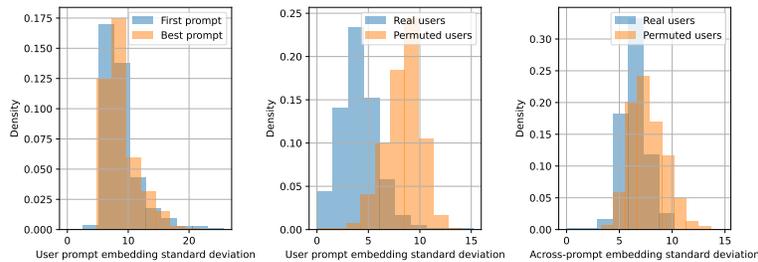


Figure 6: Left: Users submit diverse (across users) prompts, both at beginning and end of interaction. Center: Individual users do not submit diverse prompts. Right: Users have different prompting styles.

4 MODEL STEERABILITY

Model steerability refers to the ability of a user to steer a model towards a desired outcome. There is no current consensus on how to measure AI steerability. A common approach is to simply measure performance of a model on standardized dataset evaluations Jahanian et al. (2019); OpenAI (2023). While this can enable comparisons between tasks and models, this approach does not allow for the feedback loop present when humans interact with a model. Steerability can also be measured qualitatively based on user assessment of their experience interacting with the AI Chung et al. (2022). We create a simple yet informative measure of model steerability. We then analyze this measure across different subgroups of images and across two different Stable Diffusion models—SDv2.1 and the older SDv1.5 Rombach et al. (2022a).

4.1 MEASURING STEERABILITY

As discussed in Section 3.2, users typically engage with the model through clusters of similar prompts. They typically start with an initial base prompt and proceed to make multiple incremental modifications to it. We use this observation as a basis for creating a steerability metric. We define a Markov chain between scores. Each node is a score with edges connecting to the subsequent score. To make this tractable for empirical analysis, we bin scores into five groups: $[0, 20]$, $[21, 40]$, $[41, 60]$, $[61, 80]$, $[81, 100]$. We use the expected time taken to reach the last score bin, $[81, 100]$, as our steerability score (i.e., the stopping time to reach an adequate score).

For each target image, we calculate the empirical transition probability matrix between binned scores using all the players’ data for that image. We then calculate the *steerability score* for the given target image by running a Monte Carlo simulation to estimate stopping time, as defined above. To assess steerability across a group of images, we average steerability score across all images in the group.

4.2 ANALYSIS

In Figure 7, we plot the steerability score across image groups. Error bars show the standard error. For examples of steerability scores for individual images, see Appendix A.11. We find that images containing famous people or landmarks, real images (not AI generated), contain cities, or contain nature are the most steerable. AI-generated images, fantasy images, and images of human art are the least steerable. There are a few possible explanations. The model we are assessing here, SDv2.1, as well as its text encoder OpenCLIP, are trained on subsets of LAION5B Schuhmann et al. (2022). The contents of LAION5B are predominantly real world images, indicating why these images may be more steerable (i.e., text describing these types of images may have a better encoding). Moreover, the prompts for AI-generated images and fantasy images generally include specific internet artists and/or art styles which may not be known to most users making achieving the desired target image more difficult. Another potential reason is the distribution of images chosen for each category. Clearly, there are “easier” and “more difficult” images in each category; part of the reason for smaller stopping time may be the sample of images chosen rather than the actual image category.

Using the *ArtWhisperer-Validation* data, we also compare steerability across two models: SDv2.1 and SDv1.5. Across most image categories, we observe a similar steerability. Images of nature, sci-fi or space, and real images have the largest differences in steerability between the two models; SDv2.1 is more steerable in all three cases. This suggests that SDv2.1 may be more steerable for natural images as well as sci-fi images, and is similarly steerable for other kinds of images including AI-generated artwork. One explanation may be that most of our users were not aware of certain prompting strategies that help models generate more aesthetic images or certain art styles; it is possible that for experienced users, AI art images may be more steerable, and differences between models may be magnified if, for example, a user is experienced working with one particular model. More discussion is provided in Appendix A.10.

4.3 JUSTIFICATION FOR AUTOMATED SCORE

One limitation of our steerability metric comes from the method of scoring user-submitted prompts. Ideally, we would assess steerability with a user’s personal preferences. As mentioned in Section 2.2, the scores and human ratings have a positive correlation. Here, we use the human ratings from

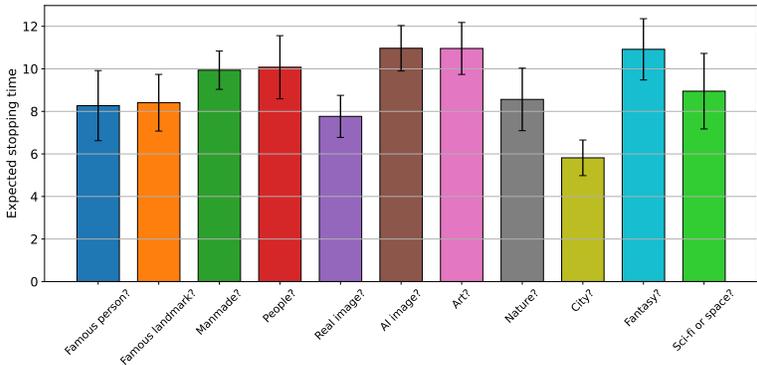


Figure 7: Steerability across image groups (smaller indicates increased steerability). Bars show average expected stopping time across images in the image group; error bars show standard error.

ArtWhisperer-Validation instead of our score function to assess steerability. We compute the steerability score across both models and across image groups. Generally, the steerability scores change little. In all but two cases (SDv2.1 on sci-fi and space images; SDv1.5 on nature images), the human rating-based steerability score remains within a 95% confidence interval of the score-based steerability score. While our score function may not perfectly capture human preferences, the steerability score we generate appears to be robust to these issues. Further discussion is included in Appendix A.13.

5 DISCUSSION

As demonstrated in our analysis, the *ArtWhisperer* and *ArtWhisperer-Validation* datasets can provide insights into user prompting strategies and enables us to assess model steerability for individual tasks and groups of tasks. What makes our dataset particularly useful is the controlled interactive environment, where users work toward a fixed goal, that we capture data in.

One of the most exciting use cases we see for our dataset is to create synthetic humans for prompt generation. For example, similar to the method described in Promptist Hao et al. (2022), we imagine fine-tuning a large language model with our dataset to generate prompt trajectories (i.e., rather than an optimized prompt) using similar exploration strategies as a human prompter. These synthetic prompters could be based on multimodal models like OpenFlamingo Awadalla et al. (2023) or text-only models and use score-feedback to condition the trajectory generation. As an initial proof-of-concept, we fine-tuned a MT0-large model Muennighoff et al. (2022) model on our dataset and found the fine-tuned model can indeed behave similarly to human users (see Appendix A.14). These synthetic prompters have several potential use cases:

1. Automating measurement of text-to-image model steerability by using synthetic users in place of real human prompters. While we believe our proposed steerability metric is effective, its main limitation currently is the requirement for human annotations.
2. Incorporating steerability in the objective function for text-to-image models. By representing steerability as a function of synthetic users, it becomes possible to explicitly optimize a model for steerability.
3. Generating human readable image captions that are compatible with a Stable Diffusion model by using the synthetic prompter to optimize the token representation of the prompt.

Additionally, our dataset can be used for further analysis on human prompting strategies beyond what we discussed in the paper. For example, one question we only touched upon is whether we can compare human prompters to automated prompt optimization methods (e.g., do humans behave similar to some gradient-based optimization approach in the prompt embedding space?). There are also potential uses for crafting better image similarity metrics using the human ratings we collected.

REFERENCES

- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. Palm 2 technical report, 2023.
- Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, Jenia Jitsev, Simon Kornblith, Pang Wei Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*, 2023.
- Stephen H Bach, Victor Sanh, Zheng-Xin Yong, Albert Webson, Colin Raffel, Nihal V Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, et al. Promptsources: An integrated development environment and repository for natural language prompts. *arXiv preprint arXiv:2202.01279*, 2022.
- Bard. <https://http://bard.google.com>, 2023.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Marco Cascella, Jonathan Montomoli, Valentina Bellini, and Elena Bignami. Evaluating the feasibility of chatgpt in healthcare: an analysis of multiple clinical and research scenarios. *Journal of Medical Systems*, 47(1):33, 2023.
- Eva Cetinic and James She. Understanding and creating art with ai: review and outlook. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2):1–22, 2022.
- chatGPT. <https://chat.openai.com>, 2023.
- John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. Talebrush: sketching stories with generative pretrained language models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–19, 2022.
- Arghavan Moradi Dakhel, Vahid Majdinasab, Amin Nikanjam, Foutse Khomh, Michel C Desmarais, and Zhen Ming Jack Jiang. Github copilot ai pair programmer: Asset or liability? *Journal of Systems and Software*, 203:111734, 2023.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

- 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 preprint arXiv:2208.01618*, 2022.
- Yaru Hao, Zewen Chi, Li Dong, and Furu Wei. Optimizing prompts for text-to-image generation. *arXiv preprint arXiv:2212.09611*, 2022.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Daphne Ippolito, Ann Yuan, Andy Coenen, and Sehmon Burnam. Creative writing with an ai-powered writing assistant: Perspectives from professional writers. *arXiv preprint arXiv:2211.05030*, 2022.
- Ali Jahanian, Lucy Chai, and Phillip Isola. On the "steerability" of generative adversarial networks. *arXiv preprint arXiv:1907.07171*, 2019.
- Apoorv Khandelwal, Luca Weihs, Roozbeh Mottaghi, and Aniruddha Kembhavi. Simple but effective: Clip embeddings for embodied ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14829–14838, 2022.
- Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *arXiv preprint arXiv:2305.01569*, 2023.
- Lexica. <https://lexica.art/>, 2023.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.
- Vivian Liu and Lydia B Chilton. Design guidelines for prompt engineering text-to-image generative models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–23, 2022.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927*, 2022.
- Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Midjourney. <https://www.midjourney.com/home/?callbackUrl=%2Fapp%2F>, 2023.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. Crosslingual generalization through multitask finetuning. *arXiv preprint arXiv:2211.01786*, 2022.
- Nhan Nguyen and Sarah Nadi. An empirical evaluation of github copilot’s code suggestions. In *Proceedings of the 19th International Conference on Mining Software Repositories*, pp. 1–5, 2022.
- OpenAI. Gpt-4 technical report, 2023.
- Jonas Oppenlaender. Prompt engineering for text-based generative art. *arXiv preprint arXiv:2204.13988*, 2022.

- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- John David Pressman, Katherine Crowson, and Simulacra Captions Contributors. Simulacra aesthetic captions. Technical Report Version 1.0, Stability AI, 2022. url <https://github.com/JD-P/simulacra-aesthetic-captions>.
- Prolific Academic. <https://www.prolific.co>, 2023.
- Junaid Qadir. Engineering education in the era of chatgpt: Promise and pitfalls of generative ai for education. In *2023 IEEE Global Engineering Education Conference (EDUCON)*, pp. 1–9. IEEE, 2023.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- 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.
- 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 (CVPR)*, pp. 10684–10695, June 2022a.
- 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*, pp. 10684–10695, 2022b.
- Gustavo Santana. Stable-diffusion-prompts. Huggingface Datasets, 2022. URL <https://huggingface.co/datasets/Gustavosta/Stable-Diffusion-Prompts>.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. URL <https://openreview.net/forum?id=M3Y74vmsMcY>.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Karen Sloan. A lawyer used chatgpt to cite bogus cases. what are the ethics? *Reuters*, May 2023. URL <https://www.reuters.com/legal/transactional/lawyer-used-chatgpt-cite-bogus-cases-what-are-ethics-2023-05-30>.
- 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. *arXiv preprint arXiv:2210.14896*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022.
- Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*, 2023.

- Tongshuang Wu, Ellen Jiang, Aaron Donsbach, Jeff Gray, Alejandra Molina, Michael Terry, and Carrie J Cai. Promptchainer: Chaining large language model prompts through visual programming. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pp. 1–10, 2022a.
- Tongshuang Wu, Michael Terry, and Carrie Jun Cai. Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–22, 2022b.
- Xiaoshi Wu, Keqiang Sun, Feng Zhu, Rui Zhao, and Hongsheng Li. Better aligning text-to-image models with human preference. *ArXiv*, abs/2303.14420, 2023.
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation, 2023.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022a.
- Yongchao Zhou, Andrei Ioan Muresanu, Ziwon Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers. *arXiv preprint arXiv:2211.01910*, 2022b.

A APPENDIX

A.1 DATASET LIMITATIONS

A potential limitation in our dataset is the diversity of unique images. However, as shown in our analysis (Sections 3, 4.2), we have multiple illuminating findings despite having <200 unique images in our dataset. In collecting this dataset, there was a tradeoff – given a fixed budget to collect data, we could choose to collect more unique images or collect more user interactions per image. We opted to collect data on a fewer number of unique images with more users interacting with each image, as we believed this data would contain more insights into human interaction with the AI.

A.2 INFORMATION ON WIKIPEDIA PAGES SCRAPED

Table 2 presents a list of the Wikipedia pages used to select real-world target images (see Section 2.1). We extracted images from each listed Wikipedia page. We then uniformly sample a category and subsequently sample an image from a page in that category. This ensures a diverse set of images, which is important given that some of the Wikipedia pages contain many more images than others (e.g., Paris has 10 times more usable images than Social_documentary_photography).

Table 2: Wikipedia images used.

Category	Wikipedia Pages
Art	Art, Fine_art, Fine-art_photography, History_of_art, Painting
Astro	Astrophotography
Buildings	Architectural_photography, Architecture, Real_estate
City	New_York_City, Paris, San_Francisco, Seoul
Fashion	Fashion_design, Fashion_photography, Model_(person)
General	Aerial_photography, Culture, Documentary_photography, Photography, Social_documentary_photography
Landscape	Landscape_photography
Nature	Nature_photography
Plants	Flower
Portrait	Mug_shot, Portrait_photography, Selfie
US	Americans, President_of_the_United_States, United_States
Wildlife	Aquatic_ecosystem, Macro_photography, Marine_habitats, Wildlife_observation, Wildlife_photography

A.3 SCORING FUNCTION DETAILS

In Section 2.2, we defined the scoring function, $score(x, t)$ as

$$score(x, t) = \max(0, \min(100, \alpha \cdot \left\| \frac{C(x)}{\|C(x)\|_2} - \frac{C(t)}{\|C(t)\|_2} \right\|_2 + \beta)),$$

α and β are constants used to scale the embedding distance prior to clipping the score. To select α and β , we used a small dataset of interactions collected by the authors prior to the main dataset collection (this data is not included in the released dataset). This dataset contains groups of images paired with scores in the range $[0, 1]$. For each target image in this small dataset (5 in total), we add the following images to the dataset:

1. AI-generated images that use the target prompt but with a different seed. These images are assigned a score of 1.

2. The images corresponding to the human-generated prompts. These images are assigned a score of 0.5.
3. AI-generated images that use a different prompt than the target prompt. These images are assigned a score of 0.

The intuition here is that with the AI-generated images, we can assume using the same target prompt with a different seed should generate a similar image hence the highest score possible (1). Using a different prompt (from our prompt dataset Santana (2022)), however, should result in an entirely different image hence the lowest score possible (0). Images generated by people are assumed to be somewhere in between, hence the score of 0.5.

We then fit a linear regression model to this dataset (to predict score given the CLIP image embedding), using balanced sampling across the image groups. This linear model has parameters

$$\begin{aligned}\alpha' &= -1.503 \\ \beta' &= 1.791\end{aligned}$$

Since our score range is $[0, 100]$, we scale the model parameters by 100, resulting in

$$\begin{aligned}\alpha'' &= -150.3 \\ \beta'' &= 179.1\end{aligned}$$

We also add a “score adjustment” term that attempts to normalize image difficulty. In particular, for each target image, we compute the un-clipped score for the target image t_k and the image generated using the target prompt with, x_k :

$$\text{unclipped_score}(x_k, t_k) = \alpha'' \cdot \left\| \frac{C(x_k)}{\|C(x_k)\|_2} - \frac{C(t_k)}{\|C(t_k)\|_2} \right\|_2 + \beta''$$

This score assess the score a user would receive if they exactly entered the target prompt. We fix this value as 100 (i.e., a perfect score prior to clipping), and set the score adjustment parameter, c_k , to appropriately normalize this score. In particular,

$$c_k = \frac{100}{\text{unclipped_score}(x_k, t_k)},$$

and then we obtain target specific parameters,

$$\begin{aligned}\alpha_k &= -150.3 \cdot c_k \\ \beta_k &= 179.1 \cdot c_k\end{aligned}$$

Note that this means each target image may have slightly different parameters as the c_k values vary across images.

A.4 ADDITIONAL INFORMATION ON RUNNING THE GAME

Why we limit user input? We deliberately limited user input to only a prompt, as opposed to giving users access to the random seed or other model hyperparameters. This was done for a few reasons:

1. We wanted all users who enter the same prompt for a given image to see the same output.
2. We wanted to limit the complexity of the task for users less/unfamiliar with text-to-image models.
3. We wanted users to generate new prompts and not just resample new seeds until getting lucky. While users could still employ a version of this “random resampling” strategy by making small changes to their prompts, we did not want to encourage this practice through a seed parameter. This random resampling strategy, while useful in practice, is not such an interesting result for research purposes as it is easy to simulate random resampling strategies without any user input.

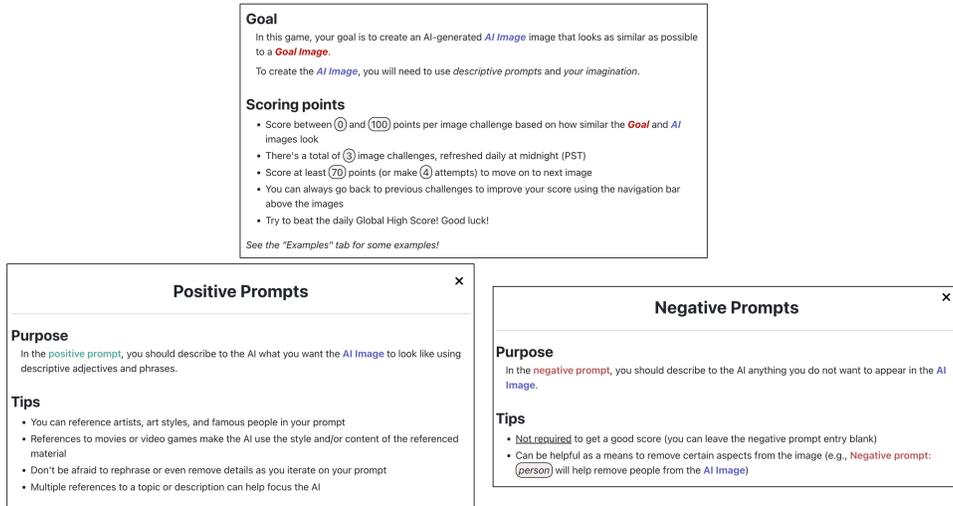


Figure 8: Game instructions

Additional Technical Details For the generative model, we use SD v2.1 Rombach et al. (2022b) with the DPM Multi-step Scheduler Lu et al. (2022) and run the model for 20 iterations. AI-generated target images use the same parameters but run for 50 iterations. 20 iterations was selected to limit latency for players to 1-3 seconds depending on the player’s internet connection. All images are generated at size 512×512 .

Game instructions Game instructions are provided in Figure 8. Here we show the main instructions provided on how to play (top), as well as the tool-tips given for positive prompts (lower left) and negative prompts (lower right).

Crowd workers Crowd workers are adults from the US. They were paid at a rate of \$12.00 per hour for roughly 20 minutes of time. Additionally, they were provided bonus payment of between \$0.10 – \$0.50 per image they received a perfect score on (depending on the image difficulty). In total, we paid about 600 for recruiting the crowd workers.

A.5 ADDITIONAL STATS

In Figure 9, we plot the distribution of user scores across target images. The left plot shows the initial and final scores. The shift in score to the right indicates the improvement in generated image similarity to target image. The average initial score is 56.7 (median is 57.0), and the average final score is 73.0 (median is 75.0). On the left, we plot the distribution of score improvement (the difference between the final and initial scores). The large density of 0 or close to 0 improvement is due to the “easiest” target images that users were able to generate similar examples of on their initial attempt (for example, see the first two rows of Figure 16). Most interactions result in a score difference of at most 50 points, as most users score above 50 points on their initial attempt, limiting the amount of improvement possible (the median initial score is 57.0).

In Figure 10, we plot how often a user’s score decreases, increases, or remains constant between prompts. Each point represents a series of user interactions to generate a single target image. We note that the ratio of score increases to score decreases is similar, with the number of score increases being slightly higher. This indicates that score improvement (and so, an increase in the similarity of the generated to target image) may be effectively represented by a stochastic process with a slight bias towards score increases. This observation is the basis on which we construct our proposed steerability metric (i.e., representing score change as a stochastic process).

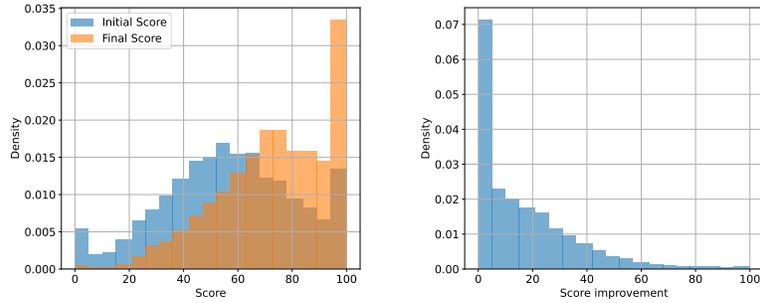


Figure 9: On the left: distribution of initial and final scores (across users and targets). Average initial score is 56.7; average final score is 73.0. On the right: distribution of image improvement (difference between final and initial score). Average score improvement is 16.3.

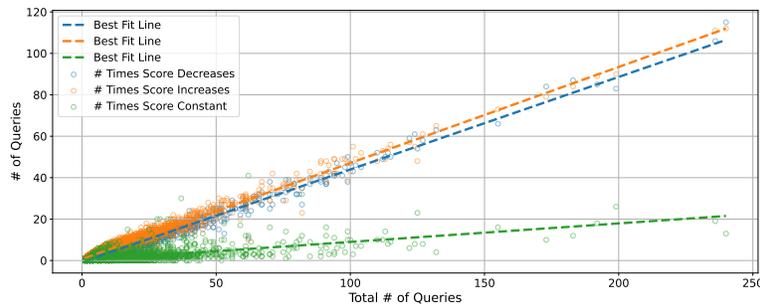


Figure 10: How score changes between queries. Each circle represents a series of user interactions to generate a single target image. In blue, we plot the number of times a user’s score decreases over that interaction (between two queries); similarly, in orange and green we plot the number of times the user’s score increases or stays constant between prompts. We also plot best fit lines (fit using linear least squares).

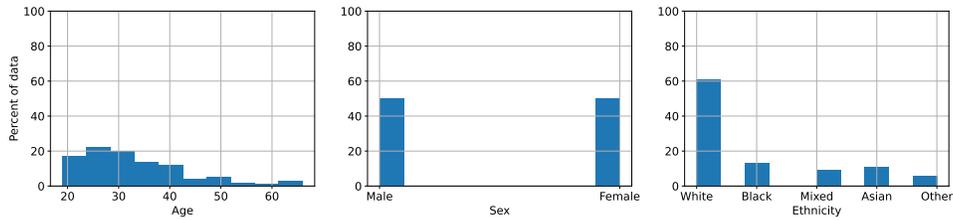


Figure 11: Demographic information for the paid crowd workers in the ArtWhisperer-Validation dataset.

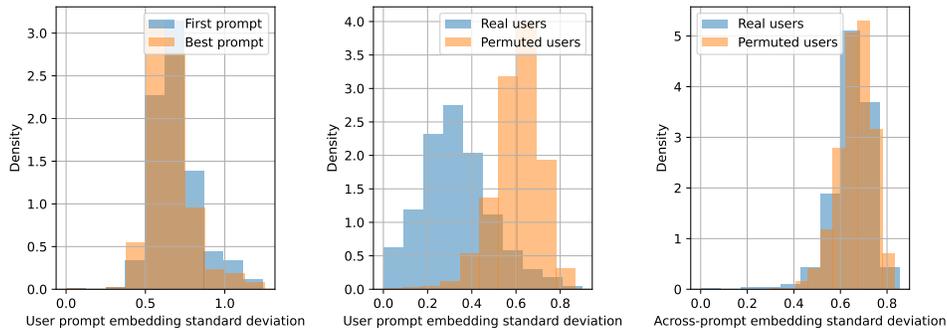


Figure 12: Same plot as Figure 6 but using the BERT Devlin et al. (2018) text embedding instead of CLIP.

A.6 ARTWHISPERER-VALIDATION STATISTICS

In Table 3, we provide general statistics about the *ArtWhisperer-Validation* Dataset. This information mirrors that statistics provided in Table 1 for the *ArtWhisperer-Validation* dataset. We also provide demographic information collected through Prolific in Figure 11. Users were sampled through Prolific to guarantee an even split of male and female users.

Table 3: *ArtWhisperer-Validation* Dataset Overview. Each row contains summary data for a different subset of the dataset. Subsets may overlap.

# Players	# Target Images	# Interactions	Average # Prompts	Average Score	Median Duration	Category
140	51	4572	8.14	54.77	24 s	Total
73	9	762	7.47	47.55	20 s	Famous person?
134	42	3937	8.19	54.58	23 s	Manmade?
124	26	2154	7.69	54.61	24 s	Real image?
123	24	2424	9.32	50.96	24 s	Art?
94	13	1044	7.05	57.93	25 s	Famous landmark?
72	9	858	8.94	50.33	26 s	Nature?
70	8	730	8.11	60.25	25 s	Sci-fi or space?
94	16	1501	8.25	54.21	24 s	People?
123	25	2418	8.57	54.90	24 s	AI image?
75	9	786	7.42	59.11	23 s	City?
75	9	808	8.16	57.56	26 s	Fantasy?

A.7 RESULTS FOR ALTERNATIVE TEXT EMBEDDINGS

In Figure 12, we replicate the analysis from Figure 6 but using a BERT Devlin et al. (2018) embedding rather than a CLIP embedding. We note similar findings and significance to the CLIP embedding analysis. We also repeated this analysis using the GLOVE Pennington et al. (2014) text embedding and again see similar results. This indicates that our analysis is robust across text representations.

A.8 ADDITIONAL EXAMPLE IMAGES

We provide additional examples of image trajectories and diverse images in Figures 13 and 14.

A.9 ALGORITHM FOR STEERABILITY

We describe the algorithm for assessing steerability in more detail in Algorithm 1. Here, we define three procedures. `EstimateStoppingTime` estimates the steerability of a target image. We define a set of score bins; we chose 5 equally sized bins so that there was sufficient data to cover each bin. We also use a regularizer, ϵ . What ϵ essentially does is encode a prior that from any given score, the transition to a new score is uniformly random. Then for each target image, we find the empirical score transition probabilities in `EstimateMarkovChain` and use Monte Carlo simulation to estimate the stopping time in `RunMonteCarloEstimation`.

Algorithm 1: Algorithm for estimating image steerability.

Input: Set of Images

Output: Steerability estimate

Procedure `EstimateStoppingTime` (*targetImages*)

```

Initialize array steerability to track stopping times;
Initialize scoreBins  $\leftarrow$   $[[0, 20], [21, 40], [41, 60], [61, 80], [81, 100]]$ ;
Initialize  $\epsilon \leftarrow 1$ ;
for imagei  $\in$  targetImages do
    | markovi  $\leftarrow$  EstimateMarkovChain (imagei, scoreBins,  $\epsilon$ );
    | steerability[i]  $\leftarrow$  RunMonteCarloEstimation (markovi);
end
return  $E[\textit{steerability}]$ 

```

Procedure `EstimateMarkovChain` (*targetImage*, *bins*, ϵ)

```

Initialize bin pair count as counts[(bini, binj)]  $\leftarrow$   $\epsilon$ , using transitions from a dummy node to
model the first prompt submitted;
for useri  $\leftarrow$  user1 to usern do
    | for scorei,j  $\leftarrow$  scorei,1 to scorei,d do
        | | Convert scorei,j to bin number, bini,j;
        | | Increment counts[(bini-1,j, bini,j)] by 1;
    | end
end
Normalize empirical transition counts to find empirical node transition probabilities;
Define markovtargetImage using the node transition probabilities;
return markovtargetImage

```

Procedure `RunMonteCarloEstimation` (*markovChain*)

```

Run Monte Carlo simulation to estimate time to reach the last bin for markovChain, starting
from the dummy (initial) node;
return Estimated stopping time;

```

A.10 STEERABILITY ACROSS MODELS

In Figure 15, we plot the steerability across SDv2.1 and SDv1.5. As described in Section 4.2, images of nature, sci-fi or space, and real images have the largest differences in steerability between the two models. “Nature” is the only image group with a steerability difference greater than the standard deviation of the mean. Other image groups seem to have similar performance across both SDv2.1

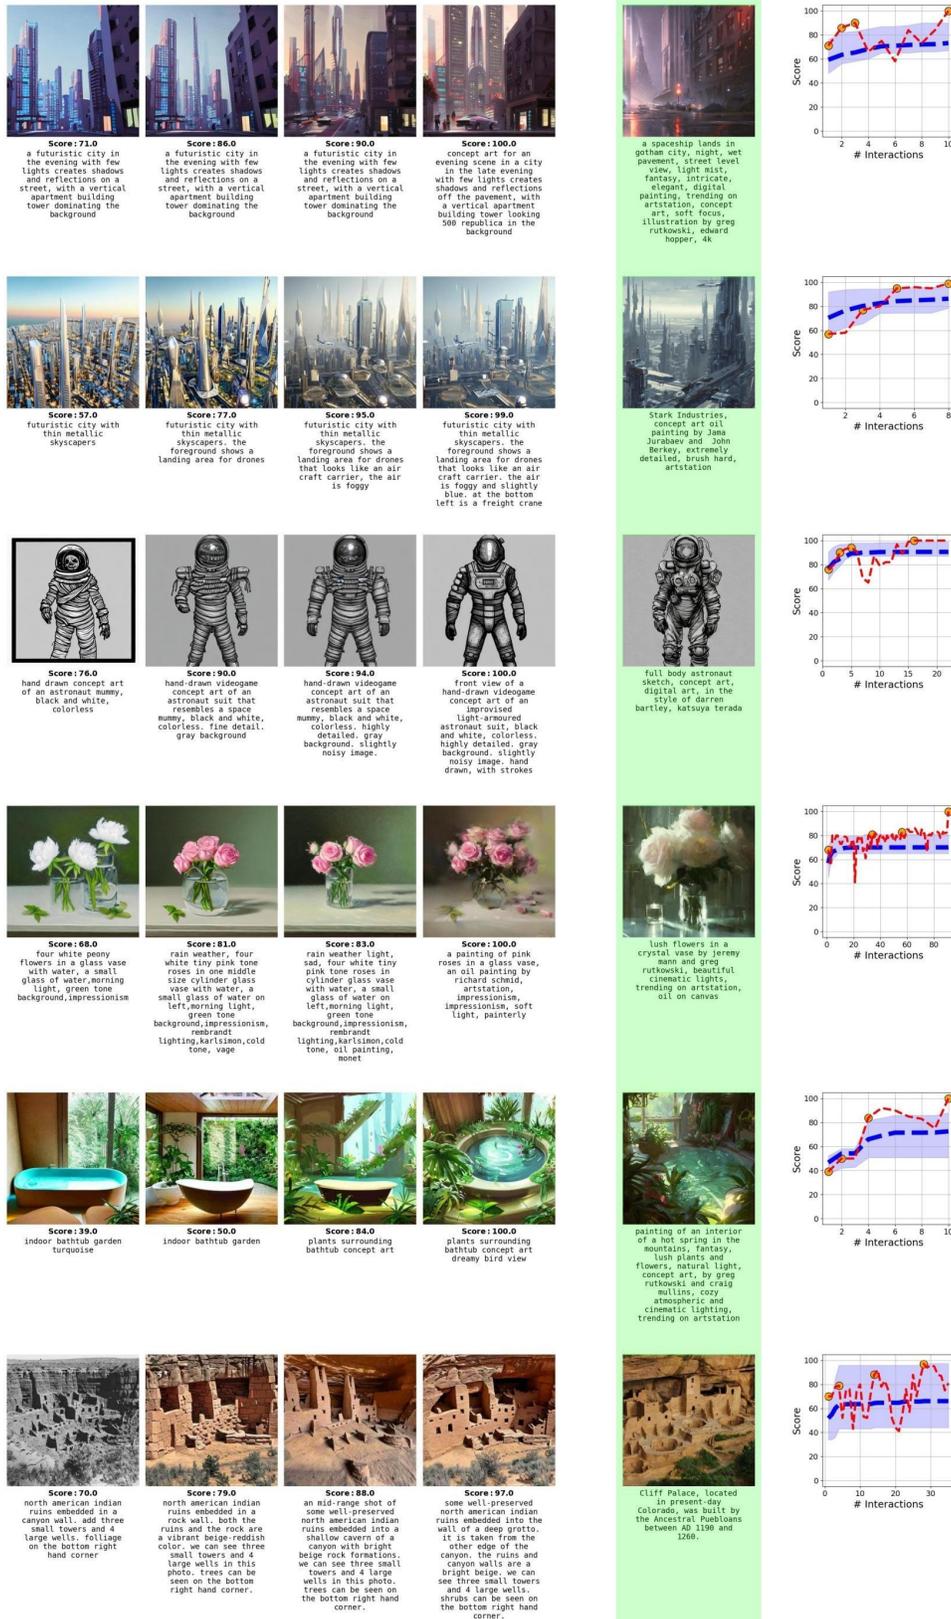


Figure 13: More examples of user trajectories, as in Figure 2.

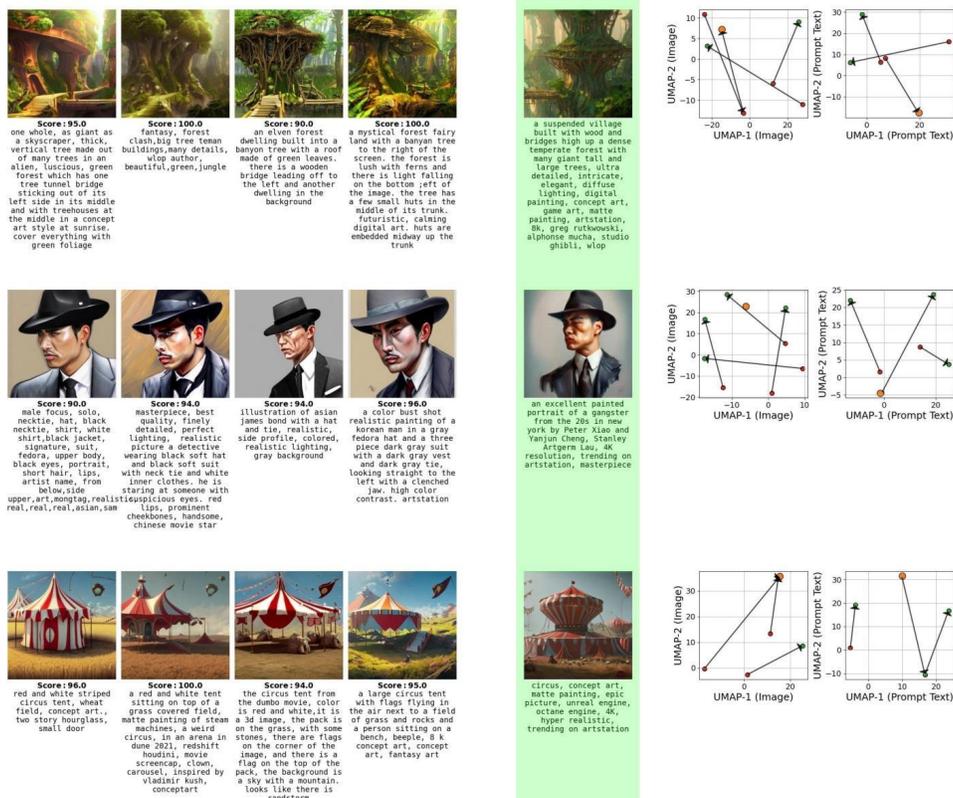


Figure 14: More examples of diverse high-scoring prompts, as in Figure 4.

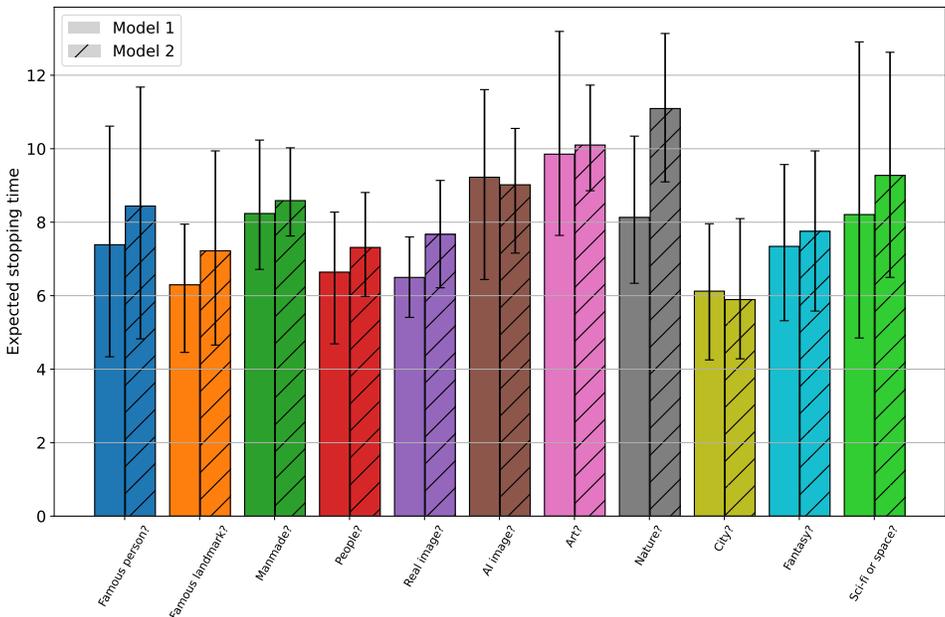


Figure 15: Steerability across image groups, comparing between models using a subset of the images. "Model 1" is Stable Diffusion v2.1 and "Model 2" is Stable Diffusion v1.5. Note only a subset of the images are used to compare the two models, hence the differences between "Model 1" and the results in Figure 7.

and SDv1.5. This suggests that SDv2.1 only makes minor improvement over SDv1.5 across most image categories.

A.11 STEERABILITY SCORES FOR INDIVIDUAL IMAGES

In Figure 16, we provide some example images along with their steerability scores. Note that more simple images with well-defined content that likely has high presence in the model’s training data (e.g., the first two rows—a fly on a leaf; a drawing of Barack Obama, well-known public figure) have smaller steerability values, indicating they are easier to steer. However, more complex content that is also more ambiguous for users (e.g., the last three rows), have larger steerability indicating greater difficulty in steering.

We also find that image steerability is negatively correlated with image variance (across seeds for a fixed prompt). Consider fixing a target prompt and sampling the model across many seeds. When the model outputs a wide variety of images, the variance of the output increases (i.e., computed in the image embedding space). We find that target images that have higher variance are also less steerable (indicated by a higher expected stopping time). In other words, if repeatedly sampling a model with the same prompt can produce a wide range of outputs (high variance), then the model is likely less steerable for the content in those generated images. We plot this result in Figure 17.

A.12 CONTENT KNOWLEDGE INCREASES STEERABILITY

Knowledge of an image’s subject matter increases steerability. In Figure 18, we plot the steerability across two groups of individuals – those who had subject matter knowledge and those would did not. To define this split, we only consider the subset of images that contain either a famous person or landmark. Users were defined to have knowledge of the subject matter when they referenced a key word (e.g., the famous person’s name) in at least one of their submitted prompts. We then computed steerability across these two groups of users. As is indicated in the plot, users with subject matter knowledge found the model to be significantly more steerable.

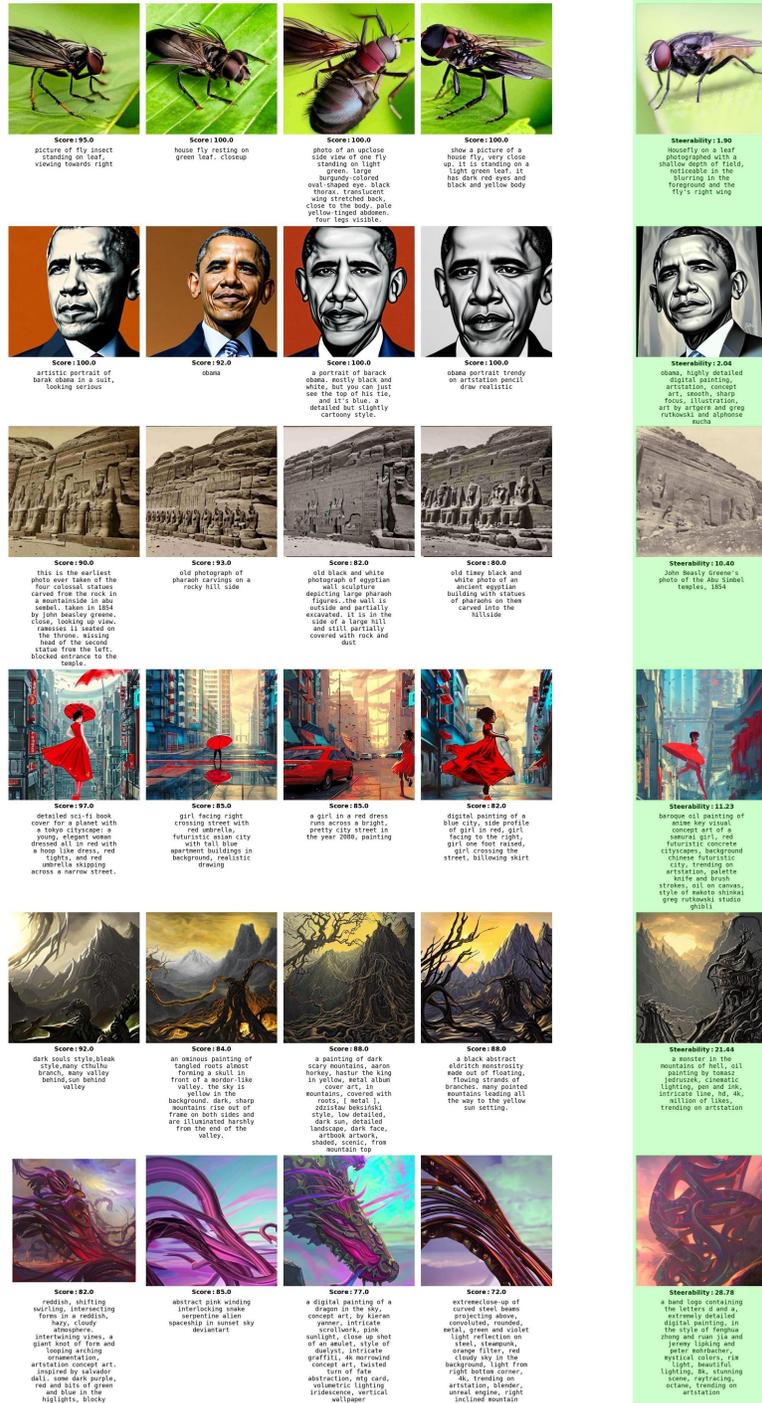


Figure 16: Steerability of individual target images. Each row is a different target image. The first four columns are examples of user submitted prompts with the corresponding AI-generated images. Right column (highlighted in green) shows the target image along with the *steerability* of that image. A smaller steerability value corresponds to the image being easier to steer.

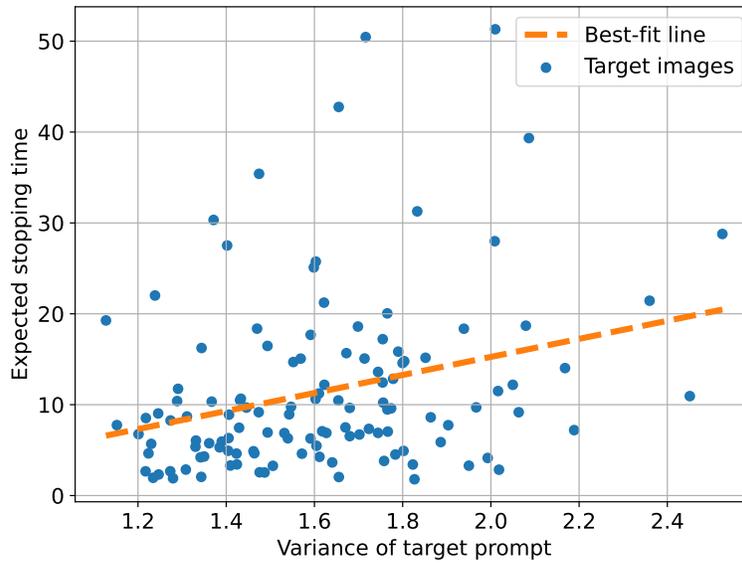


Figure 17: Steerability decreases (expected stopping time increases) when a model has higher variance given a target prompt.

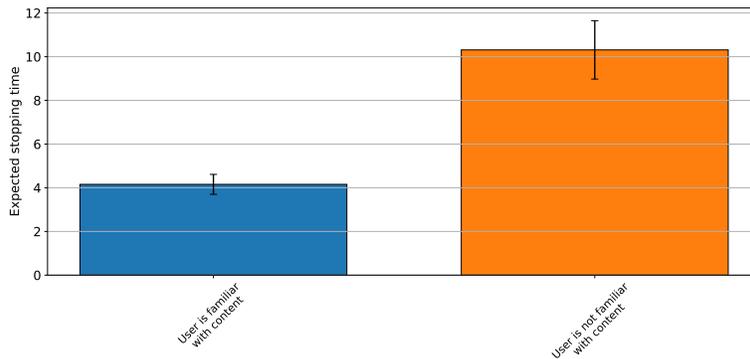


Figure 18: Steerability across familiarity with target image content. When a user is more familiar with the subject matter of the target image, it is significantly easier to steer the model.

A.13 ADDITIONAL DISCUSSION AROUND HUMAN RATINGS

Here we provide plots for the analysis using human ratings. In Figure 19, we show a scatter plot of scores and human ratings. We also plot a best fit line which has a correlation of 0.597, indicating that our score function produces values that are indeed similar to the human ratings.

We also provide a plot comparing the steerability value calculated using our score function and calculated using the human ratings in Figure 20. Error bars indicate the standard error. Images depicting sci-fi or space have the greatest difference (humans seem to be harsher judges of their generated images’ similarity to the target. However, for most image groups, the two steerability scores are quite close and generally exceed a 95% confidence interval.

A.13.1 ON OUR CHOICE OF SCORING METRIC AND ITS RELATION TO HUMAN RATINGS

These results validate our scoring metric based on CLIP. As is well known, embeddings extracted from deep models are reasonable at assessing image similarity Zhang et al. (2018). We chose to base our scoring metric using the CLIP embedding in particular, however, for two reasons: (1) its recent usage in the literature for effectively representing semantic meaning in images Radford et al. (2021); Khandelwal et al. (2022) and (2) after testing a number of methods on a small subset of trial data before launching our data collection, we found the CLIP embedding was comparable or better than any of the other methods including ensemble-based approaches. (The methods we tested include embeddings extracted from other deep learning architectures (including ResNet He et al. (2016) and VGG Simonyan & Zisserman (2014)) as well as classical image embeddings (like color histograms).)

Despite this, as is seen in Figure 19, the CLIP-based score does not perfectly correlate with the human ratings. There are two reasons the Pearson correlation coefficient is not larger: (1) noise in the human rating responses and (2) deficiencies in the CLIP embeddings, as you mentioned. While improving the image similarity metric is desirable, it was not the focus of our paper. For running ArtWhisperer, we only required a metric that worked reasonably well to generate a “reasonable score” (i.e., which we define as showing at least moderate correlation with recorded human rating assessments). Any subsequent analysis can of course use specialized similarity metrics on our data as we release all target and generated images.

The results plotted in Figure 20 also affirm CLIP as a reasonable choice for a metric – the score (based on the CLIP embedding) and human rating, while resulting in different steerability scores, offer similar conclusions across most image categories (e.g., the images that are AI-generated are less steerable than real images).

A.14 SYNTHETIC PROMPTERS

In this section, we describe a proof-of-concept Synthetic prompter based on the ArtWhisperer dataset. We fine-tuned a pretrained MT0-large model Muennighoff et al. (2022) using the IA3 method Liu et al. (2022), training for 30 epochs and using a linearly decaying learning rate starting from 10^{-3} with the AdamW optimizer Loshchilov & Hutter (2017).

Our train dataset is based on the main ArtWhisperer dataset. We first randomly sampled 80% of the ArtWhisperer dataset for training, with sampling taken over the target images (to ensure there are unseen images we can test on). For each user trajectory in the training set, we sampled a sliding window of user prompts. This is done to ensure the total prompt length of the model isn’t too long to enable efficient training. Recalling that 50% of the trajectories in the ArtWhisperer dataset have ≤ 5 total prompts, we select a window size of 6 as a reasonable cutoff (allowing a final prompt generation after seeing the 5 previous prompts).

Given that MT0-large is a text-only model, we need a way of encoding the image information for the model to simulate a user prompting on the given image. We opt to simply provide the model with the target prompt; after fine-tuning, the model uses the target prompt to condition its generation but does not repeat it verbatim. In addition to the target prompt, we also provide the model with a history of user prompts and scores indicating how well the image generated by a given prompt matches the target image. For training, we define a loss function on the next-prompt token probabilities to encourage the model to learn to predict subsequent prompts given the prompt and score histories.

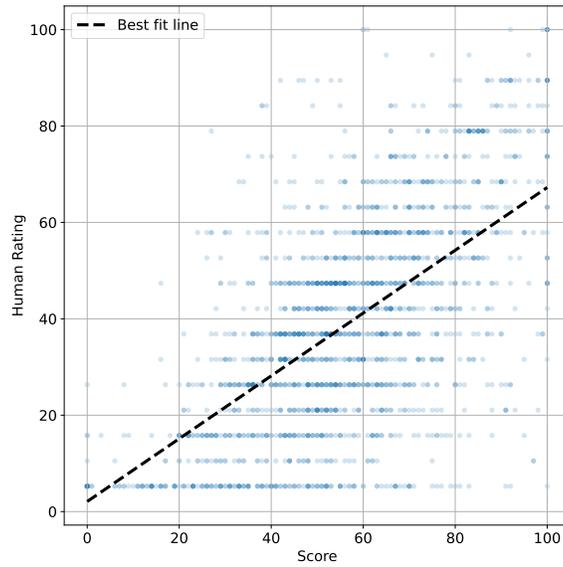


Figure 19: Scatter plot of scores and human ratings. Best fit line (linear regression) is shown in black.

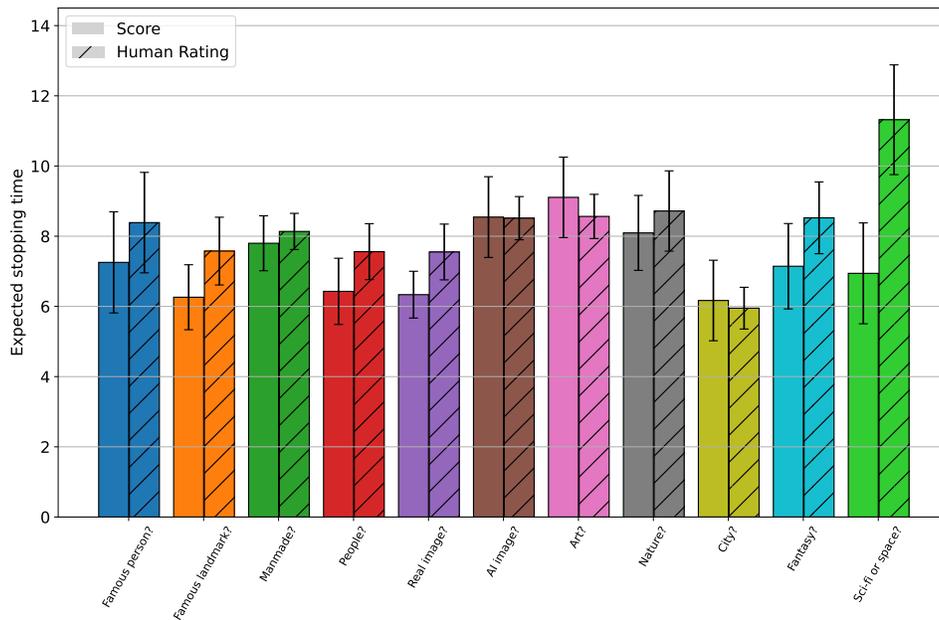


Figure 20: Comparing steerability across image groups when computing steerability using the score function and when using the human rating function.

```
[Goal score]: 100 [Goal prompt]: "{goal_prompt}"
1. [User score]: {user_score_1} [User prompt]: "{user_prompt_1}"
2. [User score]: {user_score_2} [User prompt]: "{user_prompt_2}"
...
N. [User score]: {user_score_N}
```

Figure 21: Prompt format for training and querying the synthetic prompter. $N \leq 6$ in our training and evaluation.

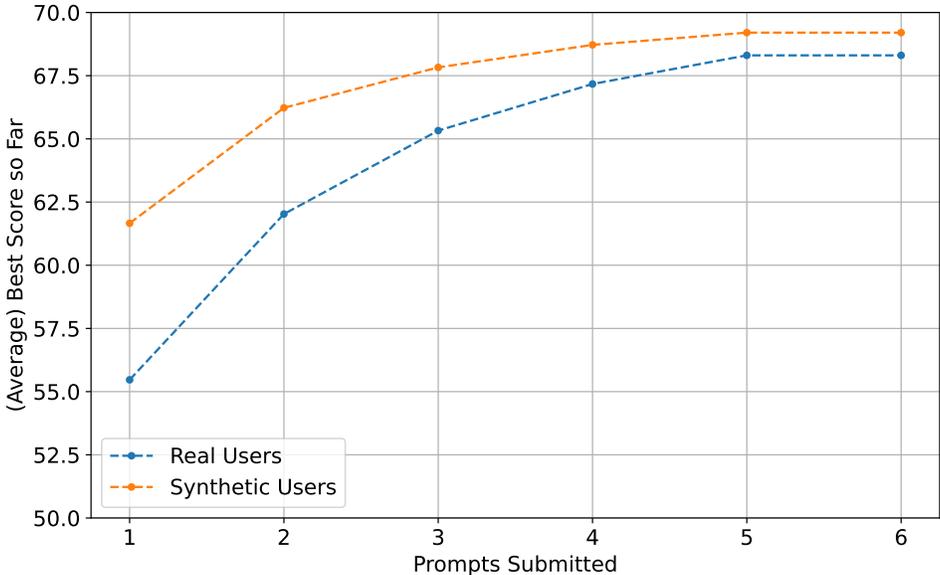


Figure 22: Comparing score progression for the synthetic prompter to real users.

The prompt is shown in Figure 21. The “goal prompt” (the target prompt) conditions the model on the target image content. The “user_score_i” values are set based on evaluating the distance between the target and generated image, though in general, we could use any metric of our choosing; the “user_prompt_i” values are set based on previously generated prompts. “user_score_N” is used to condition generation of a new prompt, and enables easy re-sampling by changing the “score” of the new prompt. Note this score is entered pre-evaluation, so we can set it to any value we like; in our evaluations, we set it to a randomly sampled value between 60 and 90.

A.14.1 ANALYSIS OF SYNTHETIC PROMPTERS

Here we present an initial analysis that demonstrates the effectiveness of our dataset for training synthetic human prompts. After fine-tuning, we evaluate the synthetic prompter on the held-out test set. The target images in this test set (39 images) were excluded from all training data. For each image, we generate 10 sample trajectories. In Figure 22, we plot the average best score for the synthetic prompter and the real human prompts on the same test data. Note that while the synthetic prompter starts with a higher average, both approach a similar average score after 6 prompts. This suggests that the synthetic prompter indeed shares some similarities with the human prompters.

In Figure 23, we plot two sample trajectories generated by the synthetic prompter. Note that across images, the prompter makes incremental changes; these changes are not restricted to appending phrases, but can involve substitutions and deletions throughout the prompt; this is in contrast to other automated prompters that exclusively complete text (these methods are not designed to simulate human behavior) Hao et al. (2022).

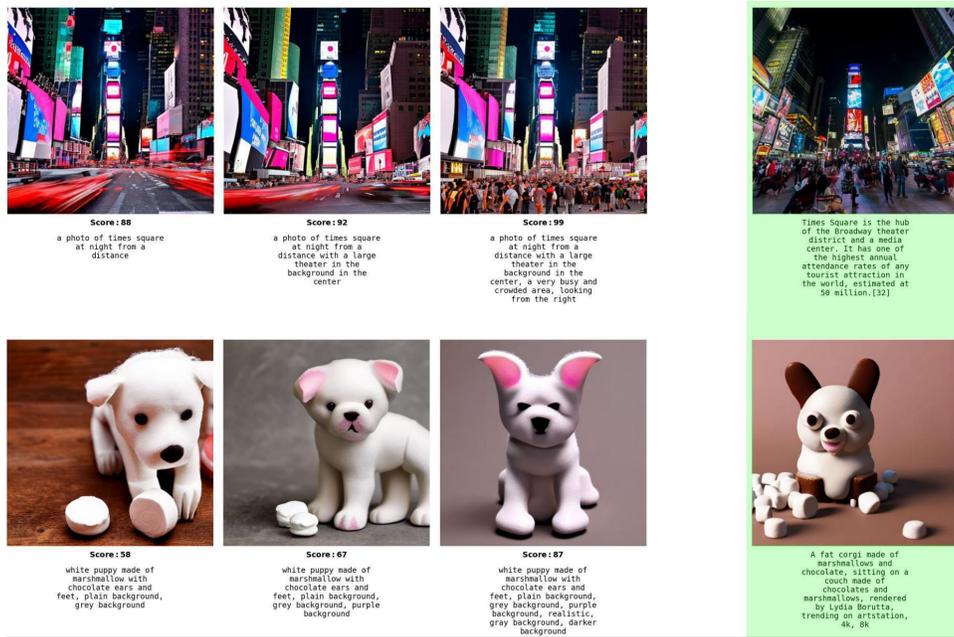


Figure 23: Example synthetic prompter trajectories. First 3 images show prompt progression for the synthetic prompter. The target image is shown in the green column.