

# Evaluating Text-to-Image Synthesis: Survey and Taxonomy of Image Quality Metrics

Sebastian Hartwig<sup>1</sup>, Dominik Engel<sup>1</sup>, Leon Sick<sup>1</sup>, Hannah Kniesel<sup>1</sup>, Tristan Payer<sup>1</sup>,  
Poonam Poonam<sup>2</sup>, Michael Glöckler<sup>1</sup>, Alex Bäuerle<sup>1</sup>, Timo Ropinski<sup>1</sup>

An astronaut floating in the ocean holding a piña colada with orange-colored decoration



<b>Object Accuracy</b>	astronaut ocean piña colada decoration	✓ ✓ ✓ ✓	✓ ✓ ✓ ✓	✓ ✓ ✗ ✓	✓ ✓ ✗ ✗
<b>+ Spatial relation</b>	astronaut in ocean	✓	✓	✓	✗
<b>+ Non-Spatial relation</b>	astronaut holding piña colada piña colada with decoration	✓ ✓	✓ ✓	✗ ✓	✗ ✗
<b>+ Attribute binding</b>	orange-colored decoration	✓	✓	✗	✗
<b>= Compositional Quality:</b>		High	High	Low	Low
<b>Realism</b>		✓	✗	✓	✗
<b>+ Aesthetic</b>		✓	✗	✓	✗
<b>+ No Artifacts</b>		✓	✗	✓	✗
<b>= General Image Quality:</b>		High	Low	High	Low

Fig. 1. Evaluating text-to-image synthesis models includes two types of quality measures that contribute to an overall image quality. *Compositional Quality* measures how well the image reflects the composition defined in the text prompt. *General Image Quality* measures the overall quality of the image. For both types, several different aspects can be considered, e.g., realism might be important only for some applications. After measuring each aspect of the two categories, an aggregated quality score can be computed.

**Abstract**—Recent advances in text-to-image synthesis enabled through a combination of language and vision foundation models have led to a proliferation of the tools available and an increased attention to the field. When conducting text-to-image synthesis, a central goal is to ensure that the content between text and image is aligned. As such, there exist numerous evaluation metrics that aim to mimic human judgement. However, it is often unclear which metric to use for evaluating text-to-image synthesis systems as their evaluation is highly nuanced. In this work, we provide a comprehensive overview of existing text-to-image evaluation metrics. Based on our findings, we propose a new taxonomy for categorizing these metrics. Our taxonomy is grounded in the assumption that there are two main quality criteria, namely compositionality and generality, which ideally map to human preferences. Ultimately, we derive guidelines for practitioners conducting text-to-image evaluation, discuss open challenges of evaluation mechanisms, and surface limitations of current metrics.

**Index Terms**—Artificial Intelligence, Surveys, Machine learning, Text analysis, Image Generation

## 1 INTRODUCTION

THE rapidly evolving landscape of text-conditioned image generation has emerged as a pivotal area in com-

puter vision and natural language processing [1]–[3]. As the demand for seamless integration between textual and visual information intensifies, understanding the intricate mechanisms behind accurate text-to-image (T2I) alignment becomes imperative. Robust metrics for T2I alignment are essential for objectively evaluating high-quality image generation that aligns with human judgment and facilitate advancements in research and technology.

In this work, we provide an overview of such metrics to help practitioners make informed decisions when eval-

- S. Hartwig, D. Engel, L. Sick, H. Kniesel, T. Payer, P. Poonam, M. Glöckler, T. Ropinski are with Visual Computing Group located at Ulm University  
E-mail: {forename}.{surname}@uni-ulm.de
- Alex Bäuerle is Postdoctoral Researcher at Carnegie Mellon University  
E-mail: alex@a13x.io
- Website: <https://huggingface.co/spaces/kopetri/text-to-image-evaluation>

Manuscript received April 5, 2024.

uating their T2I models. In doing so, we consider the T2I synthesis process as a black box where a model produces an image for a given textual input (see Figure 1). Furthermore, we define the overall quality of an image conditioned on a text prompt, as a combination of *general image quality* and *compositional quality*, where the latter measures the degree of alignment between the text and the image. Consequently, a high compositional quality score can only be achieved if all details described in a text prompt are visually represented in the image.

To help practitioners evaluate their T2I systems, we delve into the current state of research on measuring T2I alignment, explore the strengths and limitations of existing metrics, review T2I optimization methods, and ultimately aim to contribute guidelines that propel advancements in the synergy between measurement and optimization of T2I alignment. Through a comprehensive analysis of recent developments, we endeavor to elucidate key challenges and avenues for future exploration in the dynamic field of text-conditioned image generation. By improving the selection of appropriate metrics, we support the development of T2I applications and enhance our understanding of T2I models.

The importance of such alignment metrics extends even beyond the text-to-image domain. In fact, it is fundamental for application areas such as text-to-video [4]–[6], where multiple frames are generated for a single text prompt, and text-to-3D, where image-based NeRF approaches and diffusion models produce 3D representations for textual scene descriptions [7]–[11] as well. Moreover, an increasing number of web-based applications, like DALL-E, ImageFX, DreamStudio, Midjourney, and civitai.com, offer simple user interfaces and capable hardware, making modern generative AI accessible not only to researchers but also to novices. In these scenarios, the integration of appropriate T2I quality metrics could foster the further advancement of such application areas of T2I synthesis.

Within the remainder of this work, we systematically address essential components to provide a comprehensive overview of evaluation strategies for T2I synthesis. Therefore, we first derive a taxonomy in Section 2, which is based on the current state of the art (SOTA) for T2I evaluation. Next, we categorize recently published metrics utilizing our proposed taxonomy in Section 3, which are used for the assessment of text-image synthesis, before we discuss promising approaches for T2I optimization in Section 4. Finally, we address the remaining challenges of current evaluation frameworks in Section 5, and propose guidelines for comprehensive and consistent evaluation of future frameworks in Section 6. We conclude our survey in Section 7.

## 2 TAXONOMY

To provide a comprehensive overview of T2I evaluation metrics, surface open challenges, and derive research guidelines, we developed a taxonomy that categorizes existing methods based on their operating data structure, measured aspects, scope, and conditions used. Defining requirements for metrics that can benchmark image synthesis models is not straightforward. Realism is undoubtedly the major aspect targeted by researchers [12]. However, how we

interpret realism is highly dependent on the text conditioning, e.g., image can be photorealistic, realistic in the context of a manga or realistic in the style of Pablo Picasso. Other aspects that contribute to high-quality images include aesthetics [13], human preferences [14]–[16], naturalism, and the principles of photography, such as balance, harmony, closure, movement, color, pattern, contrast, negative space, and grouping. While some of these aspects may be quantitatively measurable, many are abstract, complex, and therefore hard to measure. However, natural language can describe such aspects in great detail, and there are many talented authors who generate creative descriptions of sceneries. Hence, detecting and measuring the quality of these abstract, yet well-described aspects, presents a challenge to researchers in the field of text-conditioned image synthesis.

We consider both pure image-based evaluation metrics, which are designed to measure image quality, as well as text-image evaluation metrics, which measure the alignment between the textual prompt and the generated image. In particular, the input to a text-image evaluation metric is a (human written or synthetic) prompt together with an image. Pure image-based metrics, on the other hand, solely judge the generated image.

### 2.1 Image Metrics

Image-based metrics do not consider the textual input to image generation at all. We include the palette of pure image-based quality measures in our taxonomy despite the absence of such textual condition since, quality assessment of T2I synthesis includes assessing image quality, which does not necessarily rely on text-conditioning. Moreover, image quality is evaluated independent of text-image quality, e.g., an image that looks photorealistic but ignores textual facts contributes to a high image quality score but a low alignment score. Similarly, an image that exactly depicts all objects and their relations described by the prompt can simultaneously look artificial, resulting in low image quality but high alignment quality. In Figure 1, we present such examples. The proposed taxonomy differentiates between image quality metrics that rely solely into distribution-based metrics and single image-based metrics.

#### 2.1.1 Distribution-based Metrics

Distribution-based metrics fundamentally rely on statistical measures to evaluate the differences between the distribution of generated data and the distribution of given, usually real-world, data. This provides a quality score for the performance of the generator, e.g., Inception Score (IS) [17] and Fréchet Inception Distance (FID) [18]. Distribution-based of metrics consider the image generator as a black box and base their quality computation on a sample distribution. In general, these metrics benchmark models rather than providing quality scores for individual output samples.

#### 2.1.2 Single Image Metrics

Single image metrics measure quality for individual images by analyzing an image based on its structural and semantic composition. These metrics extract features from the image

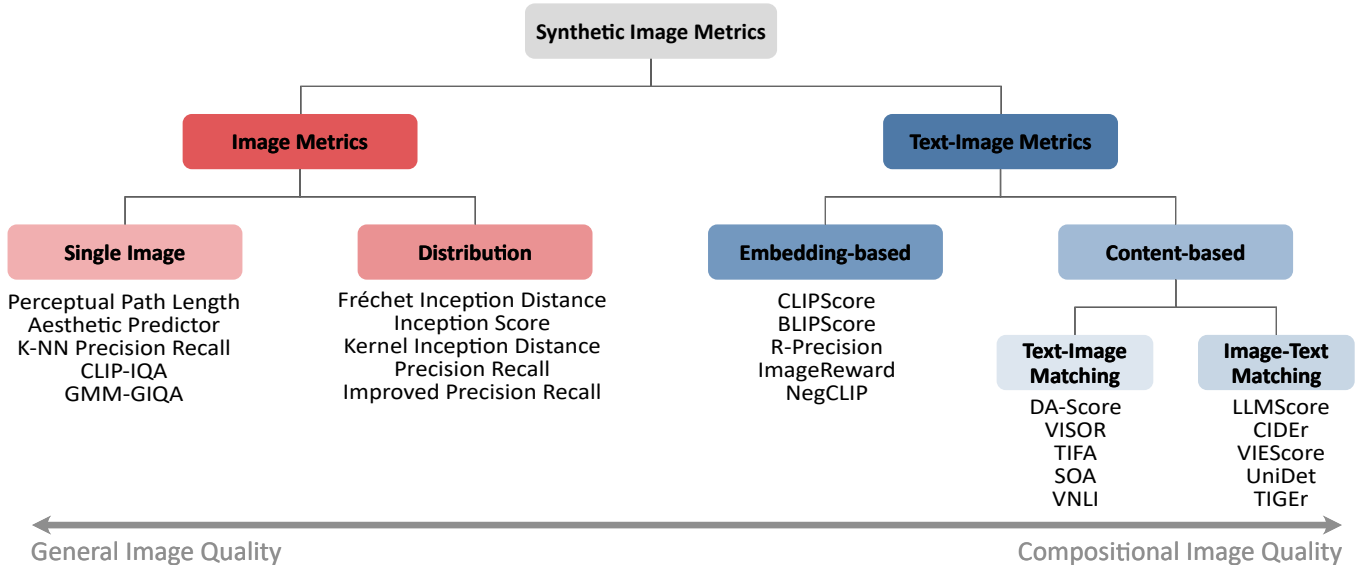


Fig. 2. Proposed taxonomy and examples for text-to-image (T2I) evaluation metrics. Two categories of metrics need to be distinguished: pure image-based and text-conditioned image quality metrics. Synthetic image metrics can measure two different qualities which correlate with this categorization, namely general image quality and compositional quality.

and subsequently infer quality. In contrast to distribution-based metrics, single image metrics are computed without any target image or distribution. Recent iterations of such metrics often rely on a fine-tuned image model that is trained to predict human judgments, e.g., LAION Aesthetic Predictor [13], perceptual artifact localization (PAL) [19], [20] or human viewpoint preferences (HVP) [21].

## 2.2 Text-Image Alignment Metrics

Text-image metrics measure the degree to which a textual prompt is aligned with a generated image. In the group of text-image metrics, we distinguish between embedding-based metrics, which quantify image generation quality based on text-image alignment and content-based metrics, which examine the content of both the generated image and the text prompt.

### 2.2.1 Embedding-based Metrics

In the context of embedding-based metrics, the evaluation of quality is based on learned embedding representations for vision and language inputs. Therefore, text prompts get tokenized by a tokenizer and are then transformed to an embedding vector using a text encoder model, e.g., a transformer [22]. Similarly, the image gets transformed into an image embedding representation using an image encoder model, e.g., ViT [23]. These embedding vectors have a fixed size and carry compressed information of both representations. Since the employed foundation models are trained via representation learning in order to output meaningful embeddings, one can compute the cosine similarity between text and image embeddings to measure alignment.

When training a powerful text-to-image model, embedding vectors for text-image pairs are aligned via vision-and-language pre-training strategies, e.g., CLIP [24], BLIP [25], BLIP-2 [26]. The extracted embedding vectors from these models encode valuable information, resulting in superior performance for multiple zero-shot scenarios [25]–[27].

However, several works have shown that pre-trained representations can be further fine-tuned on human-annotated data. This way, human judgments can be incorporated into models, tweaking embedding alignment quality measures towards missing nuances of human judgments, as done for instance by PickScore [15] and ImageReward [14].

### 2.2.2 Content-based Metrics

Content-based metrics analyze language and visual representations with respect to their semantic content, whereby the actual measurements of such metrics are computed for decomposed components separately. Content-based metrics are based on the way in which humans would compare content across the text and image domains, e.g., reading words in a prompt and matching them to regions depicted in the image, and vice versa. Hence, content-based quality metrics are comprehensible for human observers due to their relatable behavior.

**Text-Image Content Matching.** To relate parts of the text prompt to image regions, the text prompt needs to be dissected into substrings, where each substring describes distinct details, e.g., an object, the relation between two objects, scene settings, etc. This decomposition into distinct assertions is elementary for text-image content matching. Some benchmark datasets [28] synthetically compose prompts utilizing prompt templates in order to generate object relations, e.g., “{objectA} {spatial relation} {objectB}”. Using such a prompt dissection, the resulting set of assertions is compared to corresponding regions in the image. This can be done by utilizing a visual question answer model (VQA) [26], [29], where questions are generated based on the assertions together with the generated image. The VQA model is interrogated for the presence or absence of specific relations, objects, attributes etc. Another way is to extract meaningful regions from the image. Therefore, candidate regions need to be found in the image. This is usually done via object detection [30] or semantic segmentation [31].

**Image-Text Content Matching.** Starting the dissection from the image side and matching image regions to corresponding positions of the prompt can be seen as the inverse to text-image content matching. However, text-image and image-text content matching is not bijective, since there might be parts in the image that are not mentioned in the text and vice versa. Although the process is quite similar, object detection and segmentation models facilitate finding meaningful regions, which are then matched with corresponding parts of the prompt. Additionally, image caption models can be used to generate captions [29], [32], which describe the presented scene. Image captioning is an ongoing topic within the vision-language model community, and for the evaluation of such models, image caption metrics are utilized [33]–[38].

### 2.3 Types of Quality Measures

In addition to this proposed taxonomy, metrics can be assigned to two different types of image quality measures (cf. Figure 1). We define *general image quality* as quantifying a certain aspect globally for a single image, e.g., realism, aesthetics, and human preferences, for which ground-truth can be collected by asking human raters for their judgments. Usually, this is done by conducting a large-scale crowd-sourced study where images are ranked by a large group of human observers, which is a high effort endeavor. Hence, practitioners use the acquired human annotations to develop deep learning based evaluation models which designed to imitate such human judgments, e.g., PickScore [15], ImageReward [14], and Human Preference Score [16], [21], [39].

In addition to *general image quality*, we further investigate *compositional quality*, which is measured by dissecting the prompt and the image into multiple text-image pairs and measuring the alignment quality of these pairs. In particular, quality is measured by analyzing the alignment between the prompt specification and content depicted in the image, e.g., through measures like **object accuracy (OA)**, **spatial relation (S)**, **non-spatial relation (NS)**, and **attribute binding (AB)** [40]–[42]. Usually, the prompt is composed of multiple distinct pieces of information that describe different parts of a scenery. These pieces of information accumulate to a rich scene description. Specifically, a complex prompt can be decomposed into a set of disjoint assertions describing different parts of the content, e.g., single or multiple objects, relations between objects, object attributes, lighting, style, and artistic reference. Thus, the composition of assertions must be known or extracted from the prompt. This concept is akin to Winoground’s [43] notion of visio-linguistic compositional reasoning. It refers to the task of understanding and reasoning about the relationships between visual and textual components in a way that requires combining them to form a coherent understanding or to make inferences. This involves tasks that require not just recognizing objects or elements in images and understanding text, but also understanding how the textual and visual elements interact and compose to convey a particular meaning. We refer to methods that can measure the relations between such individual components as compositional quality metrics.

As such, the proposed taxonomy not only differentiates between text-image and image metrics, but also considers

overall quality metrics, which measure global image properties, and compositional quality metrics, which measure text-image alignment with a finer granularity. Altogether, quality metrics for T2I systems are often based on averaged score of both compositional and overall quality measures as high compositional quality does not automatically result in high overall image quality and vice versa.

## 3 METRICS

Measuring the quality of an image has been an important research area in recent decades, especially in the field of computer vision, computer graphics and visualization [44], and it is referred to as image quality assessment (IQA). However, these metrics are designed to measure quality difference between an image  $x$  to a reference image  $y$ , where  $y$  functions as a ground truth (with perfect quality), which is commonly needed for image-to-image tasks, e.g., noise removal [45], [46] or super resolution [47], [48]. Since the advent of generative AI, initially mostly fueled by generative adversarial networks (GANs), researchers were striving for methods to measure the quality of their proposed image generators [49]. When image quality measures like peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM) [44] failed to capture “high quality”, “photo-realism”, or “cinematic” properties, the need for dedicated metrics arose. Initial work in this direction, resulted in the development of metrics that measure distribution fidelity, e.g., inception score (IS) [17] and Fréchet inception score (FID) [18]. These metrics are designed to compare the image distribution of a target image dataset to the image output distribution of an image generator.

However, they are based solely on images without adopting any text-condition. In the following, we discuss two types of text-conditioned image quality metrics, i.e., embedding-based and content-based quality measures and for each we provide details on the most adopted metrics in context of T2I synthesis. Afterward, image captioning metrics and closely related metrics, which have been utilized for evaluation of T2I alignment, are discussed. Finally, we summarize the field of pure image-based metrics, that have been adopted for assessing image quality for text-condition image generation.

See Table 1, for an overview of the reviewed metrics, where we compare metrics across four compositional aspects, cf. Section 2.3 and award points accordingly. Zero points are awarded when the metric is neither text-conditioned nor designed to capture the corresponding aspect. One point is awarded when the metric is text-conditioned but not explicitly trained to reflect compositionality. Two points are awarded for fine-tuning or other optimizations for certain aspects. Metrics assigned three points are specifically designed to reason about the aspect in question, such as object detection, segmentation, or visual question answering for specific aspects.

### 3.1 Embedding-based Metrics

Text-conditioned image quality assessment represents a novel and evolving paradigm in the field of T2I generation. In these approaches, the perceived quality of an image is



Taxonomy	Metric	Year	Cite/ Year	Fine Tuned	Compositional Ability				Human Evaluated	Rationale	
					Object Accuracy	Spatial Relations	Non-Spatial Relations	Attribute Binding			
Embedding-based	CLIPScore [50]	2021	137		●	●	●	●	✓	✗	
	BLIP-ITC [25]	2022	623		●	●	●	●	✗	✗	
	BLIP-ITM [25]	2022	623		●	●	●	●	✗	✗	
	BLIP2-ITC [26]	2023	840		●	●	●	●	✗	✗	
	BLIP2-ITM [26]	2023	840		●	●	●	●	✗	✗	
	MID [51]	2022	4		●	●	●	●	✓	✗	
	CLIP-R-Precision [52]	2021	20		●●	●●	●●	●●	✓	✗	
	NegCLIP [27]	2022	42	🔧	●●	●●	●●	●●	✗	✗	
	MosaiCLIP [53]	2023	1	🔧	●●	●●	●●	●●	✗	✗	
	CLoVe [54]	2024	1	🔧	●●	●●	●●	●●	✗	✗	
	PickScore [15]	2023	39	🔧	●●	●●	●●	●●	✓	✗	
	ImageReward [14]	2023	46	🔧	●●	●●	●●	●●	✓	✗	
	HPSv1 [39]	2023	9	🔧	●●	●●	●●	●●	✓	✗	
	HPSv2 [16]	2023	16	🔧	●●	●●	●●	●●	✓	✗	
	DreamSim [55]	2023	16	🔧	●●	●●	●●	●●	✓	✗	
	COBRA [56]	2024	1	🔧	●●	●●	●●	●●	✓	✗	
	R-Precision [57]	2018	231	🔧	●●	●●	●●	●●	✗	✗	
	RAHF [58]	2023	1	🔧	●●●●	●●	●●	●●	✓	🖼️	
Content-based	B-VQA [59]	2023	18		○●○●	○●○●	○●○●	○●○●	✓	✗	
	VISOR <sub>cond</sub> [28]	2022	13		○●○●	○●○●	○●○●	○●○●	✓	🖼️	
	PA [40]	2022	4		●	○●○●	○●○●	○●○●	(\$)	✗	
	CA [40]	2022	4		●●●●	○●○●	○●○●	○●○●	(\$)	✗	
	SOA [60]	2020	34		●●●●	○●○●	○●○●	○●○●	✓	🖼️	
	VISOR [28]	2022	13		●●●●	○●○●	○●○●	○●○●	✓	🖼️	
	VISOR <sub>N</sub> [28]	2022	13		●●●●	○●○●	○●○●	○●○●	✓	🖼️	
	TIAM [41]	2024	1		●●●●	○●○●	○●○●	○●○●	✓	🖼️	
	3-in-1 [59]	2023	18		●●●●	○●○●	○●○●	○●○●	✓	🖼️	
	ViCE [61]	2023	1		●●●●	○●○●	○●○●	○●○●	✓	✗	
	TIFA [62]	2023	25		●●●●	○●○●	○●○●	○●○●	✓	✗	
	VNLI [63]	2023	11	🔧	●●●●	○●○●	○●○●	○●○●	✓	✗	
	MQ [42]	2023	1		●●●●	○●○●	○●○●	○●○●	✓	✗	
	VQ <sup>2</sup> [63]	2023	11		●●●●	○●○●	○●○●	○●○●	✓	✗	
	DA-Score [64]	2023	2		●●●●	○●○●	○●○●	○●○●	✓	✗	
	LEIC [35]	2018	23		○●○●	○●○●	○●○●	○●○●	✓	✗	
	Image-Text	CIDEr [33]	2015	441		○●○●	○●○●	○●○●	○●○●	✓	✗
		TIGEr [36]	2019	12		●	○●○●	○●○●	○●○●	✓	✗
SPICE [34]		2016	211		●	○●○●	○●○●	○●○●	✓	✗	
T2T [65]		2023	84		●	○●○●	○●○●	○●○●	✓	✗	
ViLBERTScore [38]		2020	10		●●●●	○●○●	○●○●	○●○●	✓	✗	
VIFIDEL [37]		2019	7		●●●●	○●○●	○●○●	○●○●	✓	✗	
UniDet [59]		2023	18		●●●●	○●○●	○●○●	○●○●	✓	🖼️	
LLMScore [66]		2023	7		●●●●	○●○●	○●○●	○●○●	✓	✗	
VIEScore [67]		2023	2		●●●●	○●○●	○●○●	○●○●	✓	✗	
Distribution		IS [17]	2016	1092						✗	✗
	FID [18]	2017	1417						✗	✗	
	MiFID [68]	2021	6						✗	✗	
	KID [69]	2018	173						✗	✗	
	C2ST [70]	2016	44						✗	✗	
	PRD [71]	2018	76						✗	✗	
	CAS [72]	2019	33						✗	✗	
	DINO Metric [73]	2023	587						✗	✗	
	I-PRD [74]	2019	94						✗	✗	
	Single Image	GMM-GIQA [75]	2020	14						✗	✗
CLIP-IQA [76]		2023	39						✗	✗	
Aesthetic Predictor [13]		2022	462						✓	✗	
PAL4VST [20]		2023	2	🔧					✓	🖼️	
PAL4InPaint [19]		2022	3	🔧					✓	🖼️	
KPR [74]	2019	94						✗	✗		
PPL [77]	2019	1460						✗	✗		

🔧: text-based rationale, 🖼️: image-based rationale

TABLE 1

Comparative overview of text-to-image evaluation metrics classified according to our proposed taxonomy, indicated by color: blue for text-conditioned metrics and red for plain image-based metrics. This table also categorizes current state-of-the-art methods based on their ability to assess compositional alignment, their validation through human evaluation studies, and their provision of additional rationale beyond a mere quality score.

assessed not only based on its visual characteristics, but also in the context of accompanying textual information. However, existing image-only measures (Section 3.3) are unable to integrate textual cues that describe the content associated with an image. Text-to-image alignment acknowledges the intrinsic relationship between language and visual perception, allowing for a more nuanced evaluation that aligns with human judgment. In applications where text and image synergy is crucial, such as T2I synthesis, image captioning, content-based image retrieval, and human-computer interaction-based image generation, quantitatively measuring the alignment between text and image is mandatory. The incorporation of textual information introduces a dynamic dimension to image quality assessment, reflecting the evolving needs of multimodal systems and fostering advancements in the understanding and evaluation of visual content. In the following, we provide an overview of recent developments in quality assessment for T2I alignment.

One of the first reference-free approaches, that were used for measuring the distance between a textual and an image representation is CLIPScore [50], which is based on Contrastive Language-Image Pre-training (CLIP) [24]. The CLIP distance is computed through the cosine similarity between the text embedding vector and the image embedding vector. By pre-training on vast and diverse datasets, CLIP exhibits a remarkable capacity to generate meaningful and contextually rich embeddings for images and corresponding textual descriptions. CLIPScore was introduced as reference-free evaluation metric for image caption generation tasks together with its reference-based version RefCLIPScore.

Multimodal mixture of Encoder-Decoder (MED) was proposed by Li et al. [25]. It is used for multi-task pre-training and flexible transfer learning from image-text pairs collected from the web, and it is incorporated into BLIP, a framework which enables a wider range of downstream tasks such as T2I retrieval on COCO Captions [78] and Flickr30K [79], [80]. In the same manner as CLIPScore, the textual and image embedding vector, learned via image-text contrastive learning (ITC), returned by BLIP can be used to compute the cosine similarity and we refer to it as BLIP-ITC. Whereas, the image-text matching (ITM) version of BLIP is trained to learn a binary classification task where the model is asked to predict whether an image-text pair is positive (matched) or negative (unmatched).

One year later, BLIP is followed by BLIP2 in 2023, which is introduced as an efficient approach for vision-language pre-training, utilizing pre-trained image encoders and large language models (LLMs) with minimal trainable parameters. It sets new benchmarks across different vision-language tasks and exhibits advanced zero-shot capabilities for generating text from images. It utilizes a two stage pre-training of a query transformer (Q-Former), where the first stage bootstraps vision-language representation learning from a frozen image encoder. The second stage applies a frozen LLM for bootstrapping vision-to-language generative learning, enabling zero-shot instructed image-to-text generation. Again, the learned embedding vectors are utilized to compute an alignment score between text and image called BLIP2-ITC and identically to BLIP, there exists an image-text matching version, BLIP2-ITM.

Singh et al. [53] additionally employ scene graphs and

propose a graph decomposition and augmentation framework to learn text-image representation. They derive a pseudo image scene graph from the text caption by dividing the text-based graph into multiple subgraphs and matching them with the image. They further extend the common vision-language component of the loss by an image-to-multi-text loss, to train their model MosaiCLIP.

LXMERT was introduced by Tan et al. [81], a transformer consisting of three encoders which model separate tasks, namely object relationship, language and cross-modality encoding. The model takes as input text and image, but only uses the image encoding that belong to objects recognized by a separate object detection module. Following this, the relationship of the two modalities is learned in the cross-modality encoding through a cross-attention module. A disadvantage of their approach is the reliance on a separate object detection module, which can fail at correctly localizing all objects in the scene, therefore excluding information in the learning process.

The Unified Transformer (UniT), proposed by Hu et al. [82], learns multiple visual perception and language understanding tasks simultaneously. Each forward pass through the model has an input the text-image pair along with a task specific index, which the model embeds. The output from the text and image encoders are concatenated and further fed into a multi-head cross attention module. At the end of the process, the output representations are fed into task specific output heads. This design alleviates the need for task-specific finetune of the model, since it learns to jointly solve all tasks.

A mixture of four pre-training tasks in image and text space is proposed by Chen et al. [83] in their work on UNITER. As pre-training tasks, they use Masked Language Modeling with image information added, Masked Region Modeling with text information added, Image-Text Matching and Word-Region Alignment. They also combine four different datasets and train on the combined data. Similar to LXMERT, the image encoder is an object detector tasked with producing encodings of recognized objects. This comes with similar drawbacks of reliance on the detector. Gan et al. [84] extend the UNITER approach by employing large-scale task-agnostic adversarial pre-training along with task-specific adversarial finetuning, both on the embedding space. When combined with KL-divergence-based regularization, their model yields an embeddings space that is more invariant than the one of UNITER.

Zhang et al. [85] revisit vision-language models by improving individual components of an existing vision-language framework by Li et al. [86]. Their approach scales up the training data by combining four different datasets. With this, they improve the object detection model used as part of their approach, a design choice also made by UNITER and LXMERT, and show this improves the downstream vision-language representations. But contrary to UniT, their approach requires task-specific finetuning.

Vision-and-Language Transformer (ViLT) is proposed by Kim et al. [87] as a simplistic but effective method for vision-language pre-training. Contrary to the previously described approaches, their model does not operate on region features extracted from an object detection model. Instead, it uses a vision transformer encoder which produces representations

at patch granularity, alleviating the reliance on the detector for feature extraction. This design also reduces the model size and increases its speed. The approach also uses a shared transformer encoder with modality tokens for the model to differentiate between text and image input as well as separate token and patch positional embeddings respectively.

Li et al. [88] propose the Visual Semantic Reasoning Network (VSRN), which uses bottom-up attention to reason about relevant regions in the image. This information is further processed by a graph convolution network that generates features with semantic relationships. The output is further paired with the text encoding where the model jointly optimizes the encoding and text generation to learn the alignment of the modalities.

Another relevant metric for evaluating text-image alignment is R-Precision. Xu et al. [57] were the first to apply it for text-image alignment. For this, they aim to identify the top  $r$  relevant text captions for a given image with caption candidates  $R$ , to compute R-Precision as  $r/R$ . This is achieved by first extract global feature vectors from their pre-trained encoders for generated images and given text captions. The cosine similarity is computed between image and text vectors, and then used to rank the captions in descending similarity to identify the  $r$  most similar candidates. Park et al. [52] extend this approach by using CLIP as the encoder for images and text, and show this leads to a more human-aligned judgement and prohibits the bias that might come from using a custom model.

Kim et al. [51] propose Mutual Information Divergence (MID), a unified metric for multimodal generation, calculated through the negative Gaussian cross-mutual information between real and generated samples. Broadly formulated, their metric quantitatively measures how well one modality is aligned with the other, where both modalities are represented as encodings generated by their respective CLIP encoders. MID is shown to have consistent behavior across a variety of datasets where cosine-similarity-based techniques have shown weaknesses, especially for narrow domains like images of human faces.

Kirstain et al. [15] train a scoring function, PickScore, to estimate the user’s satisfaction for a particular generated image by finetuning CLIP-H using a large dataset comprised of generated images along with human preferences. Their objective function is designed to maximize the likelihood of a preferred image picked over an unpreferred one. Finally, they combine everything into a comprehensive benchmark they call Pick-a-Pic, which contains 500k examples with 35k distinct prompts that better represent what humans imagine to be contained in the image generation.

Xu et al. [14] build a T2I human preference reward model. To train the network, they gather and annotate 137k text-image pairs that go through a thorough annotation pipeline. The authors go one step further than PickScore and use the output of their reward model is further used to as a target to enable Reward Feedback Learning on top of a diffusion model to then optimize its image generations. With this, they are able to create a generative diffusion model that produces images which are more human-preference-aligned.

A similar approach is taken by Wu et al. in both their first [39] and second [16] iteration of their human prefer-

ence score (HPS). In the first version, they first collect a large dataset of human choices from images that have been generated by the same prompt, comprised of roughly 98k images with 25k prompts. Each prompt is assigned multiple image choices, for which the dataset contains annotations which image the user preferred. With this, they finetune a CLIP-L model a mechanism similar to the original CLIP training: The objective is to maximize the similarity between the prompt and the image chosen by the user, by minimizing the similarity to images rejected by the user. After training, the score is defined as the re-scaled cosine similarity between text and image embedding. Using their dataset, they also tune a diffusion model to generate images that are more aligned to human preferences by adding a preference identifier to the text prompt and incentivizing the model to output images in accordance to their HPS. The second version of HPS [16] introduces an even larger dataset of 798,090 human preference annotations for 433,760 image-text pairs. The HPS v2 finetuned CLIP model improves the human aligned scoring of the first version. Further, it is demonstrated to be sensitive to algorithmic improvements of the underlying T2I models.

Fu et al. [55] propose DreamSim, an extensive benchmark for evaluation of generated images w.r.t. human preference alignment. Their dataset is composed of 20k synthetic image triplets with a reference image as well as two other images, where the user decided which is more similar to the reference. Their dataset covers various aspects of similarity, such as pose, perspective, foreground color, number of items, and object shape. Using this dataset, they learn their perceptual metric using an ensemble of networks to encode each of the triplet images, calculate the cosine similarity between each image to the reference, followed by a triplet loss. They show their learned network is able to make more human-aligned judgements, compared to e.g., CLIP. On the other hand, previous methods did not rely on an ensemble configuration, increasing the computational cost of DreamSim.

In DreamBooth [73], a combination of three metrics is used to evaluate a multi-view generation of an object. To assess the image quality, they compare a generated image with the ground truth image from the same view using cosine similarity of CLIP [24] and DINO [89]. Hereby the CLIP-based metric only requires the images to show the same subject to return high similarities, whereas the DINO-based metric was included to measure more fine-grained differences. Lastly they use the cosine similarity of CLIP embeddings of the generated image and the according text prompt to measure prompt fidelity.

### 3.2 Content-based Metrics

Content based metrics evaluate the generated image directly based on its content, rather than the images projection into an embedding space (see Section 2.2.1). This also allows for a decomposition of the evaluation of single aspects of the image quality like object accuracy (OA), spatial relationships (S), non-spatial relationships (NS) or attribute bindings (AB). In the following we will discuss multiple content-based metrics.

### 3.2.1 Text-Image Content Matching.

The SeeTRUE(VNLI) metric was introduced by [63]. It relies on fine-tuning multimodal models (like BLIP2 [26], PaLI-17B [90]). The training dataset consists of 110K text-image pairs labeled with binary alignment annotations. The model receives the image as well as a text prompt consisting of "Does this image entail the description: prompt?" The model is trained to answer the question with "yes"/"no". During inference the relative ratio between predicting "yes" or "no" is used to compute the alignment score. The more "yes" is predicted by the finetuned model, the better the alignment is assumed to be. A major limitation of this finetuned approach is its black box nature, making improvements of the validated model difficult as well as trust in the evaluation hard.

Mismatch Quest [42] especially tackles this issue by introducing an end-to-end trainable approach for providing visual and textual feedback in T2I models, aiming to identify and explain alignment discrepancies between generated images and textual prompts. They introduce a method for generating a comprehensive training set that includes aligned image-text pairs and negative examples with misalignment. This is achieved by leveraging LLMs and visual grounding models to create synthetic examples with misalignment, based on aligned image-text pairs, and corresponding textual and visual feedback. In detail, to generate this dataset the authors use Part-of-Speech (POS) Tagging by labeling each word with its grammatical category (such as noun, verb, adjective). Based on the POS tags misaligned text-image pairs are then generated based on aligned ones. The authors use this approach to generate a training set (TV-Feedback) to finetune feedback models in such way that they are able to provide visual (bounding box) as well as textual explanations for misalignment. For evaluation the authors introduce a SeeTRUE Feedback dataset with 2,008 human annotated instances highlighting textual and visual feedback. The authors also show in a study that the proposed model aligns well with human judgements. However, the proposed model possibly struggles with the detection of multiple misalignments.

In difference to the fine-tuned metrics of Mismatch Quest and SeeTRUE(VNLI), other metrics are utilizing VQA models to generate an evaluation score for generated images. They especially utilize these VQA models for evaluation disjoint parts of the image prompt, making them belong to the category of compositional metrics.

Yarom et al. [63] propose the SeeTRUE(VQ<sup>2</sup>) metric: First, the authors extract answers from the image prompt using SpaCy [91], which is based on POS and dependency parse tree annotations. For each answer  $a_j$  a question  $q_j$  is generated leading to a set of question-answer pairs. Next, the generated pairs are reformulated to match yes-no questions, following the schema of "is  $a_j$  true for  $q_j$  in this image?". The assumption is that if the VQA model is able to answer all question answer pairs with yes, the semantic alignment is perfect. Hence, the final score is computed by an average over all 'yes' answer probabilities.

Similar to the SeeTRUE(VQ<sup>2</sup>) metric, Compositional-Alignment-Score (DA-Score), by [64], decomposes the image prompt to evaluate the T2I alignment. The authors found that misalignment's are often left undetected by pre-trained

multimodal models, such as CLIP, especially when there are more complex prompts. Hence, DA-Score decomposes the prompt into a set of disjoint assertions using a LLM. Each assertion is then evaluated individually using a VQA model (BLIP). Finally, the scores are combined to give the T2I alignment score. The individual evaluation of the assertion allows to draw conclusions about the strength and weaknesses of the generative model. Moreover, it allows the optimization during the forward diffusion process by iteratively increasing the cross-attention strength of low scoring assertions. The authors show a higher correlation with human ratings over prior evaluation metrics, like CLIP [50], BLIP [25] and BLIP2 [26] especially when the prompt becomes more complex.

However, in order for the VQA based compositional metrics like SeeTRUE(VQ<sup>2</sup>) or DA-Score to be applicable to the defined image aspects like object accuracy (OA), spatial relationships (S), non-spatial relationships (NS) or attribute bindings (AB), thoughtful prompt-engineering is required, which means that a prompt needs to capture different image aspects. Selection of an adequate dataset for evaluation is hence crucial. In our supplementary material, we review existing text-image datasets and provide a comparison between their different levels of compositionality.

Hinz et al. [60] propose the usage of their Semantic Object Alignment (SOA) metric which is able to specifically address the challenges posed by multi-object and other complex scenes represented by the generated images. The authors propose to use a pre-trained detection model to infer the existence of prompted objects in an input image. To do so, the authors sample image captions from the COCO validation set, explicitly mentioning one of the 80 main object categories. They use a pre-trained object detector to check if the generated images contain the objects specified in the captions. The paper conducts a user study comparing several T2I models using SOA and demonstrates that SOA aligns well with human rankings. In contrast, other metrics like Inception Score do not exhibit the same alignment.

In the work of Grimal et al. [41], they follow a similar approach introducing their Text-Image Alignment Metric (TIAM). It captures the alignment between image and prompt by characterizing the contents of a set of generated images based on a pre-trained segmentation model. For generating prompts the authors propose to follow a simple template that is enriched by word labels (like "elephant" or "cat") and additionally optional attributes (like color). For the evaluation of the attribute color, the authors propose to count a correct assignment when at least 40% of the attribute color is detected within the segmentation mask of the object. The authors propose the usage of YoloV8 which was trained on 80 COCO classes. In order to make the segmentation model perform well on the generated images, they require to be realistic looking, as COCO was used for training YoloV8. Hence, to bias the generator towards the generation of realistic looking images, the authors propose to use a text prompt, which starts with "a photo of". They also evaluate the correlation to human evaluation and show their superiority compared to [24] and [25].

Limitations of the two previous works [41], [60] include their applicability to generative models which are trained to produce different styled images like cartoons or sketches.



Additionally, the pre-trained models are limited to detecting classes which have been presented during training, and evaluation of domain specific generative models is thus not possible. Additionally, another significant drawback of metrics relying on pre-trained models is the potential overlap between the models used for generation and evaluation. A recent study highlighted this issue [40], revealing that SOA employs the same pre-trained detector as CPGAN [92] during image generation. This overlap can result in overfitting during evaluation, giving CPGAN an unfair advantage (even over real images) when assessed using the SOA metric. To address this problem, an effective solution is to swap the detection model used in the evaluation process.

The authors of VISOR [28] found that many existing models struggle with the challenge of generating multiple objects, and even when successful, they often fall short in capturing spatial relationships described in the input text prompts. They propose three variants of the VISOR metric: VISOR, VISOR<sub>N</sub>, and VISOR<sub>cond</sub>. These metrics first rely on detecting objects that have been mentioned in the text prompt using a pre-trained object detector and the centroids of detected bounding boxes for deriving depicted relationships. The VISOR metric returns 1 if all objects are present in the image with correct spatial relationships, otherwise it returns 0. VISOR<sub>N</sub> adopts a distribution-based approach, assessing the model’s capability to generate at least  $n$  spatially correct images based on the VISOR score for a given text prompt mentioning spatial relationships. Finally, VISOR<sub>cond</sub> evaluates the conditional probability of generating correct spatial relationships, given that all objects are generated accurately. This means that the object accuracy does not influence the VISOR<sub>cond</sub> metric.

[40] also propose a metric for evaluating the Positional Alignment (PA). The metric evaluates how well the generated images align with the positional information conveyed in text descriptions. Therefore, they define a set of positional words ( $W$ ) that convey positioning information, such as “above,” “below,” “on top of,” etc. For each word in  $W$ , the authors filter captions in the evaluation set of the COCO dataset containing that word to create a matched caption set ( $Q_w$ ). A mismatched caption ( $P_w$ ) is then created by replacing the positional word with its antonym while keeping other words unchanged. For each word in  $W$ , a set ( $D_w$ ) is created, including triplets ( $R_{wi}, P_{wi}, Q_{wi}$ ), where  $R_{wi}$  is a generated image from the matched caption  $P_{wi}$ . The authors use these sets to query the input captions from the binary query set [ $Q_{wi}, P_{wi}$ ], marking a query as successful if the matched caption is correctly identified using CLIP.

Additionally, the authors introduce Counting Alignment (CA), which quantifies how accurately a T2I synthesis model aligns with counting aspects in input text descriptions, specifically focusing on the number of objects in generated images. The evaluation process begins by constructing a test dataset, filtering captions from the MS-COCO validation set that mention counting using terms like “a”, “one”, “two”, etc. ground truth counting information is annotated for each selected caption, emphasizing countable object types to avoid penalizing uncountable categories. Text-to-image models then generate images from each caption in the counting test set. An off-the-shelf object counting model is utilized to count objects for each class in the generated

images. The CA value is computed by comparing predicted and ground truth counts, measuring counting errors with root mean squared error, and averaging over all test images and object classes.

[40] state that the evaluation of T2I generation should be evaluated based on a variety of criteria. They hence propose a framework to combine multiple aspects for the image evaluation which they call “bag of metrics”. In their study they are able to show a more consistent ranking with real images and human evaluation, using the proposed “bag of metrics”. The authors further provide a python package called TISE for direct application.

In a similar vein, [59] put forward the innovative 3-in-1 metric, designed specifically to assess attribute bindings, spatial relations, and non-spatial relations, such as “look at,” “hold,” and “play with,” in text-to-image generation models. As its name implies, this metric amalgamates three distinct evaluation criteria to comprehensively analyze the entirety of image content. To evaluate attribute bindings, the authors introduce the “Disentangled BLIP-VQA” approach, recognizing that conventional VQA-based assessments often yield inaccurate results due to the model’s inability to discern correct object-attribute relationships. Consequently, they decompose complex prompts into independent questions, ensuring each question pertains to only one attribute-object pair, thus circumventing confusion in the VQA model’s understanding. For assessing spatial relations, they employ the UniDet model to evaluate relations like “next to,” “near,” “on the side of,” as well as directional relations such as “left,” “right,” “top,” and “bottom.” Finally, they employ the CLIPScore [50] to evaluate non-spatial relations, completing the triad of evaluations in the 3-in-1 metric framework.

In their paper Yuksekgonul et al. [27], the authors aim to elucidate how Visual Language Models (VLMs) encode the compositional relationship between objects and attributes. To achieve this goal, they introduce the Attribution, Relation, and Order benchmark. This benchmark evaluates the VLM’s comprehension of object properties and relations using the Visual Genome Attribution and Visual Genome Relation datasets, respectively. They evaluate the order sensitivity using COCO [93] and Flicker30k [94]. The authors emphasize a critical issue regarding contrastive pretraining in VLMs, which tends to prioritize learning low-level features over higher-level compositional structures. To tackle this challenge, the authors propose composition-aware hard negatives, which they integrate into CLIP’s contrastive objective [24]. These hard negatives are generated by altering linguistic elements such as nouns and phrases in negative captions. During training, when assembling a batch of images and their corresponding captions, the authors include not only the original images but also strong alternatives. Through their evaluations, the authors assert that integrating the proposed alternatives improves VLMs’ comprehension of composition and order.

The Visual Instruction-guided Explainable Score (VI-EScore) proposed by Ku et al. [67] is composed of the perceptual quality (PQ) and the semantic consistency (SC) score. Both scores are based on instructing a LLM using handcrafted prompt templates to reason about a given image. In their experiments LLaVA [95] and GPT-4v [29]

are utilized for image analysis. Perceptual quality is measured by two sub-scores, the first is measuring image naturalness and the second measures degree of distortions / artifacts. Note, that PQ is based on the image only, the VLM is not informed about the prompt. Also, SC is composed of sub-scores, which depend on the evaluation objective. SC is demonstrated for 5 objectives: text-guided image generation, text/mask-guided image editing, control-guided image generation, subject-driven image generation and subject-guided image editing.

The first method leveraging LLMs for automatic T2I evaluation is called LLMscore and was proposed by Lu et al. [66], which is applied in an image-to-text manner. First, they utilize BLIP2 for image captioning and producing a global image description followed by an object-centric local reasoning. Therefore, Grit [32] detects image crops of objects and generates a textual description of that particular region. In order to fuse global and local text description, GPT-4 [29] transforms them into an object-centric visual description. Finally, the evaluation objective of LLMscore can be retargeted and in their work Lu et al. demonstrate overall score and error counting objectives. The visual description together with the evaluation instruction are forward to GPT-4 and the final LLMscore with rationale is returned. However, LLM-generated captions may contain additional details fabricated by the LLM rather than originating from the image captioning process, which may lead to not sufficiently incorporating requirements and inputs from the original prompt.

Visual concept evaluation (ViCE) is a metric that was designed to possess an understanding of visual concepts similarly as humans do, who are able to directly generate the visual concepts once they receive a prompt to inspect. Similar to other VQA-based methods, ViCE initially generates questions-answer pairs using GPT-3.5-turbo [29] based on the prompt. An initial set of 15 questions are questioned to an LLM in order to parse visual concepts from the prompt. A unique feature of this method is, that the model may seek additional information to refine its understanding of the image. Therefore, after the initial response phase, the LLM is asked in an iterative process whether it requires further information until the model is satisfied with its comprehension of the image. In such a way, the model is able to validate whether if the objects are in a correct semantic relationship. Finally, the actual visual image analysis is then performed by a BLIP2-based VQA model assessing the image based on the previous generated question-answer pairs.

In the work of Hu et al. [62] they propose a metric called TIFA leveraging VQA models to measure the faithfulness of a generated image. To do so, they generate multiple-choice question-answer pairs utilizing GPT-3 [96] via in-context learning and apply verification of generated questions using a multi-task question-answering model called UnifiedQA [97]. TIFA adopt an open-domain pre-trained vision-language models (the authors recommend to use mPLUG-large [98]) as VQA model, rather than closed-class classification models fine-tuned on VQAv2 [99] enabling it to perform well on diverse set of visual elements. However, limitations of TIFA is the dependency to 12 categories: object, activity, animal, food, counting, color, material, spatial, location, shape, attribute, and other, which are considered to generated question-answer pairs.

### 3.2.2 Image-Text Content Matching.

Finally, delving into the evaluation of T2I alignment, one needs to look into its inversion. In the time before the invention of CLIP, T2I alignment was measured in reversed manner, the so called image-to-text alignment, which is a closely related task found in the field of image caption generation (or image description generation), where for a given image a generated textual description needs to be evaluated. There are several datasets like Flickr8K [100], Flickr30K [79], MS-COCO [78], [93] or Pascal 50S [33], that provide human ratings of captions for a given image functioning as benchmark datasets for image to text evaluation, and several visio-linguistic metrics were developed. In the following, we provide an overview of image captioning metrics and machine translation metrics that shaped the field of text-conditioned image synthesis evaluation before compositional quality metrics emerged.

The SPICE (Semantic Propositional Image Caption Evaluation) metric [34] is also developed in the context of image captioning evaluation, while focusing on the semantic content of the generated text. SPICE first parses both the generated sentence and the reference sentences into scene graphs that capture the underlying meaning of the text in a structured form. These scene graphs represent objects, attributes of objects, and relationships between objects mentioned in the captions. To compute the SPICE score, the candidate and reference scene graphs are compared to assess the semantic accuracy of the generated sentence with respect to the ground truth. This comparison is achieved through an F-score.

In the work of Cui et al. [35] an image captioning metric is proposed, that we referred to as LEIC (Learning to Evaluate Image Captioning). It is a discriminative evaluation metric designed to distinguish between human-generated and machine-generated sentences. The metric utilizes a CNN-based network for image encoding and an LSTM for text encoding. A binary classifier acts as the critic, assessing the quality of the generated sentence in relation to the image and reference sentences. This approach is assumed to help the metric mimic human judgment more closely than traditional metrics.

TIGER (Text-to-Image Grounding for Image Caption Evaluation) [36] is a metric specifically designed for the evaluation of image captioning models. Unlike purely lexical based methods, TIGER also takes the image content directly into account. This is done via a process called text-image grounding. TIGER compares text and image representations in vector space. In their ablations the authors show that TIGER achieves a higher correlation with human judgement than BLEU, ROUGE and METEOR.

Moreover, there are several metrics based on the comparison between detected objects in the image and the textual description of such in an image description [67], [101]. Visual fidelity of image descriptions is measured by VIFIDEL [37], which is based on the word's mover distance (WMD) [102], measured between objects in the image, and its corresponding description in the text. VIFIDEL is specified by the inverse of the minimum cumulative cost required to move semantic labels (e.g., object categories) from an image to words in the description. This converts WMD from a distance measure to a similarity measure. Category names

of detected objects in the image are compared to the words from the description. Additionally, object importance can be applied to weight more frequently used words in the reference descriptions.

In the work of Lee et al. [38] they propose a metric called ViLBERTScore, which is similar to BERTScore [103] that computes textual embeddings for a reference and a generated caption. Additionally, the computation of textual embeddings is conditioned with the target image using the model proposed by Lu et al. [104]. Hereby, contextual embeddings are computed by applying an object detector to the target image and feeding pairs of image region features and text embeddings to the pre-trained ViLBERT model. Finally, the ViLBERTScore is defined by the cosine similarity between reference caption embeddings and candidate caption embeddings.

**Text-to-text metrics.** Following, several machine translation metrics, which are frequently used together with image captioning models for evaluation of text-image retrieval tasks. The CIDEr (Consensus-based Image Description Evaluation) score [33] is based on measuring the similarity of a generated sentence to a set of reference sentences written by humans for a given image. It aims to capture how closely the generated description matches the consensus of what most people would say about the image. It applies Term Frequency-Inverse Document Frequency (TF-IDF) weighting, which gives higher importance to  $n$ -grams that are unique to the particular image and less importance to  $n$ -grams that are common across many images in the dataset, and calculates the cosine similarity between the TF-IDF vectors of the candidate sentence and each of the reference sentences. Scores calculated for  $n$ -grams of different lengths (from 1 to 4) are combined to capture a wide range of alignment, from exact word matches to more complex grammatical structures and semantic alignments. The final CIDEr score for a candidate sentence is an average of its similarity scores across all reference sentences, normalized by the number of reference sentences.

The BLEU (bilingual evaluation understudy) metric [105] was initially used to evaluate the performance of machine translation systems. BLEU measures the precision of  $n$ -grams between a machine generated candidate translation and a corpus of high quality human translations. It is important to know that only the count of matching  $n$ -grams is taken into account and the order of  $n$ -grams is ignored. The BLEU metric is always between 0 and 1, where a score of 1 would indicate that the candidate translation would directly match a translation from the text corpus. Therefore, it is highly unlikely to achieve a score of 1. Nevertheless the BLEU metric has been shown to have a high correlation with human judgments. Because of this and because it is an inexpensive metric it still remains popular to date.

ROUGE [106] is a software package as well as a set of metrics that was originally intended for the evaluation of text summaries. The paper proposes multiple metrics that show high correlation with human judgments. ROUGE-N is very similar to the BLEU metric but instead of precision it measures recall between  $n$ -grams. The other versions incorporate longest common sequences (LCS), weighting consecutive LCSes and co-occurrence of skip-bigrams.

METEOR (Metric for Evaluation of Translation with Ex-

PLICIT ORDERING) [107] is another machine translation metric that tries to improve on the BLEU metric. It works by matching unigrams between a candidate machine translation and a corpus of reference translations. As has been shown in previous work [108], including the recall and not only using precision yields a higher correlation with human judgement. Additionally, METEOR uses different modules where the unigram matches are not only done based on direct matches but also on stemmed versions and synonyms. The authors of the METEOR paper show in an ablation study that METEOR achieves a higher correlation with human judgement than the BLEU metric.

BertScore [103] is a metric used to evaluate the quality of machine-generated text. It is based on BERT (Bidirectional Encoder Representations from Transformers) [109] and computes the cosine similarity of contextualized word embeddings from generated and reference obtained from BERT. This method allows BertScore to capture not only the semantic similarity between words but also their context within the sentence. BertScore has been shown to correlate well with human judgments of text quality.

### 3.3 Image-based Metrics

#### 3.3.1 Distribution of Images Metrics

One set of popular evaluation metrics assumes the generative model to be a black box and operates only on samples of the generated distribution  $q$  and compares it to samples from the target distribution  $p$ . The most commonly used metrics then rely on comparing features produced by pre-trained neural networks. Inception Score (IS) [17] uses an Inception network pre-trained on ImageNet to compare class predictions for a set of generated samples  $x \sim q$ . Hereby the score rewards low entropy in class predictions  $p(y|x)$ , i.e., generated images that can be clearly classified as one of the classes, as well as high entropy in marginal class distribution  $p(y)$ , i.e., a large diversity among generated samples. Due to its short-comings [110], the IS has recently lost popularity. The MODE score [111] improves upon the IS by adding another term that rewards similar distribution of class predictions for the generated and target images. IS-based metrics are not suited for T2I, as the marginal class distribution  $p(y)$  is typically not available.

The Fréchet Inception Distance (FID) [18] compares the means and co-variances of the features, extracted by the Inception network from samples of the generated and target distributions, using the Fréchet (or Wasserstein-2) distance. FID proved to be a more consistent quality measure than IS and is still commonly used. MiFID [68] extends FID by incorporating a term that penalizes memorization of training samples, by computing the minimum cosine distance of Inception features to the training dataset. This penalty was introduced to avoid bogus submissions for an image generation competition.

All the Inception-based metrics share the downside of relying on the weights of the Inception network. Those weights are the result of supervised ImageNet classification training and many of the Inception metrics are not robust to different sets of weights obtained from similar trainings [110]. Further, with the increasing scale of modern T2I models and datasets [13] far beyond the ImageNet domain,



features trained to classify this comparatively narrow domain may be insufficient for quality assessment. Adoption of a more capable and general feature extractor, such as semi-supervised models, could improve the reliability of metrics like FID, especially for models exceeding the ImageNet domain.

The distributions  $p$  and  $q$  can also be compared using kernel embedding of distributions, such as with the popular maximum mean discrepancy (MMD) metric, which measures the distance between kernel embeddings of samples. Unlike information theoretic approaches, MMD has the benefit of not requiring density estimation or bias correction. Most implementations use a fixed kernel for MMD, which introduces problems with complex natural images. The parzen window estimate [112] can be seen as an MMD approach. The Kernel Inception Distance (KID) [69] computes the squared MMD between Inception representations, eliminating FID’s bias for the amount of samples.

Another common means of comparing two distributions is using two-sample tests. Lopez-Paz and Oquab introduce C2ST [70] and propose using a binary classifier to distinguish between samples of the generated and target distribution. This binary classifier should achieve an accuracy of  $\approx 50\%$  for large numbers of samples. Using a nearest neighbor classifier allows for further insights in the generated data. For example, if the majority of nearest neighbors for generated images are also generated images, this indicates possible mode collapse. C2ST can be used with both, nearest neighbor classifiers and neural network based classifiers, including such based on pre-trained feature extractors like ResNet-34 [113].

The previously mentioned image-based metrics reduce the problem of evaluating generated images to a scalar score. Sajjadi et al. [71] find that such evaluation benefits from a separate notion of precision and recall for the distributions (PRD). Hereby precision measures the fraction of generated images that are in the support of the target distribution  $p$ . Recall measures the fraction of real images that are in the support of the generated distribution  $q$ . Sajjadi et al. cluster Inception embeddings of samples from  $p$  and  $q$  using  $k$ -means, before comparing the histogram over cluster assignments. Finding meaningful clusters that are mostly occupied by only generated samples would therefore decrease precision and clusters mostly occupied by target distribution samples decrease recall. They use multiple randomized clusterings to compute final precision and recall of the distributions. Kynkäänniemi et al. [74] further improve on this idea by changing how the support of a distribution is estimated (I-PRD). While Sajjadi et al. reduce this estimation of support to a one dimensional histogram comparison, the improved method explicitly models the support manifold. The support is estimated by placing hyperspheres around each of the sample embeddings with a radius equal to the distance to its nearest neighbor within the distribution. After finding the portion of samples from  $p$  and  $q$  lying in each others’ supports, the precision and recall can be directly computed.

In the work of Ravuri et al. [72] the Classification Accuracy Score (CAS) is proposed. It is based on predictions for real images of a ResNet image classification model trained on synthetic data. The performance accuracy for the set of

real images is referred to CAS, and it is demonstrated that CAS can identify classes for which a GAN failed to correctly learn its data distribution.

### 3.3.2 Single Image Metrics

Gu et al. [75] propose GMM-GIQA, which models the embeddings of the target distribution  $p$  using a Gaussian mixture model. A generated image can then be assigned a score based on the probability density of its embedding. The authors note however, that the metric may fail for too complex distributions, as they cannot be sufficiently modeled using a Gaussian mixture model.

With CLIP at the center, Wang et al. [76] propose the CLIP Image Quality Assessment (CLIP-IQA) benchmark. In their work, they improve CLIP’s ability to assess text-image alignment through antonym prompt pairing and removing the positional embedding from the image encoder. The resulting model is significantly better for evaluating quality and abstract perception.

With the introduction of the LAION Aesthetics dataset [13], the authors trained models<sup>1</sup> to predict how aesthetic humans would rate a given generated image, resulting in an image quality metric that is aligned with human preferences.

Zhang et al. [20] collect a dataset containing human annotated segmentation of artifacts. They then train binary segmentation models to automatically detect such artifacts in images. They further propose a related metric for evaluating in-painting using generative models, the Perceptual Artifact Ratio (PAR) [19], also known as PAL4InPainting, which measures the relative area occupied by artifacts. This metric is also generally applicable to full images, not just regions for in-painting.

Since the I-PRD method yields only a binary result for an individual sample, Kynkäänniemi et al. [74] propose a variant, KPR, which estimates how close the feature vector of a single image is to the feature vectors of  $k$ -NN real images.

Karras et al. [77] introduce the perceptual path length, a metric for latent variable models. The idea is to pair-wise compare subsequent images in a latent space interpolation using a perceptual image metric. This metric measures if any drastic changes appear for close latent codes and rewards smooth transitions within the interpolation, which is an indicator for a good disentanglement.

## 4 OPTIMIZATION

While T2I quality metrics are developed in order to judge the quality of single T2I samples or entire T2I models, other techniques can also be used in order to optimize the synthesized output. In this section, we will review some of these techniques, whereby we differentiate between those which require a training process (see Section 4.1), and those which can be used without additional training (see Section 4.2).

### 4.1 Finetuning Image Generators

The StyleT2I framework [114], which utilizes a CLIP-guided contrastive loss, a semantic matching loss, and a spatial

1. <https://github.com/christophschuhmann/improved-aesthetic-predictor>



constraint to refine attribute manipulation within intended spatial regions, trains a StyleGAN [115] model to increase compositional accuracy. This approach ensures more disentangled latent representations, that can be decoded into a high-fidelity image aligned with the input text. In order to find optimal latent codes, a Text-to-Direction module is employed to predict the sentence direction that is aligned with the input text, which is trained using a CLIP-guided contrastive loss. To enhance attribute alignment, a *Attribute-to-Direction module* gets optimized by the semantic matching loss that strives to identify attribute directions of the latent codes. To mitigate the issue of changing multiple regions during attribute alignment, a spatial constraint avoids spatial variations outside of a pseudo-ground-truth mask generated by a segmentation model.

In the work of Dong et al. [116], called RAFT, a Stable Diffusion model is fine-tuned on high quality samples ranked by a reward model. In a decoupled data generation process high-quality training data is sampled from the generator by discarding low-quality data points, which is done by utilizing a reward model to filter out those that exhibit undesired behavior. In their experiments they adopt CLIPScore and LAION Aesthetic Predictor as reward models, however fine-tuning on training data produced by the generator hinders the model to overcome problems within distribution towards high-quality compositionality, which remains concealed outside the generator’s distribution. After training a reward model utilizing their proposed dataset, see Section 3.2, Liang et al. [58] demonstrate to fine-tune a Muse image generator and comparing it to the pre-trained Muse version based on 100 generated test prompts. In their experiments, they utilize Muse to generate a set of 100, 512 images for 12, 564 generated prompts, then they apply their reward scores to filter out images below a certain threshold, and fine-tune Muse. Finally, they quantify the gain from Muse finetuning by conducting a study, where they present the two images side-by-side originating from the baseline and finetuned model. They were able to show, that the plausibility finetuned model produces significantly fewer artifacts/implausibility than the original Muse. A similar approach by Lee et al. [117] collects binary decisions from human annotators for a large set of synthetic images to train a reward model. Thus, they are able to finetune Stable Diffusion via reward-weighted likelihood maximization to better align it to human feedback using 27K image-text pairs. Also, AlignProp [118] is following the idea of utilizing a reward model to supervise finetuning of Stable Diffusion with human feedback.

After such recent developments of reward models opening a way towards incorporating human feedback into the diffusion process, further approaches adopted reward models, applying supervision during reinforcement learning. Thus, Fan et al. [119] propose *DPOK*, diffusion policy optimization with KL regularization, which utilizes KL regularization to stabilize RL finetuning and aligning T2I. Unlike traditional supervised finetuning, which often degrades image quality, the RL finetuning with *DPOK* optimizes ImageReward [14] a feedback-trained reward model online, leading to better alignment between text prompts and generated images while maintaining high image fidelity. The paper demonstrates that *DPOK* outperforms super-

vised fine-tuning methods in experiments, showcasing its effectiveness in enhancing T2I diffusion models. However, their work studies KL-regularization and primarily focuses on training a different diffusion model for each prompt. Denoising diffusion policy optimization (DDPO), proposed by Black et al. [120] broadens the approach by training using multiple prompts and showcasing the model’s ability to generalize to unseen prompts. The adaptability of DDPO to various reward functions, including those derived from vision-language models, marks its advance in enhancing prompt-image alignment beyond the scope of human feedback optimization.

## 4.2 Training free Optimization

In their work, Liu et al. [121] introduce an approach called Composable Diffusion Models, targeting the challenge of accurately align compositional text prompts to their image representation. Their method proposes a structured strategy where an image is generated through the composition of a set of diffusion models, each modeling different visual concepts. By treating diffusion models as energy-based models, they enable the explicit combination of data distributions defined by these models. This approach assumes the visual concepts are conditionally independent given the image. Sampling from the resultant distribution involves using a composed score function that integrates the contributions of each concept to the denoising process, allowing for the generation of images that faithfully represent the composed concepts. This allows for the generation of scenes that are more complex than those encountered during training, effectively combining object relations, and attributes accurately. Feng et al. [122] introduce *Structured Diffusion Guidance*, a method aimed to enhance compositional T2I synthesis through the use of scene graphs, which are derived from the prompt. Instead of computing a text embedding for a single sequence, this method extracts noun phrases from the prompt corresponding to visual concepts and entities, and encode such noun phrases separately to achieve region-wise semantic guidance. Finally, for each cross-attention map the average of all noun phrase activations denotes the corresponding output. The method *Attend and Excite* by Chefer et al. [65] proposes a training-free concept, which they call *Generative Semantic Nursing* (GSN) exploiting cross-attention maps of a Stable Diffusion model [123] to incorporate all subject tokens in the prompt and strengthen their activations, ultimately steering the model to better represent objects and their attributes and relations in the image. GSN is applied during inference, in particular at each denoising step attention values are maximized at the subject’s corresponding region in the attention map. This results in following denoising steps to better incorporate subject tokens into the latent representation. The maximization of attention values is based on a loss objective that enforces high activations for at least one patch in the attention map per subject token. However, the authors observe, that the timing of the optimization is crucial, since during the final timesteps of denoising the spatial locations of objects in the generated image do not alter anymore. Also, at each denoising step an iterative latent refinement is adopted, and the authors show that a gradual refinement helps to prevent frequent updates

of latent updates which would lead to degraded images. While the proposed strategy provides promising results for simple prompts, its efficacy declines for prompts with increasing complexity, e.g., multiple entities with bound attributes. A follow-up work, named *Divide and Bind* [124] by Li et al. , picks up the idea of *GSN* and applies total variation maximization, which opens the room for activated regions and increases the amount of local changes in the attention maps. As a result, diverse object regions are encouraged to emerge, enabling concurrently competing objects during the generation process. Therefore, they introduce two novel objectives, one that attends to the object tokens and the other for attribute binding regularization. The formulation of the attendance loss is based on the finite differences approximation of the total variation along the spatial dimension, which enables activations of neighboring locations in the attention maps resulting in activation patterns effectively dividing the image in to different regions. In order to realize attribute binding, object tokens and attribute tokens, which tend to share large overlapping regions in the attention map, are normalized. Thus, the normalized attention maps are considered as two probability mass functions, whose symmetric similarity is then maximized by minimizing the Jensen-Shannon divergence [125]. The metric proposed by Singh et al. [64], see Section 3.2 for further details, is able to detect parts of the image that are not aligned with the prompt, and when passing relative importance of such parts to a modified reverse diffusion process, it becomes possible to improve T2I alignment. In their experiment, they can show that by incorporating such relative importance in the form of weighting factors through a combination of prompt weighting and cross-attention guidance (*Attend and Excite*), they can optimize T2I alignment. Layout control with cross-attention guidance is achieved by Chen et al. [126], by introducing two techniques to incorporate object bounding boxes to steer the diffusion process optimizing for spatial relation alignment. This method enables to control image generation by providing object bounding boxes to encourage the diffusion model to generate corresponding objects within the bounding box region of the image. The work proposes two ways of layout guidance: forward and backward guidance. The former method applies a smoothed window function to the cross-attention maps, which increase activation for corresponding object tokens inside the bounding box. The second method, backward guidance applies an energy-loss function which is computed via back-propagation to update the latent vector and therefore indirectly alter the cross-attention maps. During image generation they alternate between denoising steps and gradient updates. While this technique improves overall layout and provides control over spatial relations, problems of standard diffusion approaches for object attribution remain and the question of manual bounding box placement needs to be considered.

## 5 CHALLENGES

Measuring text-image alignment focuses on relations between objects described by a text. That includes spatial and non-spatial relations between objects and their bound visual attributes. However, these alignment-focused metrics are targeted to sense the presence or absence of certain

compositions in image space, but are unable to quantify the quality of such detected components (if present). Quality scores provided by VLM-based metrics are defined on their visio-linguistic capability providing quantitative reasoning in form of class probabilities for *Yes* or *No* answers of closed questions. However, such probability score merely indicates the degree of uncertainty rather than actual alignment quality. Future measures should be designed to compute quantities for detected compositions enabling to rank alignment quality on a component-level rather than on a basis of uncertainty scores. Also, VLMs tend to behave like bags-of-words [27], which is a phenomenon that describes a model’s insensibility of word order and permutations of object relations, e.g., the text-image alignment of the sentences “the goldfish is swimming in the aquarium” and “the aquarium is swimming in the goldfish” are scored similarly. Such behavior is caused by the training objective applied to pre-train VLMs. The contrastive pre-training optimizes for image-text retrieval on large datasets, which does not acknowledge compositional information and thus fails to learn unique representations [27]. A step towards a solution of this problem are hard negative samples [27], [54], where existing prompts are transformed to represent negative compositional semantics by word or relation swapping, and are included to the training set for fine-tuning. Many of the content-based T2I metrics rely on the outputs of VLMs or LLMs that may contain additional details fabricated by a language model rather than actually represented by the image. Additionally, VLMs show limited capability of understanding inputs of multiple images, which may result in low correlation scores on image editing tasks [67]. VLMs are good at generation task evaluation, but fail at image to image evaluation due to high level feature focus [67].

Existing T2I datasets [79], [93], [127]–[129] mainly originate from various online sources, where image-text pairs are collected by applying heuristics to filter the data, thereby often trading quality for quantity. Otherwise, high quality image descriptions need to be crowdsourced by human annotators, which is time consuming and costly. With increasing focus towards the evaluation of visio-linguistic compositionality the necessity of compositional datasets intensifies [43], [130]–[132]. While the evaluation on such complex datasets fosters the development of compositional metrics, the active research in this field seems to stick to a limited set of four compositional aspects: object accuracy, spatial relations, non-spatial relations and attribute binding. However, we consider this to be a subset of a greater set which is yet to be explored, thus, in the work of Dehouche [133] they apply GPT-3 [29] to explore a set of 20 topics, e.g., medium, technique, genre, mood, tone, lighting, artistic reference, which are derived from human generated prompts taken from Lexica<sup>2</sup>.

Benchmarking image synthesis is lacking of comparability due to evaluation on individually proposed datasets providing insights on specific topics. Although, there are widely adopted compositional datasets [43], the size of such datasets limits the assumptions that can be made regarding generalizability. However, creating a comprehensive benchmark for compositionality evaluation should be targeted in

2. <https://lexica.art/>

the near future.

## 6 GUIDELINES

In the following, we provide guidelines for evaluating text-to-image synthesis models based on our findings surveying the literature [62], [134]–[137]. These guidelines can help practitioners in the field make more informed decisions about which metrics and benchmark datasets to use when working with text-to-image synthesis.

### 6.1 Select Metrics based on Relevant Characteristics

Benchmarking T2I synthesis involves measuring general image quality and compositional quality (cf. Section 2). However, what defines the image quality might depend on the target application. For instance, in the domain of artistic image synthesis (e.g., comics, anime, mangas, and paintings) an image has to reflect certain art styles, drawing characteristics, shapes and colors. However, images do not need to be photo-realistic and naturalistic. In order to capture and measure such a large variety of abstract concepts, there exist many visual quality metrics, see Table 1. Each of these metrics is equipped with unique reasoning capabilities, such as aesthetic and human preference prediction, perceptual artifact localization, object recognition, object counting, spatial relations, object attribute recognition, and many more. Hence, reasoning skills for evaluation techniques need to be selected carefully and the calibration of their priority is crucial. Considering the use case of generating synthetic images for pre-training object detection networks of real images, one would need to ensure that the image generator produces correct visual representations of described objects. This necessitates metrics with strong object recognition and object counting capabilities. As shown in Section 3, the current state-of-the-art does not include a general purpose metric satisfying a comprehensive evaluation of T2I synthesis. We provide a classification of metrics and their capabilities, which can be used to make informed decisions about which metric to use in a specific application context.

### 6.2 Select appropriate Evaluation Prompts

The underlying text prompts are fundamental for evaluating T2I synthesis, as they form the input to the image generator. In Section 6.1, we recommend selecting metrics based on which characteristics are most relevant in a user’s application context. In their entirety, a combination of relevant metrics measures the performance of the examined generator. However, it is equally important to ensure evaluation prompts include rich descriptions that cover a broad set of visual concepts. Otherwise, there is no way to obtain comprehensive benchmark results. In our supplementary material, we provide an overview for the state-of-the-art datasets containing image-text pairs with different levels of complexity. Textual descriptions that originate from image captioning datasets usually lack the range of visual concepts needed for the evaluation of T2I synthesis. Using prompts that do not cover the visual depictions to be measured can help outperform other methods but render test results meaningless. Therefore, the collection of evaluation prompts

needs to represent authenticity, complexity, compositionality and representativity of textual descriptions with respect to the target application.

### 6.3 Normalize Prompts

The most recent diffusion-based image generation models [138] can be used to synthesize realistic-looking images of impressive quality. The data these models were trained of may be subject to language bias, which results in a biased image generator, e.g., specific sentence formats, such as the absence of grammatical structure, certain keyword constellations, or artist names that are known only by some models. In order to mitigate such bias, it can help to normalize evaluation prompts by adopting the strong rephrasing, summarization, and completion capabilities of modern LLMs. In particular, LLMs can be used to transform prompts to natural language, complete sentences, remove keywords. On top of this, further normalization protocols may be applied. For some applications it may be beneficial to normalize prompt length since some text encoder networks have limited token vector lengths. Hypernymization, where a word is replaced by its hypernym (i.e., another word that describes it in a more general way, e.g., daisy and rose would be replaced by flower), is a method to semantically normalize prompts [139]. However, this may lower the variety of evaluation prompts. Furthermore, the representation of numbers and dates can be brought into a consistent format, e.g., *Two dogs are playing with a ball.* and *2 dogs are playing with a ball.*

### 6.4 Set Model Parameters

As diffusion-based image generation is sensitive to the selected seed for the initial noise sampling during early diffusion steps, it is crucial to fix such seeds to guarantee reproducibility. Further, some benchmarks compare image generators that share an identical training protocol and have only small architectural differences or vice versa. Utilizing identical seeds clarifies the contribution of these changes. The image resolution used during training can have a strong influence on image quality. Thus, it should be configured appropriately and consistently throughout all evaluated methods. This extends to the sampling method, sampling steps, and guidance parameters. When the image generation pipelines are properly configured for each prompt in the evaluation set, a fixed number of  $N$  images is generated, where  $N$  is equal to the number of model parameter configurations. Higher numbers for  $N$  provide increasingly robust performance results in exchange for computational cost.

### 6.5 Measure Performance

After generating images using the pipeline configured as described in Section 6.4, evaluation performance is computed using the metrics as defined based on Section 6.1 for each synthesized image. Hereby, some measures are defined for the image distribution of the entire test dataset (cf. Section 2.1), while most of the metrics are computed for individual images. Recently developed compositional quality metrics are able to provide not only score associated



with image quality, instead, they additionally output a rationale for deeper explainability of the score, cf. Section 2.2.2. In order to extract comprehensive benchmark results, a close inspection of individual compositional aspects is necessary. Thus, both general image quality and compositional quality should be captured and investigated separately by the selected evaluation metrics. Drawing conclusions from results based on a single metric lacks sufficient variation in perspective, leaving research and usability questions only partially answered.

## 7 CONCLUSION

This survey provides an overview of the current SOTA in evaluation metrics for T2I synthesis. First, we introduce our taxonomy to categorize quality measures based on the data they evaluate (images alone vs. text-conditioned images), their scope (distribution of images vs. single images), their operating data structure (embeddings vs. content), and what they measure (general quality vs. compositional quality). Many widely adopted T2I metrics lack the ability to assess the alignment between text and image, missing important details limits a comprehensive evaluation. Second, we collect, review, and compare both established and emerging evaluation metrics, acknowledging the trend towards compositional quality metrics. Such compositional quality metrics are sensitive to the prompt definition and can detect and judge the model's alignment quality between image and text. Third, we review and discuss optimization-based approaches aimed at increasing T2I alignment that fine-tune the generators or optimize cross-attention layers during inference. Finally, based on these findings, we provide guidelines for the development of comparable and meaningful evaluation protocols, enabling consistent quality assessment and, thus, representative T2I synthesis evaluation.

## REFERENCES

- [1] S. Luo, "A survey on multimodal deep learning for image synthesis: Applications, methods, datasets, evaluation metrics, and results comparison," in *Proceedings of the 2021 5th International Conference on Innovation in Artificial Intelligence*, ser. ICIAI '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 108–120. [Online]. Available: <https://doi.org/10.1145/3461353.3461388>
- [2] C. Zhang, C. Zhang, M. Zhang, and I. S. Kweon, "Text-to-image diffusion models in generative ai: A survey," 2023.
- [3] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, "Diffusion models in vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 9, pp. 10 850–10 869, 2023.
- [4] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet, "Video diffusion models," 2022.
- [5] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, and T. Salimans, "Imagen video: High definition video generation with diffusion models," 2022.
- [6] G. Zhang, J. Bi, J. Gu, and V. Tresp, "Spot! revisiting video-language models for event understanding," *arXiv preprint arXiv:2311.12919*, 2023.
- [7] C.-H. Lin, J. Gao, L. Tang, T. Takikawa, X. Zeng, X. Huang, K. Kreis, S. Fidler, M.-Y. Liu, and T.-Y. Lin, "Magic3d: High-resolution text-to-3d content creation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 300–309.
- [8] G. Metzger, E. Richardson, O. Patashnik, R. Giryes, and D. Cohen-Or, "Latent-nerf for shape-guided generation of 3d shapes and textures," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 12 663–12 673.
- [9] R. Chen, Y. Chen, N. Jiao, and K. Jia, "Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 22 246–22 256.
- [10] S. Wang, C. Saharia, C. Montgomery, J. Pont-Tuset, S. Noy, S. Pellegrini, Y. Onoe, S. Laszlo, D. J. Fleet, R. Soricut, J. Baldridge, M. Norouzi, P. Anderson, and W. Chan, "Imagen editor and edit-bench: Advancing and evaluating text-guided image inpainting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 18 359–18 369.
- [11] Z. Wang, C. Lu, Y. Wang, F. Bao, C. Li, H. Su, and J. Zhu, "Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=ppJuFSOAnM>
- [12] X. Wu, K. Xu, and P. Hall, "A survey of image synthesis and editing with generative adversarial networks," *Tsinghua Science and Technology*, vol. 22, no. 6, pp. 660–674, 2017.
- [13] S. Schuhmann, R. Beaumont, R. Vencu, C. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman *et al.*, "Laion-5b: An open large-scale dataset for training next generation image-text models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 25 278–25 294, 2022.
- [14] J. Xu, X. Liu, Y. Wu, Y. Tong, Q. Li, M. Ding, J. Tang, and Y. Dong, "Imagereward: Learning and evaluating human preferences for text-to-image generation," *arXiv preprint arXiv:2304.05977*, 2023.
- [15] Y. Kirstain, A. Polyak, U. Singer, S. Matiana, J. Penna, and O. Levy, "Pick-a-pic: An open dataset of user preferences for text-to-image generation," *ArXiv*, vol. abs/2305.01569, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258437096>
- [16] X. Wu, Y. Hao, K. Sun, Y. Chen, F. Zhu, R. Zhao, and H. Li, "Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis," *ArXiv*, vol. abs/2306.09341, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:259171771>
- [17] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen, "Improved techniques for training gans," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2016/file/8a3363abe792db2d8761d6403605aeb7-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/8a3363abe792db2d8761d6403605aeb7-Paper.pdf)
- [18] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 6629–6640.
- [19] L. Zhang, Y. Zhou, C. Barnes, S. Amirghodsi, Z. Lin, E. Shechtman, and J. Shi, "Perceptual artifacts localization for inpainting," in *European Conference on Computer Vision*. Springer, 2022, pp. 146–164.
- [20] L. Zhang, Z. Xu, C. Barnes, Y. Zhou, Q. Liu, H. Zhang, S. Amirghodsi, Z. Lin, E. Shechtman, and J. Shi, "Perceptual artifacts localization for image synthesis tasks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 7579–7590.
- [21] S. Hartwig, M. Schelling, C. v. Onzenoodt, P.-P. Vázquez, P. Hermosilla, and T. Ropinski, "Learning human viewpoint preferences from sparsely annotated models," *Computer Graphics Forum*, vol. 41, no. 6, pp. 453–466, 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14613>
- [22] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [23] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=YicbFdNTTy>
- [24] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [25] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.



- [26] J. Li, D. Li, S. Savarese, and S. C. H. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," in *International Conference on Machine Learning*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:256390509>
- [27] M. Yuksekgonul, F. Bianchi, P. Kalluri, D. Jurafsky, and J. Zou, "When and why vision-language models behave like bags-of-words, and what to do about it?" in *The Eleventh International Conference on Learning Representations*, 2022.
- [28] T. Gokhale, H. Palangi, B. Nushi, V. Vineet, E. Horvitz, E. Kamar, C. Baral, and Y. Yang, "Benchmarking spatial relationships in text-to-image generation," *ArXiv*, vol. abs/2212.10015, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:254877055>
- [29] R. OpenAI, "Gpt-4 technical report. arxiv 2303.08774," *View in Article*, vol. 2, p. 13, 2023.
- [30] M. Minderer, A. Gritsenko, A. Stone, M. Neumann, D. Weissenborn, A. Dosovitskiy, A. Mahendran, A. Arnab, M. Dehghani, Z. Shen *et al.*, "Simple open-vocabulary object detection with vision transformers. arxiv 2022," *arXiv preprint arXiv:2205.06230*, vol. 2, 2022.
- [31] D. Reis, J. Kupec, J. Hong, and A. Daoudi, "Real-time flying object detection with yolov8," *arXiv preprint arXiv:2305.09972*, 2023.
- [32] J. Wu, J. Wang, Z. Yang, Z. Gan, Z. Liu, J. Yuan, and L. Wang, "Grit: A generative region-to-text transformer for object understanding," *arXiv preprint arXiv:2212.00280*, 2022.
- [33] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [34] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "Spice: Semantic propositional image caption evaluation," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 382–398.
- [35] Y. Cui, G. Yang, A. Veit, X. Huang, and S. Belongie, "Learning to evaluate image captioning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [36] M. Jiang, Q. Huang, L. Zhang, X. Wang, P. Zhang, Z. Gan, J. Diesner, and J. Gao, "Tiger: Text-to-image grounding for image caption evaluation," *arXiv preprint arXiv:1909.02050*, 2019.
- [37] P. Madhyastha, J. Wang, and L. Specia, "VIFIDEL: Evaluating the visual fidelity of image descriptions," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, A. Korhonen, D. Traum, and L. Màrquez, Eds. Florence, Italy: Association for Computational Linguistics, Jul. 2019, pp. 6539–6550. [Online]. Available: <https://aclanthology.org/P19-1654>
- [38] H. Lee, S. Yoon, F. Dernoncourt, D. S. Kim, T. Bui, and K. Jung, "ViLBERTScore: Evaluating image caption using vision-and-language BERT," in *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, S. Eger, Y. Gao, M. Peyrard, W. Zhao, and E. Hovy, Eds. Online: Association for Computational Linguistics, Nov. 2020, pp. 34–39. [Online]. Available: <https://aclanthology.org/2020.eval4nlp-1.4>
- [39] X. Wu, K. Sun, F. Zhu, R. Zhao, and H. Li, "Human preference score: Better aligning text-to-image models with human preference," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 2096–2105.
- [40] T. M. Dinh, R. Nguyen, and B.-S. Hua, "Tise: Bag of metrics for text-to-image synthesis evaluation," in *European Conference on Computer Vision*. Springer, 2022, pp. 594–609.
- [41] P. Grimal, H. Le Borgne, O. Ferret, and J. Tourille, "Tiam - a metric for evaluating alignment in text-to-image generation," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2024, pp. 2890–2899.
- [42] B. Gordon, Y. Bitton, Y. Shafir, R. Garg, X. Chen, D. Lischinski, D. Cohen-Or, and I. Szepkter, "Mismatch quest: Visual and textual feedback for image-text misalignment," 2023.
- [43] T. Thrush, R. Jiang, M. Bartolo, A. Singh, A. Williams, D. Kiela, and C. Ross, "Winoground: Probing vision and language models for visio-linguistic compositionality," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 5238–5248.
- [44] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [45] J. Chen, J. Chen, H. Chao, and M. Yang, "Image blind denoising with generative adversarial network based noise modeling," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [46] L. D. Tran, S. M. Nguyen, and M. Arai, "Gan-based noise model for denoising real images," in *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020.
- [47] C. Tian, X. Zhang, C.-W. Lin, W. Zuo, and Y. Zhang, "Generative adversarial networks for image super-resolution: A survey," *ArXiv*, vol. abs/2204.13620, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:248426817>
- [48] W. Ahmad, H. Ali, Z. Shah, and S. Azmat, "A new generative adversarial network for medical images super resolution," *Scientific Reports*, vol. 12, no. 1, p. 9533, 2022.
- [49] S. Frolov, T. Hinz, F. Raue, J. Hees, and A. Dengel, "Adversarial text-to-image synthesis: A review," *Neural Netw.*, vol. 144, no. C, p. 187–209, dec 2021. [Online]. Available: <https://doi.org/10.1016/j.neunet.2021.07.019>
- [50] J. Hessel, A. Holtzman, M. Forbes, R. L. Bras, and Y. Choi, "Clip-score: A reference-free evaluation metric for image captioning," *arXiv preprint arXiv:2104.08718*, 2021.
- [51] J.-H. Kim, Y. Kim, J. Lee, K. M. Yoo, and S.-W. Lee, "Mutual information divergence: A unified metric for multimodal generative models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 35072–35086, 2022.
- [52] D. H. Park, S. Azadi, X. Liu, T. Darrell, and A. Rohrbach, "Benchmark for compositional text-to-image synthesis," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. [Online]. Available: <https://openreview.net/forum?id=bKBhQhPeKaF>
- [53] H. Singh, P. Zhang, Q. Wang, M. Wang, W. Xiong, J. Du, and Y. Chen, "Coarse-to-fine contrastive learning in image-text-graph space for improved vision-language compositionality," *arXiv preprint arXiv:2305.13812*, 2023.
- [54] S. Castro, A. Ziai, A. Saluja, Z. Yuan, and R. Mihalcea, "Clove: Encoding compositional language in contrastive vision-language models," *arXiv preprint arXiv:2402.15021*, 2024.
- [55] S. Fu, N. Tamir, S. Sundaram, L. Chai, R. Zhang, T. Dekel, and P. Isola, "Dreamsim: Learning new dimensions of human visual similarity using synthetic data," *arXiv preprint arXiv:2306.09344*, 2023.
- [56] Z. Ma, C. Wang, Y. Ouyang, F. Zhao, J. Zhang, S. Huang, and J. Chen, "Cobra effect in reference-free image captioning metrics," *arXiv preprint arXiv:2402.11572*, 2024.
- [57] T. Xu, P. Zhang, Q. Huang, H. Zhang, Z. Gan, X. Huang, and X. He, "Attgan: Fine-grained text to image generation with attentional generative adversarial networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, jun 2018, pp. 1316–1324. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CVPR.2018.00143>
- [58] Y. Liang, J. He, G. Li, P. Li, A. Klimovskiy, N. Carolan, J. Sun, J. Pont-Tuset, S. Young, F. Yang *et al.*, "Rich human feedback for text-to-image generation," *arXiv preprint arXiv:2312.10240*, 2023.
- [59] K. Huang, K. Sun, E. Xie, Z. Li, and X. Liu, "T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation," *arXiv preprint arXiv: 2307.06350*, 2023.
- [60] T. Hinz, S. Heinrich, and S. Wermter, "Semantic object accuracy for generative text-to-image synthesis," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 3, pp. 1552–1565, 2020.
- [61] F. Betti, J. Staiano, L. Baraldi, L. Baraldi, R. Cucchiara, and N. Sebe, "Let's vice! mimicking human cognitive behavior in image generation evaluation," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 9306–9312.
- [62] Y. Hu, B. Liu, J. Kasai, Y. Wang, M. Ostendorf, R. Krishna, and N. A. Smith, "Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 20406–20417.
- [63] M. Yarom, Y. Bitton, S. Changpinyo, R. Aharoni, J. Herzig, O. Lang, E. O. Ofek, and I. Szepkter, "What you see is what you read? improving text-image alignment evaluation," *ArXiv*, vol. abs/2305.10400, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258740893>
- [64] J. Singh and L. Zheng, "Divide, evaluate, and refine: Evaluating

- and improving text-to-image alignment with iterative vqa feedback," *arXiv preprint arXiv:2307.04749*, 2023.
- [65] H. Chefer, Y. Alaluf, Y. Vinker, L. Wolf, and D. Cohen-Or, "Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models," *ACM Transactions on Graphics (TOG)*, vol. 42, no. 4, pp. 1–10, 2023.
- [66] Y. Lu, X. Yang, X. Li, X. E. Wang, and W. Y. Wang, "LLMScore: Unveiling the power of large language models in text-to-image synthesis evaluation," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=OJ0c6um1An>
- [67] M. Ku, D. Jiang, C. Wei, X. Yue, and W. Chen, "Viescore: Towards explainable metrics for conditional image synthesis evaluation," *arXiv preprint arXiv:2312.14867*, 2023.
- [68] C.-Y. Bai, H.-T. Lin, C. Raffel, and W. C.-w. Kan, "On training sample memorization: Lessons from benchmarking generative modeling with a large-scale competition," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 2534–2542.
- [69] M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton, "Demystifying mmd gans," in *International Conference on Learning Representations*, 2018.
- [70] D. Lopez-Paz and M. Oquab, "Revisiting classifier two-sample tests," in *International Conference on Learning Representations*, 2016.
- [71] M. S. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, and S. Gelly, "Assessing generative models via precision and recall," *Advances in neural information processing systems*, vol. 31, 2018.
- [72] S. Ravuri and O. Vinyals, "Classification accuracy score for conditional generative models," *Advances in neural information processing systems*, vol. 32, 2019.
- [73] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, "Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 22 500–22 510.
- [74] T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, and T. Aila, "Improved precision and recall metric for assessing generative models," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [75] S. Gu, J. Bao, D. Chen, and F. Wen, "Giga: Generated image quality assessment," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*. Springer, 2020, pp. 369–385.
- [76] J. Wang, K. C. Chan, and C. C. Loy, "Exploring clip for assessing the look and feel of images," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 2, 2023, pp. 2555–2563.
- [77] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4401–4410.
- [78] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick, "Microsoft coco captions: Data collection and evaluation server," *arXiv preprint arXiv:1504.00325*, 2015.
- [79] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier, "From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions," *Transactions of the Association for Computational Linguistics*, vol. 2, pp. 67–78, 2014. [Online]. Available: <https://aclanthology.org/Q14-1006>
- [80] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik, "Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models," in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 2641–2649.
- [81] H. Tan and M. Bansal, "Lxmert: Learning cross-modality encoder representations from transformers," *arXiv preprint arXiv:1908.07490*, 2019.
- [82] R. Hu and A. Singh, "Unit: Multimodal multitask learning with a unified transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 1439–1449.
- [83] Y.-C. Chen, L. Li, L. Yu, A. El Kholy, F. Ahmed, Z. Gan, Y. Cheng, and J. Liu, "Uniter: Universal image-text representation learning," in *European conference on computer vision*. Springer, 2020, pp. 104–120.
- [84] Z. Gan, Y.-C. Chen, L. Li, C. Zhu, Y. Cheng, and J. Liu, "Large-scale adversarial training for vision-and-language representation learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6616–6628, 2020.
- [85] P. Zhang, X. Li, X. Hu, J. Yang, L. Zhang, L. Wang, Y. Choi, and J. Gao, "Vinvl: Revisiting visual representations in vision-language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 5579–5588.
- [86] X. Li, X. Yin, C. Li, P. Zhang, X. Hu, L. Zhang, L. Wang, H. Hu, L. Dong, F. Wei et al., "Oscar: Object-semantics aligned pre-training for vision-language tasks," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*. Springer, 2020, pp. 121–137.
- [87] W. Kim, B. Son, and I. Kim, "Vilt: Vision-and-language transformer without convolution or region supervision," in *International Conference on Machine Learning*. PMLR, 2021, pp. 5583–5594.
- [88] K. Li, Y. Zhang, K. Li, Y. Li, and Y. Fu, "Visual semantic reasoning for image-text matching," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [89] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9650–9660.
- [90] X. Chen, X. Wang, S. Changpinyo, A. Piervigiani, P. Padlewski, D. Salz, S. Goodman, A. Grycner, B. Mustafa, L. Beyer et al., "Pali: A jointly-scaled multilingual language-image model," *arXiv preprint arXiv:2209.06794*, 2022.
- [91] M. Honnibal and I. Montani, "spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing," *To appear*, vol. 7, no. 1, pp. 411–420, 2017.
- [92] J. Liang, W. Pei, and F. Lu, "Cpgan: Content-parsing generative adversarial networks for text-to-image synthesis," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*. Springer, 2020, pp. 491–508.
- [93] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. Springer, 2014, pp. 740–755.
- [94] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier, "From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions," *Transactions of the Association for Computational Linguistics*, vol. 2, pp. 67–78, 2014.
- [95] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *ArXiv*, vol. abs/2304.08485, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258179774>
- [96] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell et al., "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [97] D. Khashabi, S. Min, T. Khot, A. Sabharwal, O. Tafjord, P. Clark, and H. Hajishirzi, "Unifiedqa: Crossing format boundaries with a single qa system," *arXiv preprint arXiv:2005.00700*, 2020.
- [98] C. Li, H. Xu, J. Tian, W. Wang, M. Yan, B. Bi, J. Ye, H. Chen, G. Xu, Z. Cao, J. Zhang, S. Huang, F. Huang, J. Zhou, and L. Si, "mPLUG: Effective and efficient vision-language learning by cross-modal skip-connections," in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Y. Goldberg, Z. Kozareva, and Y. Zhang, Eds. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics, Dec. 2022, pp. 7241–7259. [Online]. Available: <https://aclanthology.org/2022.emnlp-main.488>
- [99] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh, "Making the v in vqa matter: Elevating the role of image understanding in visual question answering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [100] M. Hodosh, P. Young, and J. Hockenmaier, "Framing image description as a ranking task: Data, models and evaluation metrics," *Journal of Artificial Intelligence Research*, vol. 47, pp. 853–899, 2013.
- [101] J. Wang and R. Gaizauskas, "Generating image descriptions with gold standard visual inputs: Motivation, evaluation and baselines," in *Proceedings of the 15th European Workshop on Natural Language Generation (ENLG)*, A. Belz, A. Gatt, F. Portet, and M. Purver, Eds. Brighton, UK: Association for



- Computational Linguistics, Sep. 2015, pp. 117–126. [Online]. Available: <https://aclanthology.org/W15-4722>
- [102] M. Kusner, Y. Sun, N. Kolkin, and K. Weinberger, “From word embeddings to document distances,” in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 957–966. [Online]. Available: <https://proceedings.mlr.press/v37/kusnerb15.html>
- [103] T. Zhang\*, V. Kishore\*, F. Wu\*, K. Q. Weinberger, and Y. Artzi, “Bertscore: Evaluating text generation with bert,” in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=SkeHuCVFDR>
- [104] J. Lu, D. Batra, D. Parikh, and S. Lee, “Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/c74d97b01eae257e44aa9d5bade97baf-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/c74d97b01eae257e44aa9d5bade97baf-Paper.pdf)
- [105] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ser. ACL ’02. USA: Association for Computational Linguistics, 2002, p. 311–318. [Online]. Available: <https://doi.org/10.3115/1073083.1073135>
- [106] C.-Y. Lin, “ROUGE: A package for automatic evaluation of summaries,” in *Text Summarization Branches Out*. Barcelona, Spain: Association for Computational Linguistics, Jul. 2004, pp. 74–81. [Online]. Available: <https://aclanthology.org/W04-1013>
- [107] S. Banerjee and A. Lavie, “METEOR: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, J. Goldstein, A. Lavie, C.-Y. Lin, and C. Voss, Eds. Ann Arbor, Michigan: Association for Computational Linguistics, Jun. 2005, pp. 65–72. [Online]. Available: <https://aclanthology.org/W05-0909>
- [108] A. Lavie, K. Sagae, and S. Jayaraman, “The significance of recall in automatic metrics for mt evaluation,” in *Machine Translation: From Real Users to Research: 6th Conference of the Association for Machine Translation in the Americas, AMTA 2004, Washington, DC, USA, September 28-October 2, 2004. Proceedings 6*. Springer, 2004, pp. 134–143.
- [109] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [110] S. Barratt and R. Sharma, “A note on the inception score,” *arXiv preprint arXiv:1801.01973*, 2018.
- [111] T. Che, Y. Li, A. Jacob, Y. Bengio, and W. Li, “Mode regularized generative adversarial networks,” in *International Conference on Learning Representations*, 2016.
- [112] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola, “A kernel method for the two-sample-problem,” *Advances in neural information processing systems*, vol. 19, 2006.
- [113] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [114] Z. Li, M. R. Min, K. Li, and C. Xu, “Stylet2i: Toward compositional and high-fidelity text-to-image synthesis,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 18 197–18 207.
- [115] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [116] H. Dong, W. Xiong, D. Goyal, Y. Zhang, W. Chow, R. Pan, S. Diao, J. Zhang, K. SHUM, and T. Zhang, “RAFT: Reward ranked finetuning for generative foundation model alignment,” *Transactions on Machine Learning Research*, 2023. [Online]. Available: <https://openreview.net/forum?id=m7p5O7zblY>
- [117] K. Lee, H. Liu, M. Ryu, O. Watkins, Y. Du, C. Boutilier, P. Abbeel, M. Ghavamzadeh, and S. S. Gu, “Aligning text-to-image models using human feedback,” 2023.
- [118] M. Prabhudesai, A. Goyal, D. Pathak, and K. Fragkiadaki, “Aligning text-to-image diffusion models with reward backpropagation,” 2023.
- [119] Y. Fan, O. Watkins, Y. Du, H. Liu, M. Ryu, C. Boutilier, P. Abbeel, M. Ghavamzadeh, K. Lee, and K. Lee, “Reinforcement learning for fine-tuning text-to-image diffusion models,” in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=8OTPePXzeh>
- [120] K. Black, M. Janner, Y. Du, I. Kostrikov, and S. Levine, “Training diffusion models with reinforcement learning,” in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=YCWjhGrJFD>
- [121] N. Liu, S. Li, Y. Du, A. Torralba, and J. B. Tenenbaum, “Compositional visual generation with composable diffusion models,” in *European Conference on Computer Vision*. Springer, 2022, pp. 423–439.
- [122] W. Feng, X. He, T.-J. Fu, V. Jampani, A. R. Akula, P. Narayana, S. Basu, X. E. Wang, and W. Y. Wang, “Training-free structured diffusion guidance for compositional text-to-image synthesis,” in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=PUlqjT4rzq7>
- [123] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 10 684–10 695.
- [124] Y. Li, M. Keuper, D. Zhang, and A. Khoreva, “Divide & bind your attention for improved generative semantic nursing,” *arXiv preprint arXiv:2307.10864*, 2023.
- [125] M. Menéndez, J. Pardo, L. Pardo, and M. Pardo, “The jensen-shannon divergence,” *Journal of the Franklin Institute*, vol. 334, no. 2, pp. 307–318, 1997. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0016003296000634>
- [126] M. Chen, I. Laina, and A. Vedaldi, “Training-free layout control with cross-attention guidance,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2024, pp. 5343–5353.
- [127] J. Mao, J. Xu, Y. Jing, and A. Yuille, “Training and evaluating multimodal word embeddings with large-scale web annotated images,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, ser. NIPS’16. Red Hook, NY, USA: Curran Associates Inc., 2016, p. 442–450.
- [128] P. Sharma, N. Ding, S. Goodman, and R. Soricut, “Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 2556–2565.
- [129] S. Changpinyo, P. Sharma, N. Ding, and R. Soricut, “Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 3558–3568.
- [130] Z. Ma, J. Hong, M. O. Gul, M. Gandhi, I. Gao, and R. Krishna, “Crepe: Can vision-language foundation models reason compositionally?” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 10 910–10 921.
- [131] C.-Y. Hsieh, J. Zhang, Z. Ma, A. Kembhavi, and R. Krishna, “Sugarcrepe: Fixing hackable benchmarks for vision-language compositionality,” in *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. [Online]. Available: <https://openreview.net/forum?id=jsc7WSCZd4>
- [132] A. Ray, F. Radenovic, A. Dubey, B. Plummer, R. Krishna, and K. Saenko, “cola: A benchmark for compositional text-to-image retrieval,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [133] N. Dehouche and K. Dehouche, “What’s in a text-to-image prompt? the potential of stable diffusion in visual arts education,” *Heliyon*, vol. 9, no. 6, 2023.
- [134] M. Ku, T. Li, K. Zhang, Y. Lu, X. Fu, W. Zhuang, and W. Chen, “Imagenhub: Standardizing the evaluation of conditional image generation models,” in *The Twelfth International Conference on Learning Representations*, 2024. [Online]. Available: <https://openreview.net/forum?id=OuV9ZrkQlc>
- [135] M. Otani, R. Togashi, Y. Sawai, R. Ishigami, Y. Nakashima, E. Rahtu, J. Heikkilä, and S. Satoh, “Toward verifiable and reproducible human evaluation for text-to-image generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 14 277–14 286.

- [136] T. Lee, M. Yasunaga, C. Meng, Y. Mai, J. S. Park, A. Gupta, Y. Zhang, D. Narayanan, H. B. Teufel, M. Bellagente, M. Kang, T. Park, J. Leskovec, J.-Y. Zhu, L. Fei-Fei, J. Wu, S. Ermon, and P. Liang, "Holistic evaluation of text-to-image models," in *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. [Online]. Available: <https://openreview.net/forum?id=qY9LR74O3Z>
- [137] E. Salin, S. Ayache, and B. Favre, "Towards an exhaustive evaluation of vision-language foundation models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, October 2023, pp. 339–352.
- [138] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, W. Zhang, B. Cui, and M.-H. Yang, "Diffusion models: A comprehensive survey of methods and applications," *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–39, 2023.
- [139] A. Baryshnikov and M. Ryabinin, "Hypernymy understanding evaluation of text-to-image models via wordnet hierarchy," *arXiv preprint arXiv:2310.09247*, 2023.
- [140] V. Ordonez, G. Kulkarni, and T. Berg, "Im2text: Describing images using 1 million captioned photographs," in *Advances in Neural Information Processing Systems*, J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, Eds., vol. 24. Curran Associates, Inc., 2011. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2011/file/5dd9db5e033da9c6fb5ba83c7a7e99-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2011/file/5dd9db5e033da9c6fb5ba83c7a7e99-Paper.pdf)
- [141] H. Agrawal, K. Desai, Y. Wang, X. Chen, R. Jain, M. Johnson, D. Batra, D. Parikh, S. Lee, and P. Anderson, "nocaps: novel object captioning at scale," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [142] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier, "Collecting image annotations using amazon's mechanical turk," in *Proceedings of the NAACL HLT 2010 workshop on creating speech and language data with Amazon's Mechanical Turk*, 2010, pp. 139–147.
- [143] C. L. Zitnick and D. Parikh, "Bringing semantics into focus using visual abstraction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2013.
- [144] S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh, "Vqa: Visual question answering," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [145] R. Zellers, Y. Bisk, A. Farhadi, and Y. Choi, "From recognition to cognition: Visual commonsense reasoning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [146] A. Rohrbach, A. Torabi, M. Rohrbach, N. Tandon, C. Pal, H. Larochelle, A. Courville, and B. Schiele, "Movie description," 2016.
- [147] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans *et al.*, "Photorealistic text-to-image diffusion models with deep language understanding," *Advances in Neural Information Processing Systems*, vol. 35, pp. 36 479–36 494, 2022.
- [148] G. Marcus, E. Davis, and S. Aaronson, "A very preliminary analysis of dall-e 2," 2022.
- [149] J. Cho, A. Zala, and M. Bansal, "Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 3043–3054.
- [150] P. Schramowski, M. Brack, B. Deiseroth, and K. Kersting, "Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 22 522–22 531.
- [151] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei, "Visual genome: Connecting language and vision using crowdsourced dense image annotations," *International Journal of Computer Vision*, vol. 123, pp. 32 – 73, 2016. [Online]. Available: <https://api.semanticscholar.org/CorpusID:4492210>
- [152] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," *arXiv preprint arXiv:1809.11096*, 2018.
- [153] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [154] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2008 (VOC2008) Results," <http://www.pascal-network.org/challenges/VOC/voc2008/workshop/index.html>.
- [155] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov *et al.*, "The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale," *International Journal of Computer Vision*, vol. 128, no. 7, pp. 1956–1981, 2020.
- [156] J. Kiros, W. Chan, and G. Hinton, "Illustrative language understanding: Large-scale visual grounding with image search," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018, pp. 922–933.
- [157] P. Jenkins, A. Farag, S. Wang, and Z. Li, "Unsupervised representation learning of spatial data via multimodal embedding," in *Proceedings of the 28th ACM international conference on information and knowledge management*, 2019, pp. 1993–2002.
- [158] K. Desai and J. Johnson, "Virtext: Learning visual representations from textual annotations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 11 162–11 173.
- [159] C. Chambers, A. Raniwala, F. Perry, S. Adams, R. R. Henry, R. Bradshaw, and N. Weizenbaum, "Flumejava: easy, efficient data-parallel pipelines," in *Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 363–375. [Online]. Available: <https://doi.org/10.1145/1806596.1806638>
- [160] W. Su, X. Zhu, Y. Cao, B. Li, L. Lu, F. Wei, and J. Dai, "Vlbert: Pre-training of generic visual-linguistic representations," in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=SygXPaEYvH>
- [161] J. Cho, J. Lei, H. Tan, and M. Bansal, "Unifying vision-and-language tasks via text generation," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 1931–1942. [Online]. Available: <https://proceedings.mlr.press/v139/cho21a.html>
- [162] G. A. Miller, "Wordnet: a lexical database for english," *Commun. ACM*, vol. 38, no. 11, p. 39–41, nov 1995. [Online]. Available: <https://doi.org/10.1145/219717.219748>
- [163] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *International conference on machine learning*. Pmlr, 2021, pp. 8821–8831.
- [164] N. Xie, F. Lai, D. Doran, and A. Kadav, "Visual entailment task for visually-grounded language learning," 2019.



**Sebastian Hartwig** is a Ph.D. student at Ulm University, Germany, where he previously received his B.Sc. and M.Sc. degrees in computer science. He completed his master's degree at the Institute of Media Informatics in 2017 before joining the research group Visual Computing. His research focus on human perception based machine learning for computer graphics and visualization application, as well as machine scene understanding like depth and layout estimation.





**Dominik Engel** is a Ph.D. student at Ulm University, Germany, where he previously received his B.Sc. and M.Sc. degrees in computer science. In 2018, he joined the Visual Computing research group. His research focuses on deep learning in visualization and computer graphics, differentiable and neural rendering.



**Michael Glöckler** is currently pursuing a Master's degree in Media Informatics at Ulm University, having previously earned his Bachelor's degree in the same field from Ulm University in 2022. Collaborating in projects with the Visual Computing Group, exploring areas such as natural language processing, deep learning, and text-to-image synthesis.



**Leon Sick** is a Ph.D. student at Ulm University and part of the Visual Computing Group. Before starting his Ph.D., he obtained his B.A. in International Business Administration from Aalen University of Applied Sciences and his M.Sc. in Business Information Technology from Konstanz University of Applied Sciences. His research is focused on self-supervised pre-training and unsupervised segmentation on 2D images.



**Hannah Kniesel** is a Ph.D. student at Ulm University, Germany. She finished her M.Sc. in 2022 with a focus on reconstruction of bio-medical data. Prior to that she completed her Bachelor degree in 2019, also at the University of Ulm. She joined the Visual Computing Research Group in April 2020.



**Alex Bäuerle** received Ph.D. in computer science from Ulm University in 2022 and is now working at a researcher at Carnegie Mellon University. His current research interests are at the intersection of human-computer interaction and artificial intelligence.



**Tristan Payer** has joined the Visual Computing Group in November 2020 as research associate. He received his master's degree in 2020 from Radboud University Nijmegen in the field of Artificial Intelligence. His focus in the master's program was on Natural Language Processing, Medical Image Analysis and Deep Learning. Previously, he completed his bachelor's degree at Radboud University Nijmegen in 2018.



**Timo Ropinski** is a professor at Ulm University, heading the Visual Computing Group. Before moving to Ulm, he was Professor in Interactive Visualization at Linköping University, heading the Scientific Visualization Group. He received his Ph.D. in computer science in 2004 from the University of Münster, where he also completed his Habilitation in 2009. Currently, Timo serves as chair of the EG VCBM Steering Committee, and as an editorial board member of IEEE TVCG.



**Poonam Poonam** has completed her Master in Cognitive Systems at Ulm University in January 2023 before joining the research group Visual Computing. In her Master Thesis, she focused on contrastive and continual learning.

Dataset	Size	Words	Source	Text Compositionality				Task
				QA	SR	NR	AB	
SBU Captioned Photo Dataset [140]	1,000,000		Flickr.com	●	●	●	●	Image Caption
Pinerest40M [127]	40,000,000	10	Pinerest.com	●	●	●	●	VLM Pre-Training
Conceptual Captions [128]	3,369,218	10.3	World Wide Web	●	●	●	●	Image Caption
nocaps [141]	15,100	10	Open Images V4 (Flickr.com)	●	●	●	●	Image Caption
Conceptual 12M [129]	12,423,374	20.2	World Wide Web	●	●	●	●	VLM Pre-Training
UIUC Pascal Sentence Dataset [142]	1,000	n.a.	VOC2008	●	●	●	●	Image Caption
Flickr8K [100]	8,092	n.a.	Flickr.com	●	●	●	●	Image Caption
Flickr30K [79]	31,783	n.a.	Flickr.com	●	●	●	●	Image Caption
COCO Captions [78]	204,721	11	Flickr.com	●	●	●	●	Image Caption
PASCAL-50S [33]	1,000	8.8	VOC2008	●	●	●	●	Image Caption
ABSTRACT-50S [33]	500	10.59	ASD [143]	●	●	●	●	Image Caption
VQA [144]	254,721	< 2	MS COCO & Abstract Images	●	●	●	●	VQA
VQAy2.0 [99]	204,721	n.a.	MS COCO	●	●	●	●	VQA
VCR [145]	110,000	11.8	LSMDC [146] & YT	●	●	●	●	VQA
DrawBench [147]	200	11.69	DALL-E, [148],Reddit	●	●	●	●	T2I
PaintSkills [149]	65,535	n.a.	synthetic prompts	●	●	●	●	T2I
ABC-6K [122]	6,400	n.a.	MS COCO	●	●	●	●	T2I
CC-500 [122]	500	n.a.	synthetic prompts	●	●	●	●	T2I
I2P [150]	4,703	20.56	user generated prompts	●	●	●	●	T2I
Visual Genome [151]	108,077	n.a.	MS COCO	●	●	●	●	T2I
Winoground [43]	800	8.99	Getty Images API	●	●	●	●	T2I
RichHF-18K [58]	18,000	n.a.	Pick-a-Pic [15]	●	●	●	●	T2I
T2I-CompBench [59]	6,000	8.98	generated prompts by GPT [29]	●	●	●	●	T2I

TABLE 2

Comparison of text-image datasets based on the number of prompts, prompt length, compositional aspects of the prompts, and the context of the dataset provided in their textual data.

## APPENDIX A DATASETS

This section provides an overview of datasets used to evaluate text-conditioned image synthesis, see Table 2. First, we discuss datasets that origin from the image captioning research community, which were used first to evaluate text-image generation systems. With recent developments of vision-language models, researchers started to seek for increasing complexity of evaluation data resulting in the emergence of the term, visio-linguistic compositionality. It describes the task and datasets for evaluating the ability of vision and language models to conduct reasoning of image and text which are subject to compositionality, meaning they are ensembles of several contents. Second, we discuss existing visio-linguistic compositionality benchmark datasets, and finally we inspect specifically designed datasets for development of text-image metrics and their verification on human judgments.

The ranking system in Table 2 evaluates prompts based on their source and complexity, assigning points from zero to three. Zero points are given for basic object labels, providing minimal information. One point is awarded to prompts derived from web-scraping, offering a bit more context. Two points go to prompts obtained through crowdsourcing, reflecting a higher level of detail and relevance. The highest score, three points, is reserved for prompts that accurately reflect the actual compositional intentions behind an image, showcasing the deepest understanding and context.

### A.1 Image Caption Datasets

The development of image generators requires tremendous amounts of image data [152], [153] in order to learn data statistics and fitting the output distribution of the generator to real image distributions. In the context of evaluation text-conditioned image synthesis, the necessity of text-image pairs arises. Fortunately, there already exist such datasets

collected from researchers of the image captioning research domain, e.g., MS-COCO Captions [78], Flickr30K [79], [80], PASCAL-50S [33], Abstrac-50S [33], which curate one to fifty human generated descriptions per image.

UIUC Pascal Sentence Dataset [142] and Flickr8K [100] are the first popular image caption datasets, featuring multiple image descriptions per image. The work of Hodosh et al. [100], that frames the evaluation of image description as ranking task and describes the collection of 8,092 images from Flickr and 1,000 from PASCAL VOC-2008 [154] having human annotators describe each image. The crowdsource instructions and quality control mechanism are adopted from Rashtchian et al. [142] in order to collect five image descriptions per image from human annotators utilizing the Amazon Turk Platform. Users were asked to describe images they see, using a single sentence while focusing on the main characters, settings and object relations. Further, they were asked to use adjectives for color, spacing, emotion or quantity, utilizing no more than 100 characters. A separate set of users were asked to perform quality control by conducting spelling and grammar checks in order to ensure high quality textual descriptions of images. Later, Young et al. [79] extend the size of the Flickr8K dataset incorporating 158,915 image captions, which correspond to 31,783 referring to it as Flickr30K. Another extension was proposed by Plummer et al. [80]. They collect cross-caption coreference chains, that link same entities mentioned in different captions of the same image, and meanwhile bounding boxes localize such entities in image space.

The SBU Captioned Photo Dataset [140] collects one million images from Flickr.com ensuring some quality requirements, in particular they filter the collected data for textual descriptions with satisfactory length of visual description, at least two words belonging to (objects, attributes, actions, stuff, and scenes) and at least one preposition indicating visible spatial relation. While the dataset poses a tremendous amount of image-text pairs, the content of image captions

may be visual descriptive but lacks human supervised verification resulting in many image captions being only comprehensible with personal knowledge of the caption’s author, e.g., using the given name of a dog for describing a dog playing with a ball.

Based on Microsoft COCO [93], which is a large-scale dataset consisting of images acquired through Flickr showing multiple objects in their natural context, the frequently used COCO Captions [78] dataset was created. It supplements MS-COCO by collecting 1,026,459 captions for 164,062 images, including five captions for each image in MS-COCO and a subset of 5,000 images were annotated with 40 reference sentences. Together with the actual dataset the authors released an evaluation protocol, in particular they deploy an evaluation server ensuring consistent evaluation computing numerous metrics like BLEU [105], ROGUE [106], METEOR [107] and CIDEr [33].

In the work of Vedantam et al. [33] two datasets are collected, PASCAL-50S and ABSTRACT-50S based on the UIUC Pascal Sentence Dataset and the Abstract Scenes Dataset [143], respectively. Annotations for these datasets were collected with the goal to investigate consensus between humans annotators, in particular the similarity between a candidate image description and several reference descriptions. While PASCAL-50S feature real images, ABSTRACT-50S consists of images in a clip-art style designed by humans in a different crowdsourcing study [143]. For both datasets, 50 human generated sentences are collected while annotators are instructed to provide descriptions that should help others recognize the image from a collection of similar images. Having a large set of fifty reference sentences per image facilitates research of text-image alignment, however the amount and variety of images provided by both datasets seems too few in order to provide complexity for profound text-image evaluation.

Increasing the number of object classes is achieved by the image dataset called, nocaps. It consists of over 600 object classes, and it is presented in the work of Agrawal et al. [141], which is based on OpenImages V4 [155] a large-scale human-annotated object detection dataset. Nocaps was acquired by filtering Open Images and excluding images with non-zero or unknown image rotations, instances from a single object category, less than six unique object classes and finally they apply a balancing scheme to have an even distribution of images depicting two to six unique object classes avoiding frequently occurring object classes.

Mao et al.’s [127] introduction of the Pinterest40M dataset represents a significant advancement in multimodal word embeddings, featuring over 40 million images and 300 million sentences from Pinterest.com. Far exceeding the scale of existing datasets like MS COCO, Pinterest40M’s unique blend of visual and textual data enables the development of richer word embeddings. Further, this dataset serves as a vital resource for exploring vision-language pre-training methods [156]–[158].

The Conceptual Captions [128] dataset is programmatically acquired through web-crawling billions of webpages utilizing the JavaScript framework Flume [159] developed to build efficient data-parallel pipelines. In such a way, a tremendous amount of image-text pairs can be downloaded, however stringent filtering is necessary in order to keep

only high quality images and captions. Therefore, they excluded images based on encoding format, size, aspect ratio, and offensive content. Since the accompanying Alt-text harvested from the HTML webpages are not restricted to be good image descriptions, filtering was applied by analyzing part-of-speech, sentiment/polarity, pornography/profanity annotations using the Google Cloud Natural Language APIs. To further increase textual quality, several heuristics based on the count of nouns, prepositions, token repetition, word capitalization, English Wikipedia token likelihood and predefined boilerplate prefix, e.g., “click to enlarge picture”, “stock photo”. Ultimately, image-text based filtering through Google Cloud Vision APIs is applied, which tries to map text tokens to content of the image, as well as text transformation with hypernymization replaces proper names by its hypernym. The intention of this dataset, consisting of over 3 million image-text pairs, is to serve different downstream tasks of image captioning, however it tends to be frequently utilized for vision-language pre-training [129].

To facilitate image-text pre-training, Conceptual 12M [129] was acquired by relaxing filter criteria of the collection pipeline used for Conceptual Captions. This strategy trades precision of image descriptions for increased scale of the data corpus, in particular they increase the recall of visual concept descriptions by lowering requirements of word repetitions, caption size ranges, image aspect ratios and hypernymization. Just as Conceptual Captions and Pinterest40M, such web sourced image descriptions enable vision-language pre-training but lack the quality and complexity for proper evaluation of T2I synthesis methods.

## A.2 Visual Question Answering

The Visual Question Answering (VQA) dataset [144], a pioneering resource in the field, combines 123,287 MS COCO images and 50,000 abstract scenes with over 760K million questions and 10 million answers crowdsourced via Amazon Mechanical Turk. Designed to challenge VQA models in understanding complex visual content, the dataset covers a broad spectrum of question types and answers, reflecting real-world diversity and the intricacies of human language. Each image is associated with five open-ended questions, requiring detailed visual recognition, commonsense knowledge, and inferential reasoning, with ten distinct answers per question to capture the variability of human responses. A successor work, called VQA v2.0 [99], builds on top of VQA by collecting complementary images for each question, resulting in a pair of images with two different answers per question. Goyal et al. counteract the problem of VLMs ignoring visual information by doubling the VQA dataset size to twice the number of image-question pairs. This expansion enables them to develop a model that can not only provide an answer to an image-question pair but also offer a counterexample-based explanation.

The Visual Commonsense Reasoning (VCR) dataset [145] is specifically designed to move beyond mere recognition tasks to test models on cognition-level visual understanding. It features 290,000 question-answer-rationale (QAR) triples across 110,000 unique movie scenes. Each QAR triple challenges models to not only identify objects within a scene but also to understand complex interactions and motiva-



tions. The dataset focuses on deep visual comprehension, requiring models to infer and rationalize about unseen aspects of the image, thus bridging the gap between visual perception and commonsense reasoning. VCR is frequently used as downstream task for evaluating representation learning of visual-linguistic approaches [38], [83], [160], [161].

### A.3 Compositionality Benchmarks

The work of Thrush et al. [43] introduces Winoground, a novel task and dataset designed to evaluate vision and language models’ capabilities in visio-linguistic compositional reasoning. This involves matching two images with two captions, each containing the same set of words but in a different order, requiring precise comprehension of both modalities. Winoground is collected from the Getty Image API instructing human annotators to generate captions and find a corresponding image while being as creative as possible and marking visual reasoning tags, which can be categorized into three groups: objects swaps, relation swaps or swaps of both. Winoground consists of 1,600 image-text pairs, with 800 correct and 800 incorrect pairings across 400 examples, showcasing 800 unique captions and images. This dataset emphasizes quality expert annotations over size, serving as a probing dataset for linguistic and visual analysis.

[151] The Visual Genome dataset [51], is a dataset for comprehensive scene understanding. It contains more than 108k images, each with an average of 35 objects delineated by a bounding box. However, bounding box annotated objects are not sufficient for a compressive scene understanding. Object attributes and their relationships are also needed. To obtain these, about 50 overlapping image sub-regions per image have been captured by human annotators. From those, object attributes and relationships could be extracted, which turn are used to create image scene graphs and additional question answer pairs. Object attributes and relationships are canonicalized on Word-Net synsets [162].

T2I-CompBench [59] is a compositional dataset targeting to provide complex prompt compositions in order to study attribute binding, object relations and complex compositions skills of image generation models. Therefore, they acquire a dataset consisting of 6,000 text-image pairs (1,000 for each sub-category: color, shape, texture, spatial relation, non-spatial relation, complex composition). Text prompt for color attribute binding are gathered from CC500 [122] and COCO [78], while for the remaining sub-classes prompts are generated by GPT [29] or handcrafted using prompt templates.

Rich human feedback, in short RichHF-18K [58], is a dataset consisting of 18K image-text pairs collected from Pick a Pic dataset [15]. For each image human annotations are crowdsourced consisting of two heatmaps localizing artifacts/implausibility and misalignment, four scores, i.e. plausibility, alignment, aesthetic, overall quality and one text sequence indicating misaligned keywords. In order to filter for photorealistic images and balance classes across images they adopt PaLI visual question answering model, which assesses realism and selects images of five classes: *animal*, *human*, *object*, *indoor scene* and *outdoor scene*. Heatmaps are derived from averaged key point positions across annotators for artifact and misalignment annotations. The dataset

provides complex annotations enabling fine-tuning a scoring model on human feedback, however the collected annotations origin from only 27 individual annotators which took around 3,000 rater-hours, which questions the quality of annotations as well as consensus of such a small group of raters.

DrawBench is a dataset proposed by Saharia et al. [147], developed alongside the Imagen model. It comprises a challenging set of 200 prompts designed to evaluate T2I generators across 11 categories, aimed at investigating various abilities such as colors, numbers of objects, spatial relations, text in the scene, unusual interactions between objects, misspellings, rare words, long prompts, and prompts from Reddit, Gary Marcus et al. [148], and DALL-E [163]. Saharia et al. [147] utilize DrawBench to compare different T2I models; thus, they present generated images to human raters for quantifying image quality and text-image alignment quality.

PaintSkills proposed by Cho et al. [149] constitutes a dataset collected specifically to mitigate a statistical bias towards a few common objects. Therefore, Cho et al. generate a dataset carefully controlling three aspects (skills): object recognition, object counting and spatial relations resulting in 65,535 scene configurations. By uniformly sampling from a set of relations, PaintSkills ensures equal distributed objects and relations. Finally, based on the scene configurations a 3D simulator is used to render images.

Feng et al. [122] propose two datasets, Attribute Binding Contrast (ABC-6K) and Concept Conjunction 500 (CC-500). The former dataset is derived from MSCOCO, where Feng et al. filter for sentences containing at least two color words, and by switching the position of two color words they generate additional contrastive sentences resulting in total of 6.4K sentences. CC-500 is generated by combining two objects with their attribute description, where each sentence follows the same pattern, e.g. “a red apple and a yellow banana” resulting in 500 sentences.

The Inappropriate Image Prompts (I2P) dataset [150] targets towards safe latent diffusion by mitigating the problem of models generating inappropriate images. Therefore, Schramowski et al. collect 4,703 prompts from an online source, that distributes real-world human-generated prompts together with SD [123] generated images and corresponding generation parameters. Prompts are filtered based on 26 keywords that correspond to one of seven inappropriateness concepts, e.g. hate, harassment, violence, self-harm, sexual content, shocking images and illegal activity.

Holistic Evaluation of Text-to-Image Models (HEIM) [136] is a benchmark dataset that evaluates T2I models based on 12 aspects, e.g. alignment, quality, aesthetics, originality, reasoning, knowledge, bias, toxicity, fairness, robustness, multilingualism and efficiency. It combines several existing text-image datasets like MSCOCO, DrawBench, PartiPrompts, Winoground, PaintSkills, I2P, etc. to cover evaluation of each aspect.

The goal of the SeeTRUE benchmark is to study text-image alignment evaluation. The dataset builds on top of several existing vision-language datasets: COCO Captions [78], SNLI-VE [164], DrawBench [147], EditBench [10], Winoground [43] and Pick a Pic [15]. It includes 31, 855 real and synthetic image-text pairs and corresponding human

annotations, where each binary annotation indicates alignment or misalignment of text and image.

**APPENDIX B  
HUMAN PREFERENCE METRICS**

In the following experiments, we investigate human preferences regarding text prompt generation and corresponding diffusion-based image synthesis. Consequently, our first goal is to inspect human provided text prompts, investigating preferred styles, concepts, topics, sceneries, etc., which we retrieve in a data-driven way from a large corpus of human provided data.

**B.1 Dataset Acquisition**

In order to collect a large-scale dataset consisting of human-generated prompts along with synthesized images, we downloaded messages from the official Midjourney Discord server. Each message contains information about a user-generated text prompt, the corresponding synthetic image, the author’s name, as well as links to preceding and succeeding user interactions. In this way, we were able to collect 8, 290, 132 text-image pairs.

**B.2 Data-driven Prompt Categorization**

To gain initial insights into the dataset containing various prompts, we examined a subset and discovered that users often use the same prompt multiple times with only minor modifications to generate an image. This observation led us to infer that users were dissatisfied with the initial image generated, prompting them to modify the prompt in hopes of achieving a more pleasing outcome. To understand what users added to the prompts to enhance image generation, we conducted a detailed analysis of these modified prompts. Establishing relationships between these prompts requires a preliminary comparison of their similarity. To accomplish this, we examined all prompts in the dataset and calculated CLIP embeddings for each individual prompt. With these CLIP embeddings, we were able to compare prompts using cosine similarity between the embeddings and connect prompts with high similarity in a graph structure. This enables us to understand user interactions with multiple prompts and their corresponding images produced by Midjourney. Furthermore, we assume that the most recent prompt, in terms of time, was considered the best variant by the user, as it seemingly met their satisfaction, leading them to cease modifications of the prompt. Subsequently, we examined all final prompts from the prompt graphs and created a dictionary to store the occurrences of each individual word.

Then, we proceeded to analyze the most frequently occurring words (cf. Figure 3) and defined categories for the prompts. However, we exclude common words such as “a”, “and” and “the” as well as parameter keywords specific to the Midjourney Engine, since they do not contribute to differentiating between categories. The top five most frequent words included general terms such as “style” (6096 occurrences), “background” (5179 occurrences), and “white” (5137 occurrences), as well as category-specific terms like “realistic” (6820 occurrences) and “logo” (5789



Fig. 3. This wordcloud visualize the occurrences of different words in the final prompts

occurrences). Through an examination of these frequently occurring words, we were able to define 11 categories, see Table 3. For each category, we then selected a keyword that epitomizes the category. These keywords were subsequently used to ascertain the group to which a prompt belongs.

Category	Top Words
realistic	“realistic” (6820), “real” (674)
logo	“logo” (5789)
photo	“lighting” (5056), “photography” (1940), “photo” (1778), “photorealistic” (1248), “reflections” (1137)
Art	“art” (2312)
cartoon	“cartoon” (2064)
anime	“anime” (1699), “manga” (185)
cyberpunk	“cyberpunk” (1207), “futuristic” (1061), “cyber” (182)
portrait	“portrait” (1135), “eyes” (2178)
simple	“simple” (1065), “minimalist” (795)
illustration	“illustration” (995), “painting” (970)
landscape	“landscape” (593), “mountains” (389), “sunset” (673)

TABLE 3  
Categories and Top Words occurred in different prompts

**B.3 Inferring Human Preferences**

In Section 3, several synthetic image metrics are reviewed, that are based on fine-tuning to increase alignment with human preferences (e.g. ImageReward, Human Preference Score v2 and Aesthetic Predictor). In this experiment, we aim to infer image quality scores, enabling us to investigate the differences between such metrics. As a baseline, we include metrics learned via representation learning, e.g., CLIPScore, BLIP, and BLIP2. In this experiment, we applied the six metrics to all text-image pairs collected in Section B.1 enabling us to rank pairs for each metric separately. However, since these scores use different scales for their values, for instance, CLIPScore returns values around 2.5, while the aesthetic score ranges between 0 and 10, comparing the scores is not straightforward. Therefore, we normalized all scores so that all values lie between 0 and 1. In Figure 4, we show box plots of the distribution of each image quality score. It is notable that BLIP mostly contains very high scores, with a median above 0.9 and a very low standard deviation, whereas the other scores demonstrate a more balanced distribution.




Fig. 4. Box plot visualization of the value ranges for each of the normalized image quality scores.

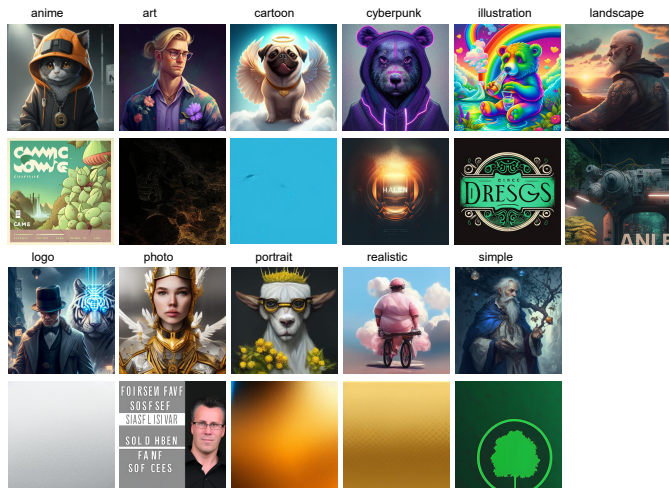


Fig. 5. For each prompt category, we display the corresponding image with the highest (top row) and lowest (bottom row) sum of six image quality scores

After taking a closer look at the value ranges for different scores, this section investigates the visual appearance of images ranked by the scores, respectively. Therefore, we attempted to identify the best and worst scored image for each category and quality metric. First, we investigate an overall quality by combining the six metrics, Figure 5 displays the images for each category based on the sum of the six image quality scores. As observed across all categories, the best image exhibits superior characteristics, featuring clear and identifiable features, whether they be animals, humans, or other depicted elements. In contrast, the worst image within each category tends to be harder to recognize, often resulting in unclear visual content or even areas of a single color. This comparison provides valuable insights into the varying quality levels across different categories, highlighting the effectiveness of the selected quality metrics in assessing and distinguishing image quality.

To gain more detailed insights, we also examined the highest and lowest-rated images for each individual image quality score. When reviewing the images, as shown in Figure 6, a noticeable difference between the worst and best

images is evident, although the contrast is not as strong as it is when considering the sum of all scores, as we did in Figure 5. In some images, the depicted scenes can still be easily identified, which may not always be possible with the lowest-rated images when considering the sum of all scores. Another observation is that CLIP, BLIP and BLIP2 tend to value a match between the prompt and the image more than the actual aesthetics of the image. The lowest-rated score for an image of these metrics does not correlate with poor visual quality. It seems to receive a lower rating because it doesn't match well with the given prompt. This observation leads to the conclusion that representation learning based image metrics do not reflect human preferences well.

#### B.4 Influence of Prompt Length

In this experiment, we delve into the correlation between the length of prompts and their evaluation performance across various categories. By analyzing the prompts associated with both the lowest-rated and highest-rated images in each category, as detailed in Table 4, a clear trend emerges: prompts for the lowest-rated images tend to be significantly longer than those for the highest-rated images. This observation suggests that prompt length may inversely affect image evaluation outcomes, potentially due to factors such as clarity, focus, and user engagement. Further analysis reveals patterns in the compositional elements and thematic content of the highest-rated prompts, indicating that optimal prompt length might also depend on the specific context and requirements of each category. Furthermore, we observe additional keywords that remained present after our initial filtering out of command parameters specific to Midjourney. We assume that these characteristics of the prompts introduce a negative bias, resulting in poor text-image alignment assessment by the adopted metrics. This issue should be investigated in future research.

To investigate the potential correlation between prompt length and the combined image quality score, we conducted a detailed examination of the prompts. It was observed that the maximum prompt length is 1800 characters, likely a limitation imposed by MidJourney at the time. Figure 7 illustrates the impact of prompt length on the combined image quality score. It was found that prompts shorter than 200 characters tend to achieve the highest scores, and as prompt length increases, the scores generally decline. However, the correlation between prompt length and image quality scores is not particularly strong, with a coefficient of approximately -0.07. Despite this, the median score remains consistently close to four across all lengths of prompts.



	Best Image	Best Image Prompt	Worst Image	Worst Image Prompt
HPS		A nendroid style priest with a long grey beard and a cross around his neck in black round priest's hat in high white sneakers and black sunglasses disassembling boxes of shoes in the warehouse. 3D, Logo, Nendroid style		patterndesigns:: blocky::1
Image Reward		The Terrible Vampire Monster in the Peach Blossom Forest		captain of the ship guided by the stars in the ocean logo
CLIP Score		crypto trading hacker goon behind desk in analley trashcan cats silhouette of man with knife in his mouth money on the floor beggar with raised hands poor homeless gold wide shot stock trading bitcoin ethereum hopium gains diamond hands help cry panic thunder loss crash stock crash bear		collection of six variations of a simplistic funkko anime-style cartoon character, resembling Digimon Agumon to full height. Each rendition exhibits distinct variations in attributes, hues and temperament exhibiting distinct variations in attributes, hues and temperament, gradient shading, uncomplicated sketches, well-defined outlines, and octane rendering. Arranged in a 3 and 3 formation. Ultra-high resolution quality at 8K UHD. The parameters used for color, saturation
Aesthetic Score		humanization of the city of Saratov Russia as Russian bakery girl, bright portrait, russian village atmosphere, high realistic, 35mm, by Norman Rockwell		area chats showing a crypto fund's great performance over the last year
BLIP Score		Logo with inscription \"HIGH QUEENS\" minimalist pink street urban		Cybernetic samurai with katana blades looking down on Tokyo from atop a skyscraper on a cloudy night. vaporwave, neon lighting, ultra-detailed, realistic, science fiction, Adobe photoshop, high-tech, metallic, award-winning, futuristic sci-fi, imaginative futuristic, abstract::-2, style:: 3, high-quality
BLIP2 Score		connor mcgregor eating a purple ice-cream in the rain		Cybernetic samurai with katana blades looking down on Tokyo from atop a skyscraper on a cloudy night. vaporwave, neon lighting, ultra-detailed, realistic, science fiction, Adobe photoshop, high-tech, metallic, award-winning, futuristic sci-fi, imaginative futuristic, abstract::-2, style:: 3, high-quality

Fig. 6. For each image quality score, we display the corresponding image and image prompt with the highest (left) and lowest scores (right).

Category	Best prompt	Worst prompt
realistic	fat boy, on bicycle, cloud, pink dress, big arms, <b>realistic</b> style	Large Luxury Single Family neighborhood with mansions + aerial drone photo + Very Detailed + Very <b>Realistic</b> + Ultra High Definition website design, UI/UX, unreal engine, detailed, ultra high definition, 2k
logo	a man with top hat and smoking suit with a white tiger at his left side and an ethereum <b>logo</b> fluctuating over his right hand in a futurist cyberpunk city	a vector <b>logo</b> for a software company, infrastructures, roads, highway, vector graphics, logotype
photo	Goddess in gold and white ivory armor ornate with golden wings on the helmet levitating above the ground, full-body portrait + <b>photo</b> -realistic renders. ultra-detailed , 8k , dramatic epic lightning, realistic texture	Please design a cover image for my upcoming video on "Simple and Free Cash Flow Spreadsheet for Businesses". The image should feature a computer or mobile screen with the spreadsheet open, displaying the main financial data of a company. There should be a button for free download of the spreadsheet in Excel or Google Sheets in the upper right corner of the screen. At the bottom of the screen, there should be a list of the main features of the spreadsheet, such as "Daily control of inflows and outflows", "Cash movement tracking", "Expense control", and "Compatibility with mobile and desktop". The title of the video should be highlighted at the top of the screen, with the following call-to-action phrase below: "Take control of your finances with the simple and free cash flow spreadsheet for businesses". The background image can be a <b>photo</b> of a desk or a financial chart to emphasize the business aspect of the spreadsheet. Please use green and blue as the main colors to convey trust and serenity. Thank you!
Art	character <b>art</b> , blonde man, long blonde hair in a bun, average build, glasses, purple floral button up shirt, highly detailed, Pixar <b>art</b> style,	giant octopus 3D sculpture, in the style of cloisonné plique-à-jour, enamel, translucent amber and teal color <b>art</b> glass, soft dim light glows from inside octopus sculpture, intricate :: center :: full body :: wide shot :: 45 degree angle :: dark background :: solid black background
cartoon	create a realistic 3d <b>cartoon</b> image of a happy pug with wings and a halo of angel in paradise	a <b>cartoon</b> character about money management, cute, blue
anime	A cat black and white, with orange parka, and a gray cap, with a chain with the bitcoin logo, <b>anime</b> cover	Tags, The explosion of cucumber and jasmine in the iced coffee turned into Hami melon, Dreamy mountains in the style of kawase hasui, Highlight the coffee jasmine cucumber, Peter Mohrbacher, James Jean, Simon Stalenhag, and CloverWorks' style, Japanese urban pop style, Japanese 1980 vintage <b>anime</b> noise, super detail, written on Japanese paper, dynamic angles, high feeling illustrations, masterpiece, 16K, Ultra HD, best quality, perfect surreal composition, decals, explode coffee bean
cyberpunk	purple bear with neon <b>cyberpunk</b> lines in the face roaring, dressed with a hoodie	electric train, light snow, traffic lights, train station, <b>cyberpunk</b> city, Long shot, hyper realistic, 4K, 8k, HD, cinematic, cinematic composition:: Nikon D750:::15 Halogen:::250
portrait	realistic <b>portrait</b> of a white goat with scientific glasses and a crown of yellow flowers on his head	stereo equipment robot, unreal engine 5, Real photography, movement, realism, detailed texture, Cinematic, Color Grading, <b>portrait</b> Photography, Shot on 50mm lens, Ultra-Wide Angle, Depth of Field, hyper-detailed, beautifully color-coded, insane details, intricate details, beautifully color graded, Cinematic, Color Grading, Editorial Photography, Photography, Photoshoot, Shot on 70mm lens, Depth of Field, DOF, Tilt Blur, Shutter Speed 1/1000, F/22, White Balance, 32k, Super-Resolution, Megapixel, ProPhoto RGB, , Good, Massive, Halfrear Lighting, Backlight, Natural Lighting, Incandescent, Optical Fiber, Moody Lighting, Cinematic Lighting, Studio Lighting, Soft Lighting, Volumetric, Contre-Jour, Beautiful Lighting, Accent Lighting, Global Illumination, Screen Space Global Illumination, Ray Tracing Global Illumination, Optics, Scattering, Glowing, Shadows, Rough, Shimmering, Ray Tracing Reflections, Lumen Reflections, Screen Space Reflections, Diffraction Grading, Chromatic Aberration, GB Displacement, Scan Lines, Ray Traced, Ray Tracing Ambient Occlusion, Anti-Aliasing, FXAA, TXAA, RTX, SSAO, Shaders, OpenGL-Shaders, GLSL-Shaders, Post Processing, Post-Production, Cel Shading, Tone Mapping, insanely detailed and intricate, hypermaximalist, elegant, realistic, super detailed, dynamic pose, photograph
simple	wizard in <b>simple</b> blue and white clothes with gray hair and a beard picking an apple from a tree	Simplicity: A hospital logo should be <b>simple</b> and easily recognizable. This is especially important since hospitals deal with patients who may be in distress and not be able to focus on complex designs. Color scheme: Choose a color scheme that is calming, soothing, and associated with health and wellness, such as blue, green, or white. Avoid using bright, bold colors that may be overwhelming. Fonts: Use clear and easily readable fonts, preferably sans-serif fonts, that are legible even from a distance. Images: Incorporate images that are relevant to the hospital's specialty or mission, such as a caduceus for medical institutions, or a heart for cardiac centers. Originality: Avoid copying other hospital logos, and strive to create a unique design that stands out and represents the hospital's values. With these considerations in mind, here are a few hospital logo design ideas: A stylized caduceus with a <b>simple</b> font in blue or green. An abstract symbol that represents the hospital's specialty, such as a stylized heart for a cardiac center. A logo that incorporates a hospital building with a calming color scheme. A logo that uses a <b>simple</b> , clean, and modern font with a small, stylized icon or symbol in the corner. A logo that incorporates a symbol of hope and healing, such as a dove or a lotus flower, in a calming color palette
illustration	Lisa Frank <b>illustration</b> of a bear smoking marijuana through a bong while sitting next to a rainbow river	The logo could feature the words Bastos Chronicles in bold, playful font. The word Bastos could be in a different color or font to emphasize its ironic contrast with the page's wholesome content. The letter Ñn Chronicles could be replaced with a simple <b>illustration</b> of a book or a scroll to represent storytelling. The color scheme could be bright and cheerful, such as blue and orange, to convey a sense of fun and joy.
landscape	An old man on top of his motorcycle, tattoos, big beard and bald, in his look a long road with a beautiful <b>landscape</b> in the background, the sun setting and the sea in the distance	ableton live :: berlin, night time, year 2030, sci-fi, cupertino, flying cars, night sky filled with stars, high res, 4K definition, Night, Cold Colors, Color Grading, Shot on 35mm wide angle lens, Ultra-Wide Angle, Depth of Field, hyper-detailed, beautifully color-coded, insane details, intricate details, beautifully color graded, Unreal Engine 5, Cinematic, Color Grading, Editorial Photography, Photography, Photoshoot, Landscape Shot, Depth of Field, DOF, Tilt Blur, Shutter Speed 1/1000, F/22, White Balance, 32k, Super-Resolution, Megapixel, ProPhoto RGB, VR, Lonely, Good, Massive, Halfrear Lighting, Backlight, Natural Lighting, Incandescent, Optical Fiber, Moody Lighting, Cinematic Lighting, Studio Lighting, Soft Lighting, Volumetric, Contre-Jour, Beautiful Lighting, Accent Lighting, Global Illumination, Screen Space Global Illumination, Ray Tracing Global Illumination, Optics, Scattering, Glowing, Shadows, Rough, Shimmering, Ray Tracing Reflections, Lumen Reflections, Screen Space Reflections, Diffraction Grading, Chromatic Aberration, GB Displacement, Scan Lines, Ray Traced, Ray Tracing Ambient Occlusion, Anti-Aliasing, FXAA, TXAA, RTX, SSAO, Shaders, OpenGL-Shaders, GLSL-Shaders, Post Processing, Post-Production, Cel Shading, Tone Mapping, CGI, VFX, SFX, insanely detailed and intricate, super detailed

TABLE 4

This table shows the prompts of the worst and best image for each category




Fig. 7. This Figure shows the correlation between prompt length and the combined image quality score. For better readability we split the data based on the prompt length into nine chunks and created a boxplot for each chunk

## B.5 Metric Correlation Analysis

To investigate the correlation among the six metrics, we applied a dimension reduction technique to the computed score vectors. This approach allows us to explore how scores vary across different categories and whether the scores of prompts within a category are similar. To this end, we reduced the six dimensions to two using t-SNE (t-distributed Stochastic Neighbor Embedding). This reduction aims to visualize high-dimensional data in a two-dimensional space, facilitating the creation of a cluster plot. By reducing the dimensionality, we hope to uncover hidden patterns and structures not immediately apparent in its original, more complex form. t-SNE excels at maintaining similarities between nearby points, making it ideal for exploring data and identifying groups of similar points. In Figure 8, we present the distribution of data points across the four largest categories: realistic, logo, photo, and art. The plots reveal that some categories, like 'realistic,' 'photo,' and 'art,' are more dispersed, forming several clusters across the space, while others, such as 'logo,' are more concentrated in specific areas. Additionally, we observe clusters representing combinations of the categories realistic, photo, and art, indicating that the prompts contain keywords from each category.

To take a closer look at these observations, we created a cluster plot shown in Figure 9 containing all four categories. The combination of categories is represented by the mixture of the categories' specific colors. The color yellow represents the combination of the category photo (with the color green) and realistic (with the color red). The results indicate that there are indeed clusters representing a category or even sub-categories, such as the orange cluster representing the combination of realistic, photo, and art, or the blue cluster representing the category logo. This suggests that the image quality metrics generally yield similar results for images within the same category.



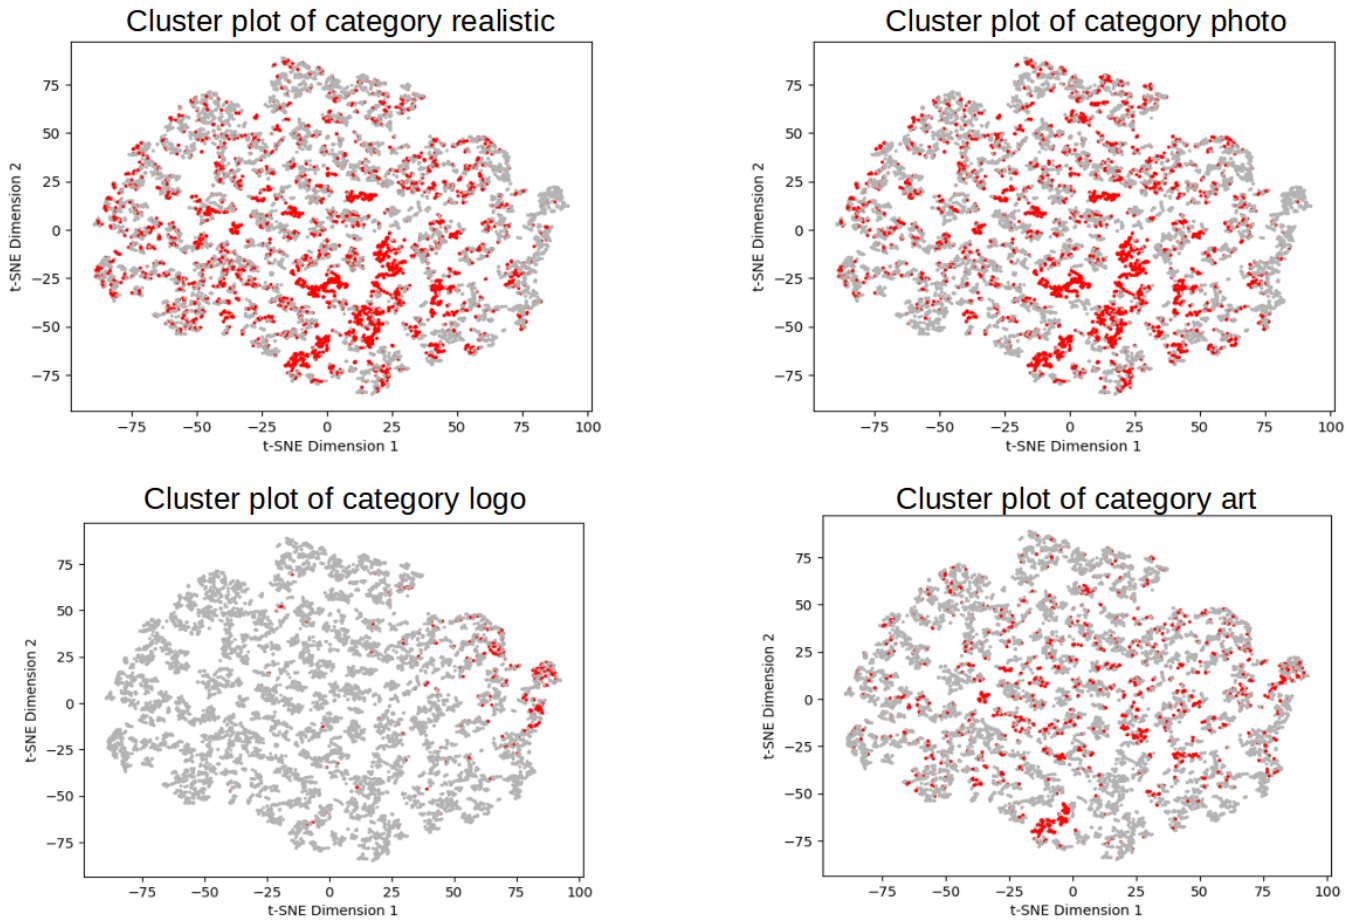


Fig. 8. These cluster plots illustrate the distribution of projected scores across four different categories. Points belonging to each category are highlighted in red, while those not belonging to any category are displayed in gray. Due to computational limitations, we considered only 10,000 prompts for this analysis.




Fig. 9. This cluster plot illustrates the distribution of projected scores across four different categories. Each category is represented by a specific color: 'realistic' in red, 'logo' in blue, 'photo' in green, and 'art' in purple. The colors for combinations of these categories result from the blending of their respective colors.