

TEMPVIZ: On the Evaluation of Temporal Knowledge in Text-to-Image Models

Carolyn Holtermann¹, Nina Krebs², Anne Lauscher¹

¹Trustworthy AI Lab, University of Hamburg

²University of St. Gallen

carolin.holtermann@uni-hamburg.de

Abstract

Time alters the visual appearance of entities in our world, like objects, places, and animals. Thus, for accurately generating contextually-relevant images, knowledge and reasoning about time can be crucial (e.g., for generating a landscape in spring vs. in winter). Yet, although substantial work exists on understanding and improving temporal knowledge in natural language processing, research on how temporal phenomena appear and are handled in text-to-image (T2I) models remains scarce. We address this gap with TEMPVIZ, the first data set to holistically evaluate temporal knowledge in image generation, consisting of 7.9k prompts and more than 600 reference images. Using TEMPVIZ, we study the capabilities of five T2I models across five temporal knowledge categories. Human evaluation shows that temporal competence is generally weak, with no model exceeding 75% accuracy across categories. Towards larger-scale studies, we also examine automated evaluation methods, comparing several established approaches against human judgments. However, none of these approaches provides a reliable assessment of temporal cues – further indicating the pressing need for future research on temporal knowledge in T2I. We publish all data and code at <https://github.com/TAI-HAMBURG/TempViz>.

1 Introduction

“Time is the longest distance between two places.”

– Tennessee Williams, 1944

An artwork from 1900 looks different than an artwork from 1960, and a dog at age 3 months looks different than a dog at age 10 years. Similarly, landscapes change their appearance across seasons, such as *winter* or *summer* – Time influences the visual appearance of the world around us.

The effects of time and their perception are intrinsically intertwined with our senses (e.g., Ma

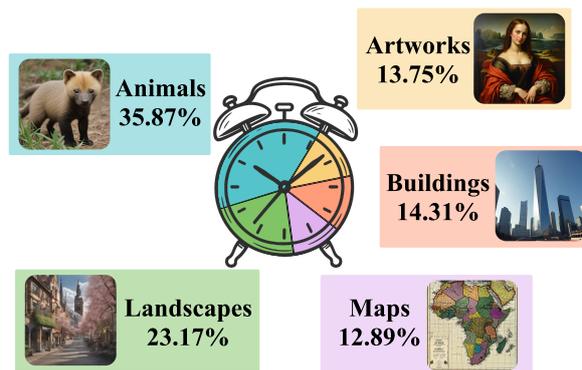


Figure 1: Distribution of the temporal categories within TEMPVIZ with example images of different T2I models.

et al., 2024), and play a major role in most human societies. For accurately representing the world, an understanding of the temporal effects (e.g., aging) and corresponding visual effects is thus crucial. Users might specify in an input prompt to image generation (e.g., to a Stable Diffusion model (Rom-bach et al., 2022a)) the particular epoch or year they are looking for, expecting the generated output to be aligned with their instruction.

In natural language processing (NLP), many studies have analyzed temporal knowledge capabilities of NLP systems, for instance, relating to the temporal reasoning capabilities of large language models (e.g., Gurnee and Tegmark, 2024; Fatemi et al., 2024, *inter alia*). However, despite the large body of existing works on temporal knowledge in text-only models, and the central role of temporal knowledge and reasoning for accurate visual representations, the landscape on resources and evaluation of temporal knowledge in text-to-image (T2I) models remains largely unexplored.

In this work, we address this research gap by posing the following research question: ***How well do existing T2I models perform at following temporal instructions across different subject categories?*** To answer this question, we introduce a novel resource for analyzing temporal knowledge

in T2I models and conduct both manual and automatic evaluations. We find that even the strongest T2I model among those evaluated correctly applies the correct temporal visual cues in at most 75% of cases for only one category, and that for some categories the temporal instruction is barely followed (for example, as low as 15% for historical maps). Strikingly, none of the automated evaluators we tested was able to reliably judge the correct application of temporal knowledge above an macro F1 of 72%. Together, these results motivate more targeted research on temporal knowledge in image generation and evaluation.

Contributions Our contributions are threefold: (1) We present TEMPVIZ, the first dataset specifically designed for analyzing temporal knowledge in T2I models. It comprises 7,940 prompts across five diverse categories (e.g., *animals*, *artworks*), each paired with category-appropriate temporal cues (e.g., *lifespans*, *art periods*, *seasons*), and includes reference images that exemplify correct generations for a subset of categories. (2) Using TEMPVIZ, we conduct a human evaluation of five recent text-to-image models, measuring image quality, subject fidelity, and adherence to the specified temporal cue. (3) We complement the manual study with an extensive assessment of automatic evaluation methods, analyzing their suitability for judging whether temporal knowledge has been applied correctly in generated images.

2 Related Work

We present an overview of datasets for evaluating temporal knowledge in NLP, as well as datasets designed to assess T2I models.

Temporal Knowledge Datasets In large language models, temporal knowledge has been examined from multiple perspectives. Early work on temporal relation extraction introduced annotated corpora such as TimeBank (Pustejovsky et al., 2003). Subsequent benchmarks probe time-varying factual knowledge and temporal reasoning, including the TempLAMA benchmark (Dhingra et al., 2022) and its derived datasets such as LAMA-TK (Zhao et al., 2022) and TempReason (Tan et al., 2023). Other work in the direction of LLM probing focuses rather on time-sensitive question answering (Chen et al., 2021; Holtermann et al., 2025).

Another line of work in the direction of temporal NLI focuses on inferring knowledge about

temporal relations and more complex temporal reasoning. This includes the ordering of temporal events (Thukral et al., 2021; Wang and Zhao, 2024) and temporal grounding (Liu et al., 2023; Fatemi et al., 2024). In addition, several studies investigate methods for improving temporal reasoning through specialized training objectives (Xiong et al., 2024; Tan et al., 2024).

Text-to-Image Evaluation Datasets Compared to the amount of benchmarks available to evaluate the capabilities of LLMs, the landscape of resources to evaluate T2I models still remains scarce. Most benchmarks focus on the evaluation of the general text-image alignment (e.g., Huang et al., 2023; Yarom et al., 2023). Park et al. (2021) specifically focus on the question of how well these models generalize to novel word compositions. Lee et al. (2023) introduce HEIM, a benchmark for the holistic evaluation of T2I models, which includes 12 evaluation aspects such as text-image alignment, image quality, aesthetics, originality, reasoning, knowledge, bias, toxicity, fairness, robustness, multilinguality, and efficiency; still neglecting temporally dependent depictions. Concurrent to our work Niu et al. (2025) present WISE, a benchmark for assessing world knowledge reasoning in T2I models with 1,000 prompts spanning cultural commonsense, spatio-temporal reasoning, and natural science. While they also cover aspects of temporal knowledge, their focus is rather broad and the scale of their dataset is modest: Of the dataset, only 167 prompts concern temporal knowledge, targeting mostly narrow aspects like period-specific instruments instead of the wider visual manifestations of temporal change. As a result, WISE provides limited insights for analyzing and improving these fine-grained temporal effects in image generation.

3 TempViz

We analyze how established T2I models interpret and represent temporal cues during image generation. To this end, we introduce TEMPVIZ, a carefully curated dataset of 7,940 instructions spanning five subject categories (see Figure 1), chosen to reflect common use cases where temporal effects are visually observable. Moreover, the prompts cover multiple dimensions of temporal variation at varying levels of difficulty, enabling a focused assessment of models' temporal reasoning. We intentionally exclude prompts that generate human individu-

als due to ethical concerns: varying both temporal cues and subjects within this category would necessarily involve prompting different demographic backgrounds, introducing complex issues of stereotype reinforcement and bias that would shift the scope from temporal reasoning to fairness considerations. We therefore chose to omit humans and focus on categories where temporal variation can be analyzed with fewer confounds. To support reliable human annotation, we also collect semantically similar real-world reference images for each prompt and provide concise textual descriptors of the visual attributes expected in the generations, when available. In §5, we demonstrate how TEMPVIZ can be used to evaluate temporal image-text alignment.

3.1 Creation of Prompts

We query all models using four paraphrases of the template prompt “[...] a photorealistic image of $\{s_i\}$ in $\{t_i\}$ ”. Each prompt instantiates a subject s_i from category $c \in S$ and a category-specific temporal cue t_i where $S = \{\textit{landscapes}, \textit{animals}, \textit{artworks}, \textit{buildings}, \textit{maps}\}$.

Landscapes A natural test of temporal knowledge in T2I models is the rendering of seasonal and daily variation in landscapes. Therefore, we select commonly known, generic natural environments that evolve over time (e.g., *a lake*). Next, we define a set of temporal cues t_n , including all four seasons (for example *in spring*, *in July*, or *07/17/1899*), and times of the day (e.g., *morning*). To reduce prompt ambiguity in seasonality and daylight, we append the suffix *in Europe* to every prompt and restrict generations to the Northern hemisphere. Correct image generation for this category requires models to integrate multiple knowledge types, specifically commonsense knowledge of seasonal and diurnal cycles, temporal grounding that maps temporal information to seasons and times of day, and geographical knowledge of European climate patterns.

Animals Temporal knowledge is also required to depict biological aging. Here, we focus on mammals, as their aging process is more visually recognizable to human observers than that of insects or amphibians. To assemble a representative set of mammals, we prompt GPT-4 (OpenAI et al., 2024a) to provide “*a set of well-known mammals that show visible signs of aging*”, then consolidate frequent suggestions, remove duplicates, and exclude extinct species. As temporal cues, we defined

four distinct life stages: young (1 month old), adult (50% of life span), very old (80% of life span), and well beyond the expected life span (200% of life span). To this end, we determine the typical lifespans of the selected animals from HAGR¹, an age database of animal aging and longevity. This category requires models to combine factual knowledge about how aging manifests visually in different species with temporal grounding that maps age expressions to species-specific lifespans.

Buildings Another form of temporal change appears in the built environment. Urban form evolves through events such as war, redevelopment, and natural disasters. As city-scale changes are hard to judge from a single image, individual buildings offer a more tractable unit for human annotation. We therefore focus on buildings and structures that have been destroyed over time (e.g., *the World Trade Center before 2001*). Using Wikidata², we identify subjects via SPARQL queries that have known inception and demolition dates, a country label, and a Wikipedia page in at least five languages. As temporal cues, we use four dates: two close to the demolition (just before and just after) and two far from it (well before and well after). To prevent ambiguity from globally recurring building names, we include country information in each prompt. Success in this category requires models to combine event-specific temporal knowledge with causal temporal reasoning (before vs. after the demolition) and factual knowledge about the building’s correct appearance at that time.

Maps Complementing previous categories, we evaluate temporal knowledge in a cartographic context. Therefore, we first use GPT-4 to identify 12 regions with documented border changes over time. Then, to select salient time points, we query the Google Autosuggest API³ for frequently searched completions for each region, and then manually validate all resulting timestamps. We list the number of suggestions per region in the Appendix A.2. This category tests factual temporal knowledge of time-stamped geopolitics, geographic reasoning about borders and land shapes, and symbolic rendering conventions for cartographic depiction.

¹<https://genomics.senescence.info/>

²<https://www.wikidata.org>

³<https://github.com/wildlyinaccurate/google-suggestqueries-api>

Category	Selection Methodology	# Prompts	Prompt Example
Landscapes	92 dates \times 12 locations	460	Create a photorealistic image of a landscape in january
Animals	4 lifespans \times 178 mammals	712	... an Alpaca that is one month old
Buildings	4 dates \times 71 buildings	284	... the Berlin Wall in Germany on 10/23/1983
Maps	64 dates \times 12 regions	256	... a map of europe in 1938
Artworks	10 dates \times 13 periods \times 2 + 13 periods by name	273	... an artwork in a style of [Ancient Art I that was popular in the Western world in 1732]
Total	x4 prompt variations	7,940	

Table 1: Composition of the TEMPVIZ dataset across the five subject categories. For *Artworks*, we include two additional category-specific templates.

Artworks Recent T2I models have demonstrated advanced performance in generating images in well-known artistic styles (Podell et al., 2023), suggesting access to factual knowledge about historical art movements. We thus include an Artworks category to evaluate temporal and period knowledge. Our premise is that, given models’ ability to imitate specific artistic styles, remaining errors primarily reflect weak temporal grounding. Again, we first prompt GPT-4 to propose popular styles (e.g., *Impressionism*), then manually assign the period during which each style was prominent (e.g., *1863–1883*). In this setting, the subject and the temporal cue jointly define the style. For each style, we sample multiple dates within its period and, unlike the other categories, employ two templates: one that names the period explicitly and one that specifies a concrete date within it. To reduce ambiguity and simplify annotation, we restrict prompts to Western art styles and choose dates that lie clearly within, rather than near the boundaries of well-known art epochs. This category requires models to ground timestamps in the correct stylistic period and apply stylistic pattern matching, requiring both temporal reasoning and factual knowledge of art history.

We show an overview of the number of prompts per category in Table 1.

3.2 Collection of Reference Images

For all categories except Landscapes, we retrieve at least one reference image to support the annotation task by providing suitable visual examples. For Buildings and Animals, we use SPARQL queries to crawl suitable images for the subjects from Wikidata. For Artworks, we obtain at least ten examples per period from Google Arts & Culture⁴. For Maps, we source examples from OpenHistoricalMap⁵, a freely usable platform maintained by a community

⁴<https://artsandculture.google.com/>

⁵openhistoricalmaps.com

of mappers and historians. Licensing and copyright considerations are discussed in the Appendix A.3.

3.3 Collection of Expected Values

To provide additional guidance for annotators, we compile short textual descriptions of what a correct image might depict. We refer to these as “expected values”. To obtain them, we prompt GPT 4 for adjectives, single words, and brief phrases that describe the appearance implied by each subject–temporal cue pair. For instance, expected values for an image of a 1-month-old mammal include *tiny*, *infant*, and *juvenile*. We provide a list of all expected values in the Appendix A.4.

4 Evaluating Temporal Knowledge in Multimodal Models

We conduct a human evaluation on a subset of the TempViz dataset, yielding carefully curated ground-truth annotations that serve two purposes: (i) to rigorously assess the extent to which text-to-image models successfully encode and express temporal information during image synthesis, and (ii) to systematically evaluate the validity and reliability of automated evaluation measures in their capacity to judge the generated images accurately.

4.1 Human Evaluation

Task Description For each instance, we present a human annotator with the input prompt and a model generation to assess. Additionally, we provide a reference image, and/or a list of expected values as discussed in §3.2. The annotator then needs to answer each of the following three questions (possible answers: *yes/ no*):

Q1 Is the image free from any visual generation errors?

Q2 Is the subject present in the image?

Q3 Did the model apply the required temporal knowledge correctly?

This way, we clearly separate concerns pertaining to the *image quality (Q1)* from concerns that relate to the models’ instruction following behavior. Moreover, here, we clearly distill general issues with depicting a *subject (Q2)* from issues with following the *temporal cues (Q3)*.

Annotator Selection The annotation task was performed by two annotators from our institution who were (a) familiar with the concept of AI-based image generation, (b) proficient in the English language, and (c) familiarized with the task both via guidelines and direct training. Both annotators labeled the created images by the different models independently from one another.

Despite employing only two annotators, we are confident in the robustness of our annotations for several reasons. First, many categories (e.g., *Animals, Landscapes*) exhibit low subjectivity and are less prone to annotator bias. Second, for more subjective categories (*Artworks, Buildings*), resource constraints prevented us from hiring domain experts, but we mitigated subjectivity through a tightly controlled annotation setup: we used on-site annotators rather than crowdsourcing, allowing direct communication and real-time clarification. We also provided reference images for subjective categories to align expectations and reduce noise.

4.2 Automatic Evaluation

Given the cost and impracticality of large-scale human evaluation, we investigate several automated methods for reliably detecting temporal cues in images. Specifically, we employ two established image–text evaluation metrics (e.g., CLIPScore and Captioning), and query a suite of VLM-as-a-judge approaches using various strategies. In the following, we present all the approaches we compare.

Baseline We compare all our results with a simple baseline computed by simply predicting the binary final label at random.

CLIPScore We adopt the established and widely used CLIPScore proposed by Hessel et al. (2021). Therefore, we use a CLIP-based model to embed both the text prompt used to generate the image and the image itself, and then compute the cosine similarity between the corresponding embeddings to quantify prompt–image alignment.

Captioning For the caption-based evaluation, we invert this procedure. We first employ a BLIP-based model to generate a descriptive caption for each generated image. We then use a sentence-embedding model to encode both the original prompt and the generated caption, and compute the cosine similarity between these embeddings.

For both CLIPScore and caption-based evaluation, we normalize the prompts prior to embedding by stripping imperative prefixes (e.g., “Generate”) so that only the semantic content describing the intended depiction remains. Because both methods yield continuous similarity scores while the annotation labels are binary, we report the Point Biserial Correlation (PBS) between the similarity values and the ground-truth labels to enable a principled comparison.

Decompositional VQA Following Hu et al. (2023), we use an auxiliary LLM to decompose the prompt into fine-grained questions, enabling VQA models to assess nuanced prompt–image alignment. Given a text-to-image prompt, the LLM extracts salient elements and generates multiple-choice and open-ended questions with corresponding correct answers for each. Depending on the prompt complexity, the LLM generates a different number of question-answer pairs, resulting in at least seven questions per prompt. Subsequently, a VLM model is used to answer each question, given the T2I-generated image. Finally, if a certain threshold of questions was answered correctly, the image can be labelled as representing the prompt correctly. We provide an example in Appendix D.

Direct VLM Judging Finally, we evaluate several variants of VLM-as-a-judge prompting: a VLM is asked whether the image correctly depicts the prompt and returns a judgment. We test multiple VLMs and compare strategies aimed at probing a VLM’s capacity to assess temporal cues, including few-shot prompting and an elaborate judging prompt in the style of Lee et al. (2024) that provides task guidance and requests structured criteria and a ranking. Further details and full prompt templates are provided in Appendix D.

5 Experiments and Results

5.1 Experimental Setup

Data With our experiments, we address two questions: (1) whether T2I models are capable of incorporating temporal knowledge within their im-

age generation process, and (2) how well current automatic evaluation approaches capture these fine-grained visual cues. To accurately answer both of these questions, we collect human annotations on a stratified sample of TEMPVIZ comprising 500 prompts with corresponding generations, balanced across categories and models.

Image Generation We evaluate five T2I models on temporal knowledge, selected for public availability, architectural diversity, usage in prior research, and generation quality. Specifically, we compare *stable-diffusion-v1-5* (SDv1.5) (Rombach et al., 2022b), relying on the original Stable Diffusion architecture, with *sdxl-turbo* (SDXL-T), *stable-diffusion-xl-base-1.0* (SDXL-B), and *stable-diffusion-3.5-large* (SDv3.5) relying on the SDXL architecture with a three times larger UNet and a combination of multiple text encoders compared to the original architecture (Podell et al., 2023). This is to analyze performance differences that might occur due to scaling effects. SDXL-T constitutes the distilled version of the original SDXL model using Adversarial Diffusion Distillation. We chose this model to investigate potential drops in performance due to the distillation process. Finally, we add *FLUX.1-dev* (FLUX) (Labs, 2023) to represent a different publisher, and thus likely a different training corpus. Detailed information about the models and their architectures can be found in Appendix B.

Automatic Evaluation To compute embeddings for CLIPScore, we use *CLIP-ViT-bigG-14-laion2B-39B-b160k*, a CLIP-based model (Radford et al., 2021) fine-tuned on the LAION-5B corpus (Schuhmann et al., 2022). For automatic image caption generation, we employ the BLIP model *blip-image-captioning-base* (Li et al., 2022). Then, to estimate semantic similarity between generated captions and prompts, we embed both with the *all-MiniLM-L6-v2* SentenceBERT model (Reimers and Gurevych, 2019) and compute cosine similarities for the resulting representations. For the standard VQA analysis, we evaluate three models with strong reasoning performance: the open model *Qwen2.5-VL-32B-Instruct* (Qwen2.5-VL) (Bai et al., 2025) and the two proprietary models *gpt-4o-mini-2024-07-18* (GPT-4o mini) and *gpt-5-2025-08-07* (GPT-5) (OpenAI et al., 2024b). Finally, for the advanced VQA approach, we use *Llama-3.3-70B-Instruct* (Grattafiori et al., 2024) to generate suit-

Category	Image Quality	Subject	Temporal
Landscape	0.8160	1.0	0.9122
Animals	0.5508	1.0	0.9189
Buildings	0.2801	0.7626	1.0
Maps	0.2412	0.9054	1.0
Artworks	0.3080	1.0	0.9192
Average	0.5016	0.8970	0.9502

Table 2: Inter Annotator Agreement of the Human Evaluation. We present Cohen’s k for each question individually and the average per category.

able question-answer pairs. Additionally, we use *Qwen2.5-VL-32B-Instruct* to predict the answer given the generated image.

5.2 Results: Human Evaluation

Inter-annotator Agreement We calculate Cohen’s κ (Cohen, 1960) for each category and question individually and show the results in Table 2.

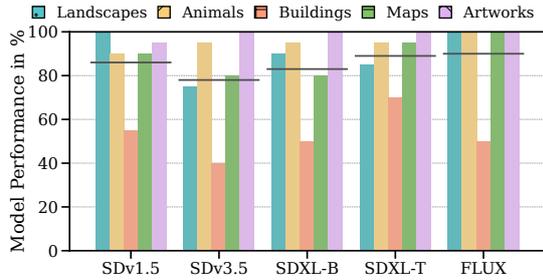
Overall, we observe an average inter-annotator agreement of 0.78 Cohen’s κ across the three questions and categories. Encouragingly, for the categories most relevant to our study, i.e., whether the subject is present in the image and whether temporal knowledge was applied correctly, we achieve an average agreement of at least 0.76. The lowest agreement, 0.5 κ , is found for the question concerning image quality. A qualitative analysis revealed differing intuitions between the two annotators: while the first annotator considered only visibly distorted images as low quality, the second also classified non-photorealistic images as visual generation errors. In the final label aggregation, we adopted the more liberal criterion of the first annotator (Q1). Disagreements among the annotators for the other questions (Q2, Q3) were resolved through guided discussions with one of the authors.

Temporal Knowledge in T2I models Utilizing the harmonized judgments obtained through our annotation study, we now discuss the T2I models’ generations regarding their quality, subject presence, and application of temporal knowledge.

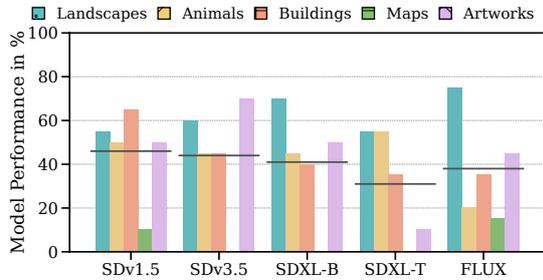
Figure 2 reports model accuracy on the three questions. For Q1 (image quality), SDXL-T attains the lowest score, with only on average 32% of images generated without visible generation errors, whereas FLUX achieves the highest average accuracy of 87% across categories. The *Maps* category emerges as the most challenging; even FLUX yields only 50% images without visual errors, and both SDXL-T and SDXL-B fail to produce any



(a) **Q1:** % of images that do not contain any errors.



(b) **Q2:** % of images that present the correct subject.



(c) **Q3:** % of images that correctly apply the required temporal knowledge.

Figure 2: Accuracy per Category in % of images that correctly represent the respective concept. We present the results for each tested T2I model separately. Grey bars indicate mean accuracy per model.

coherent outputs in this setting. By contrast, the *Artworks* category shows the largest share of images without apparent errors. This pattern could partially be explained by the fact that annotators could interpret generation errors within this category as stylistic variation within artistic imagery, which can mask underlying generation defects.

Turning to **Q2** (subject presence), we find that all models generally demonstrate high accuracy in correctly generating the requested subject (e.g., a horse) of most categories, but struggle for the *Artworks*, where correct depiction of the subject is closely coupled with the application of temporal knowledge such as era and style requirements. This observation suggests that the basic object-level complexity of the prompts is not the primary performance bottleneck for the majority of categories.

Finally, in Figure 2c we present the results for

Q3 (temporal knowledge). Performance is uniformly low across categories, indicating that models struggle to integrate temporal cues. Notably, FLUX, despite leading on **Q1** and **Q2**, attains the second temporal accuracy at 38%, whereas SD v1.5, although producing less coherent images overall, achieves the highest average at 46%. As expected, all models score lowest for the *Maps* category, underscoring the difficulty of the task. Strikingly, even for supposedly simpler categories such as *Landscapes* and *Animals*, models do not exceed 75% temporal accuracy. Overall, these findings indicate the following: (i) General image quality performance is not an indicator of the temporal knowledge encoded in the models. (ii) Different categories demand distinct types of temporal understanding, each posing varying levels of difficulty for text-to-image (T2I) models – but seemingly easier categories are not always easiest for all models. (iii) Finally, current T2I models remain limited in their ability to encode and apply temporal knowledge in a consistent and systematic manner – even when succeeding in subject placement.

In Figure 3, we show an example from SDXL-B for the subject *Dwarf Hippo* across four temporal cues. In this example, the model largely captures age progression across the typical lifespan, and presumably even tries to depict a deceased or extinct-looking specimen in the right image.

5.3 Results: Automatic Evaluation

Next, we investigate the ability of automatic methods to judge the correct use of temporal cues specified in the prompt. As ground truth, we use the labels from our human annotation study. We compare traditional similarity metrics with vision–language judges under multiple prompting strategies.

Traditional Approaches Table 3 shows the results for the two traditional approaches: ClipScore and the caption-based evaluation. Overall, both methods are largely ineffective at capturing the fine-grained temporal cues targeted in our study. CLIPScore shows no significant correlation with any of the annotated labels. We hypothesize that these results are largely caused by CLIP’s limited sensitivity to temporal expressions and their subtle visual consequences. This is consistent with prior work showing that metrics such as CLIPScore struggle with fine-grained or nuanced visual attributes (Bugliarello et al., 2023; Wiles et al., 2025). Moreover, our use of CLIPScore compares



Figure 3: Example images generated by SDXL-B. We present an example from the Animals category for the subject *Dwarf Hippo*, shown at different stages of its life cycle that serve as temporal cues.

Approach	Dimension	Question	PBS
ClipScore	Subject	Q2	0.0158
	Temporal	Q3	-0.0329
	Subj+Temp	Q2+Q3	-0.0394
Captioning	Subject	Q2	0.1571*
	Temporal	Q3	-0.1841*
	Subj+Temp	Q2+Q3	-0.0955*

Table 3: Results for traditional evaluation approaches. We present point biserial correlations with human annotations for **Q2** and **Q3**, plus their combined judgment. Asterisks indicate statistical significance.

each generated image to a single reference text, a setting that is less informative than CLIP’s original retrieval-style use case in which multiple candidate captions are ranked relative to one another, which likely also contributes to the weak association even for subject presence. The one notable exception is CLIPScore’s relatively better performance on the *Artworks* category, which we hypothesize stems from artworks being particularly well-represented in CLIP’s training data and from temporal cues being more tightly coupled with subject identity in this domain (e.g., “Renaissance artworks”).

In contrast to the CLIPScore, prompt–caption similarity yields significant correlations, but these gains primarily reflect subject presence (**Q2**), with the strongest effects in *Artworks* and *Buildings*. In contrast, correlations for the temporal cues (**Q3**) are negative, indicating that caption-based measures can mislead when correctness depends on temporal cues rather than object identity.

Decompositional VQA Figure 4 shows the macro-F1 scores of decompositional VQA across different thresholds for correctly answered LLM-generated questions. The method is evaluated against the combined human judgment for **Q2** and **Q3**, assessing whether the full prompt is accurately represented in the image. For *Animals*, the method fails to surpass the baseline at any decision threshold. Across the remaining categories, macro-

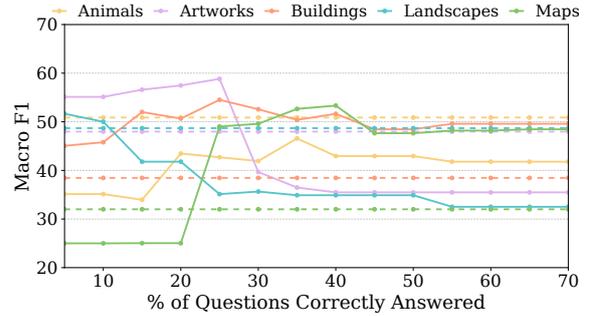


Figure 4: Decompositional VQA results. We report macro F1 scores (y-axis) across categories for varying answer correctness thresholds (x-axis). Each threshold indicates the minimum percentage of correctly answered questions required to classify an image as successfully reflecting the temporal prompt. Dashed lines indicate the random baseline for each category.

F1 never exceeds 60%, and no single question-accuracy threshold yields consistently strong performance. Taken together, these patterns indicate that the approach is not yet adequate for reliable automatic assessment of temporal cues in images.

Direct VLM Judging We report aggregated results for the direct VLM judging approach in Table 4. Finer-grained results are provided in Appendix B. We evaluate two prompting styles: *Simple*, which asks whether the image correctly depicts the prompt, and *Instructed*, which provides task guidance and requests a rating. For each style, we consider zero-shot and three-shot variants.

Encouragingly, all three models exceed the random baseline across all prompting strategies. GPT-5 attains the strongest overall performance but peaks at a macro F1 of 72%, well below reliable accuracy. For both GPT-5 and GPT-4o mini, the *Simple* prompts outperform the *Instructed* prompts, indicating that while the models grasp the task without guidance, the relevant temporal cues might be too subtle and remain difficult to assess consistently. Surprisingly, Qwen-VL even outperforms GPT-4o mini with the *Instructed* prompt. Moreover, few-

Model	Strategy	Prompt	P	R	F1
Baseline			48.4	48.3	47.7
Qwen-VL	0-Shot	Simple	64.3	60.1	60.0
		Instr.	64.8	64.4	64.6
	3-Shot	Simple	61.9	62.7	62.0
		Instr.	65.9	67.2	65.8
GPT-4o mini	0-Shot	Simple	64.4	61.2	61.5
		Instr.	61.0	59.5	59.8
	3-Shot	Simple	62.7	61.7	62.0
		Instr.	59.3	58.0	58.1
GPT-5	0-Shot	Simple	73.2	66.0	66.8
		Instr.	75.5	68.3	69.4
	3-Shot	Simple	75.7	70.9	72.0
		Instr.	74.6	69.5	70.5

Table 4: Direct VLM Judging: We present macro Precision (P), Recall (R), and F1 against human annotations for the combined judgments of **Q2** and **Q3** (full prompt).

shot prompting with out-of-distribution exemplars yields gains for all models, with the largest increase for GPT-5. Overall, even the most recent proprietary system falls well short of robust performance, indicating that the VQA-as-judge paradigm is not yet sufficient for dependable automatic evaluation of temporal knowledge in generated images.

Error Analysis Taking a closer look, Figure 5 shows GPT-5’s failure distribution to examine systematic patterns. Despite human annotations identifying *Buildings* and *Maps* as the most challenging categories, GPT-5 most frequently misjudges the presumably simpler *Animals* category across T2I models. Conversely, *Maps* exhibits the fewest errors, suggesting that its temporal cues (or their absence) are easier for the model to detect. Apart from an outlier in *Landscapes* for SDXL-T, which aligns with the low image quality noted in human evaluation, errors are otherwise relatively evenly distributed. Overall, GPT-5’s evaluation reliability is not strictly correlated with human-perceived task difficulty. Instead, the model appears more sensitive to visually ambiguous or less distinctive temporal cues (e.g., aging in animals) than to structurally constrained ones (e.g., historical maps).

6 Conclusion

We introduced TEMPVIZ, the first dataset specifically designed to evaluate the temporal knowledge in T2I across a diverse set of temporal dimensions. Using TEMPVIZ, we studied the ability of five T2I models to incorporate temporal knowledge into

Animals	6	12	8	11	6
Artworks	6	1	7	5	5
Buildings	4	2	4	6	3
Landscapes	3	4	11	5	2
Maps	3	2		2	
	FLUX	SDXL	BSDXL	T SDv1.5	SDv3.5

Figure 5: GPT-5 error distribution across categories and T2I models for the best performing prompting strategy.

image generation. To this end, we conduct a human evaluation and additionally compare the obtained judgments against several automated judging schemes. We find that all T2I models we tested show only low to moderate capabilities in applying the required knowledge across the different temporal categories of TEMPVIZ. We further find that automatic evaluation approaches fail to reliably judge the temporal cues depicted in images and that their mistakes do not necessarily align with those of humans. In sum, **there is thus a pressing need for both improving the understanding as well as the generation of temporal cues in images**. We hope that our findings and our novel TEMPVIZ resource can motivate further research in this direction.

Limitations

Our work has several limitations that need to be acknowledged. First, due to resource constraints, our work is limited to five T2I models. However, we have decided to use frequently deployed models that exhibit state-of-the-art performance and hope that our benchmark will also be used to evaluate further T2I models. Second, for the same reason, we relied on only two annotators to evaluate the generated images, which may introduce bias or limit the robustness of the results. Third, while our evaluation approach offers a starting point, more sophisticated and varied evaluation metrics could provide a deeper understanding of the models’ capabilities in integrating temporal knowledge in the image generation process. Fourth, while the reference images we provide aid dataset annotation, they serve only as a silver standard and cannot provide sufficient coverage for reference-based evaluation. Landscape prompts such as “A lake in Europe” are inherently too diverse to capture with single reference images. For Animals, we typically have only one image per species, lacking the multiple aging stages required to match our temporal

prompts. Similarly, Buildings often have reference images for only one temporal state (e.g., the historical structure before deconstruction) without corresponding images of later states. Finally, certain categories in our dataset require specific domain knowledge to judge accurately, which could impact the reliability of those evaluations. However, by providing reference images for more difficult categories, we are confident that we can minimise these effects.

Acknowledgements

The work of Carolin Holtermann and Anne Lauscher is funded by the Excellence Strategy of the German Federal Government and the Federal States.

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. *Qwen2.5-vl technical report*. *Preprint*, arXiv:2502.13923.
- Emanuele Bugliarello, Laurent Sartran, Aishwarya Agrawal, Lisa Anne Hendricks, and Aida Nematzadeh. 2023. *Measuring progress in fine-grained vision-and-language understanding*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1559–1582, Toronto, Canada. Association for Computational Linguistics.
- Wenhu Chen, Xinyi Wang, and William Yang Wang. 2021. *A dataset for answering time-sensitive questions*. *Preprint*, arXiv:2108.06314.
- J. Cohen. 1960. *A coefficient of agreement for nominal scales*. *Educational and Psychological Measurement*, 20:37–46.
- Bhuwan Dhingra, Jeremy R. Cole, Julian Martin Eisenschlos, Daniel Gillick, Jacob Eisenstein, and William W. Cohen. 2022. *Time-aware language models as temporal knowledge bases*. *Transactions of the Association for Computational Linguistics*, 10:257–273.
- Bahare Fatemi, Mehran Kazemi, Anton Tsitsulin, Karishma Malkan, Jinyeong Yim, John Palowitch, Sungyong Seo, Jonathan Halcrow, and Bryan Perozzi. 2024. *Test of time: A benchmark for evaluating llms on temporal reasoning*. *Preprint*, arXiv:2406.09170.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. *The llama 3 herd of models*. *Preprint*, arXiv:2407.21783.
- Wes Gurnee and Max Tegmark. 2024. *Language models represent space and time*. *Preprint*, arXiv:2310.02207.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. *CLIPScore: A reference-free evaluation metric for image captioning*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Carolin Holtermann, Paul Röttger, and Anne Lauscher. 2025. *Around the world in 24 hours: Probing LLM knowledge of time and place*. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22875–22897, Vienna, Austria. Association for Computational Linguistics.
- Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A Smith. 2023. *Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering*. *arXiv preprint arXiv:2303.11897*.
- Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. 2023. *T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation*. *Preprint*, arXiv:2307.06350.
- Black Forest Labs. 2023. Flux. <https://github.com/black-forest-labs/flux>.
- Seongyun Lee, Seungone Kim, Sue Hyun Park, Geewook Kim, and Minjoon Seo. 2024. *Prometheus-vision: Vision-language model as a judge for fine-grained evaluation*. *Preprint*, arXiv:2401.06591.
- Tony Lee, Michihiro Yasunaga, Chenlin Meng, Yifan Mai, Joon Sung Park, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Benita Teufel, Marco Bellagente, Minguk Kang, Taesung Park, Jure Leskovec, Jun-Yan Zhu, Li Fei-Fei, Jiajun Wu, Stefano Ermon, and Percy Liang. 2023. *Holistic evaluation of text-to-image models*. *Preprint*, arXiv:2311.04287.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. *Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation*. *Preprint*, arXiv:2201.12086.
- Jason Xinyu Liu, Ziyi Yang, Ifrah Idrees, Sam Liang, Benjamin Schornstein, Stefanie Tellex, and Ankit Shah. 2023. *Grounding complex natural language commands for temporal tasks in unseen environments*. *Preprint*, arXiv:2302.11649.

- Alex C. Ma, Ayana D. Cameron, and Martin Wiener. 2024. [Memorability shapes perceived time \(and vice versa\)](#). *Nature Human Behaviour*, 8(7):1296–1308. Publisher: Nature Publishing Group.
- Yuwei Niu, Munan Ning, Mengren Zheng, Weiyang Jin, Bin Lin, Peng Jin, Jiaqi Liao, Chaoran Feng, Kunpeng Ning, Bin Zhu, and Li Yuan. 2025. [Wise: A world knowledge-informed semantic evaluation for text-to-image generation](#). *Preprint*, arXiv:2503.07265.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024a. [GPT-4 Technical Report](#). *arXiv preprint*. ArXiv:2303.08774.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024b. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Dong Huk Park, Samaneh Azadi, Xihui Liu, Trevor Darrell, and Anna Rohrbach. 2021. [Benchmark for compositional text-to-image synthesis](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. [Sdxl: Improving latent diffusion models for high-resolution image synthesis](#). *Preprint*, arXiv:2307.01952.
- James Pustejovsky, Patrick Hanks, Roser Saurí, Andrew See, Rob Gaizauskas, Andrea Setzer, Dragomir Radev, Beth Sundheim, David Day, Lisa Ferro, and Marcia Lazo. 2003. The timebank corpus. *Proceedings of Corpus Linguistics*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *International Conference on Machine Learning*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. 2022a. [High-Resolution Image Synthesis with Latent Diffusion Models](#). In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, New Orleans, LA, USA. IEEE.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022b. [High-resolution image synthesis with latent diffusion models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. [Laion-5b: an open large-scale dataset for training next generation image-text models](#). In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Qingyu Tan, Hwee Tou Ng, and Lidong Bing. 2023. [Towards benchmarking and improving the temporal reasoning capability of large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14820–14835, Toronto, Canada. Association for Computational Linguistics.
- Qingyu Tan, Hwee Tou Ng, and Lidong Bing. 2024. [Towards robust temporal reasoning of large language models via a multi-hop qa dataset and pseudo-instruction tuning](#). *Preprint*, arXiv:2311.09821.
- Shivin Thukral, Kunal Kukreja, and Christian Kavouras. 2021. [Probing language models for understanding of temporal expressions](#). *Preprint*, arXiv:2110.01113.
- Yuqing Wang and Yun Zhao. 2024. [Tram: Benchmarking temporal reasoning for large language models](#). *Preprint*, arXiv:2310.00835.
- Olivia Wiles, Chuhan Zhang, Isabela Albuquerque, Ivana Kajić, Su Wang, Emanuele Bugliarello, Yasumasa Onoe, Pinelopi Papalampidi, Ira Ktena, Chris Knutsen, Cyrus Rashtchian, Anant Nawalgaria, Jordi Pont-Tuset, and Aida Nematzadeh. 2025. [Revisiting text-to-image evaluation with gecko: On metrics, prompts, and human ratings](#). *Preprint*, arXiv:2404.16820.
- Siheng Xiong, Ali Payani, Ramana Kompella, and Faramarz Fekri. 2024. [Large language models can learn temporal reasoning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10452–10470, Bangkok, Thailand. Association for Computational Linguistics.
- Michal Yarom, Yonatan Bitton, Soravit Changpinyo, Roei Aharoni, Jonathan Herzig, Oran Lang, Eran

Ofek, and Idan Szpektor. 2023. [What you see is what you read? improving text-image alignment evaluation.](#) *Preprint*, arXiv:2305.10400.

Ruilin Zhao, Feng Zhao, Guandong Xu, Sixiao Zhang, and Hai Jin. 2022. [Can language models serve as temporal knowledge bases?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2024–2037, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A TempViz Creation

A.1 Prompt Formulation Variations

ID	Prompt Template
1	<i>Give me a photorealistic image of {subject} in {temporal cue}.</i>
2	<i>Produce a photorealistic image of {subject} in {temporal cue}.</i>
3	<i>Generate a photorealistic image of {subject} in {temporal cue}.</i>
4	<i>Create a photorealistic image of {subject} in {temporal cue}.</i>

Table 5: Generic prompt templates used for each subject s_i and temporal cue t_i combination $s_i \times t_i$ in TEMPVIZ.

A.2 Distribution Maps Category

Figure 6 shows the distribution of regions we use to prompt the Maps category of TempViz.

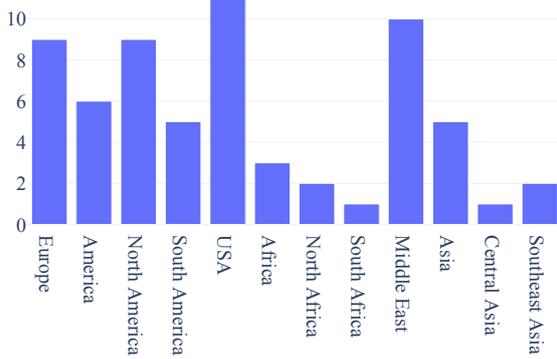


Figure 6: Number of suggestions per region in the Maps category of TempViz

A.3 Reference Images

Category	Source	License
Animals, Buildings	Wikidata https://www.wikidata.org/wiki/Wikidata:Licensing	Creative Commons C0 https://creativecommons.org/publicdomain/zero/1.0/
Maps	Open History Map https://map.openhistorymap.org/	Creative Commons BY 4.0 https://creativecommons.org/licenses/by/4.0/
Artworks	Google Arts and Culture https://artsandculture.google.com/	The licenses under which Google Arts and Culture makes the images available are very opaque, so we have decided to make these images available in our dataset only through links to the corresponding website.

Table 6: A list of the licenses under which the images in our dataset were published.

A.4 Expected Values

Art period	Expected Values
Ancient Art	symbolic, preserved, symbolism, craftsmanship, ritual, early civilization, religious and methodological themes, symbolic imagery
Baroque Art	dramatic, ornate, grandiose, lavish, opulent, vivid, ornamentation, grandeur, lavishness, opulence, vibrancy, grandeur in scale and detail, elaborate ornamentation, complex composition
Contemporary Art	abstract, experimental, minimalist, abstraction, minimalism, diverse forms, diverse cultural perspectives, experimental techniques
Impressionism	luminous, expressive, vibrant, fluid, evocative, luminosity, expression, vibrancy, fluidity, effects of light and color, impression of a scene, scenes of everyday life, landscapes, urban settings, light and weather, less detail
Medieval Art	religious, ornamental, gothic, byzantine, devotion, symbolism, ornamentation, monumentality, strong religious influence, ecclesiastical purpose, christian themes, romanesque, spiritual aesthetic, two-dimensional aesthetic
Modern Art	religious, ornamental, gothic, byzantine, devotion, symbolism, ornamentation, monumentality, strong religious influence, ecclesiastical purpose, christian themes, romanesque, spiritual aesthetic, two-dimensional aesthetic
Neoclassicism	formal, realistic, monumental, sculptural, formality, realism, monumentality, sculptural, clarity, elegant simplicity, formal composition, restrained emotional expression, ideals of harmony, proportion and disciplined technique, enlightenment
Post-Impressionism	vibrant, expressive, symbolic, structured, geometric, bold, vibrancy, symbolism, geometry, boldness, focus on light and color, emphasis on deeper symbolic content, structural forms, vivid and emotive use of line and color, emotional depth, structured compositions, expressive, sometimes geometric forms
Realism	detailed, naturalistic, mundane, detail, authenticity, observation, naturalism, depicting everyday scenes, depicting people in a true-to-life manner, meticulous detail, naturalistic depiction of subjects, still lifes
Renaissance Art	harmonious, realistic, detailed, classical, proportional, luminous, harmony, detail, realistic representation, precise anatomy, linear perspective, depth
Rococo	ornamental, lavish, fluid, delicate, sensuous, lavishness, delicacy, elaborate ornamentation, intricate detailing, elegant and decorative aesthetics, lightness, curvaceous forms, themes of love, themes of nature, soft colors, asymmetrical designs, sculptural elements
Romanticism	emotive, sublime, passionate, mystical, exotic, sublimity, passion, mysticism, exoticism, sublime beauty of nature

Table 7: Expected values for each art periods in the Artworks prompt category

Prompted Lifespan	Expected Values
1 month	tiny, soft, infant, juvenile, suckling, offspring, baby, small, bright-eyed, smooth, bright eyes, wide and open eyes, cuddly and endearing presence
0.5*lifespan	mature, vigorous, strong, healthy, adult, muscular, toned, fit, agile, vibrant, sturdy, strong and healthy, muscular and toned body, vitality and good health, vibrant and sturdy appearance, prime of their life
0.8*lifespan	elderly, aged, grey, wrinkled, signs of aging, graying fur, thinning fur, weathered appearance, wrinkled skin
2*lifespan	lifeless, deceased, dead, late, passed, perished, death, passing, decease, mortality

Table 8: Expected values for each lifespan in the Animals prompt category

Prompted Lifespan	Expected Values
demolished	ruined, dilapidated, shattered, decayed, wrecked, ravaged, devastated, demolished, destroyed, torn off, rubble, debris, ruins, wreckage, fragments, remains, state of ruin

Table 9: Expected values for the current state "*demolished*" in the Buildings prompt category

Season	Expected Values
Winter, december, january, february	frosty, snowy, icy, snow, ice, frost, frosty and icy
Spring, march, april, may	blooming, lush, blossoming, green, greenery, flowers, bloom, fresh greenery, trees and plants covered in buds and blossoms, nature's renewal
Summer, june, july, august	sunny, bright, green, warm, sunshine, greenery, warmth, blue skies, bright sunshine
Fall, september, october, november	golden leaves, colored leaves, orange leaves, warm-hued landscape

Table 10: Expected values for the different seasons in the Landscape prompt category

B Experimental Details

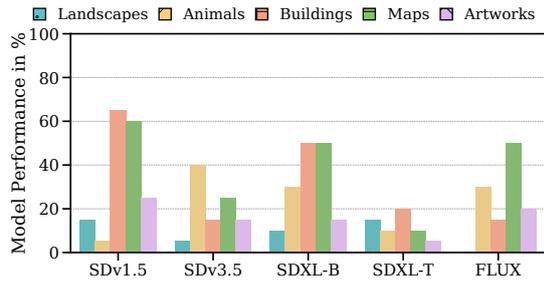
Table 11 shows all T2I models we analyze and details on their architecture. All models that we use for our experiments are loaded using the Hugging Face Transformers library⁶ and we run all experiments on 4 NVIDIA RTX A6000 GPUs.

Modelname	Text Encoder	Diffuser
FLUX Dev	CLIP-G/14, CLIP-L/14, T5 XXL	MM-DiT
Stable Diffusion XL - Base	CLIP ViT-L/14, OpenCLIP ViT-bigG/14	Latent Diffusion U-Net
Stable Diffusion XL - Turbo	CLIP ViT-L/14, OpenCLIP ViT-bigG/14	Latent Diffusion U-Net (distilled, ADD)
Stable Diffusion v1.5	CLIP ViT-L/14	Latent Diffusion U-Net
Stable Diffusion v3.5	CLIP G/14, CLIP L/14, T5 XXL	MMDiT (Diffusion Transformer)

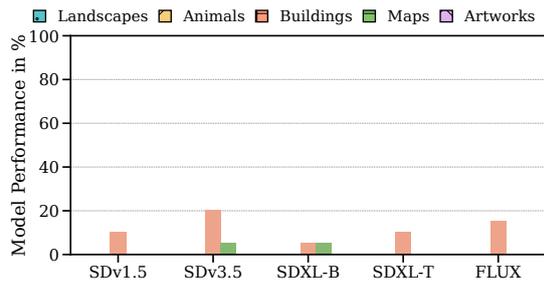
Table 11: Architectures of T2I models.

⁶<https://huggingface.co/docs/transformers>

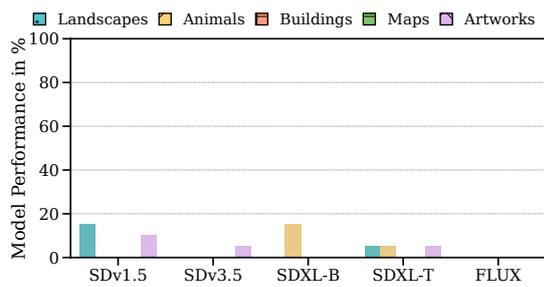
C Annotator Disagreement



(a) **Q1:** Disagreement in % for image quality errors.



(b) **Q2:** Disagreement in % for correct subject.



(c) **Q3:** Disagreement in % for temporal knowledge.

Figure 7: Annotator Disagreement in % of images for the respective annotation question.

D Automatic Evaluation

D.1 Prompting Strategies

Decompositional VQA We use the following prompt to generate suitable question-image pairs from each of the original T2I input prompts.

LLM Prompt

Given an image description, generate multiple-choice questions that can verify if the image description is correct for an image. First, extract elements from the image description. Then classify each element into a category (object, human, animal, food, activity, attribute, counting, color, material, spatial, location, shape, other). Finally, generate different questions for each element, open-ended and multiple choice and also the correct answer. For example: 'What is the main subject of the image? Answer: Pony' Description: '{}

Example Questions by Llama3.3 70B In the following, we present example questions for the prompt “Create a photorealistic image of an Aardvark that is one month old.” generated by Llama3.3 70B.

Example Questions by the LLM

[...] ### Questions for Each Element:
For Aardvark (Animal):
1. **Open-ended question:** What type of animal is the main subject of the image? -
Answer: Aardvark
2. **Multiple-choice question:** What is the main animal in the image? - **A)**
Lion - **B)** Elephant - **C)** Aardvark
- **D)** Giraffe - **Correct answer:** C)
Aardvark
For One month old (Attribute - Age):
1. **Open-ended question:** How old is the Aardvark in the image? - **Answer:**
One month old
2. **Multiple-choice question:** What is the age of the Aardvark depicted in the image? - **A)** One week old - **B)**
One month old - **C)** Three months old
- **D)** Six months old - **Correct answer:** B) One month old
[...]

Direct VLM Judging We present the two prompting strategies that we use to guide the VLMs to analyze the generated image with respect to the input prompt.

Standard Prompt

Is the image correct for the following prompt: {} Is the correct temporal knowledge applied (i.e., <temporal cue>)? Just reply with 'yes' or 'no'

Judging Prompt

As a professional “Text-to-Image” quality assessor, your task is to judge the performance of a text-image model w.r.t. a certain criterion by evaluating the image generated from a specific prompt. The criteria for evaluation are as follows: Rubrics: 'Assess how accurately the image reflects the given prompt. Check if the object attributes along with all temporal elements—such as signs of aging in animals, historical map borders, seasonal indicators in landscapes, lighting and shadows for time of day, architectural features that change over time, and stylistic traits of artworks from different epochs—are correctly represented.' Please analyze step by step and provide the RATING using the following scale: [“Not accept”, “Accept”]. In this scale, “Not accept” represents images that do not accurately reflect the prompt. “Accept” represents images that do accurately reflect the prompt, along and also the temporal elements. Please do not generate any other opening, closing, or explanations. The output of the analysis and rating should strictly adhere to the following format: ANALYSIS: Provide your analysis here RATING: Only provide your rating here. Now, proceed with evaluating the image based on the prompt: '{}'

D.2 Additional silver-standard T2I Evaluation Results

We provide results on an additional set of 500 images for the combined questions Q2 + Q3 (that is, whether an image matches the prompt completely). This subset was labeled by two team members and one crowdsourced annotator to create a potential training set and was not included in the main paper because the annotations were less stringent and were collected only at the full prompt level. Given the limited size of the evaluation set, we release this additional subset as a silver standard benchmark to complement the primary annotations and further support our findings.

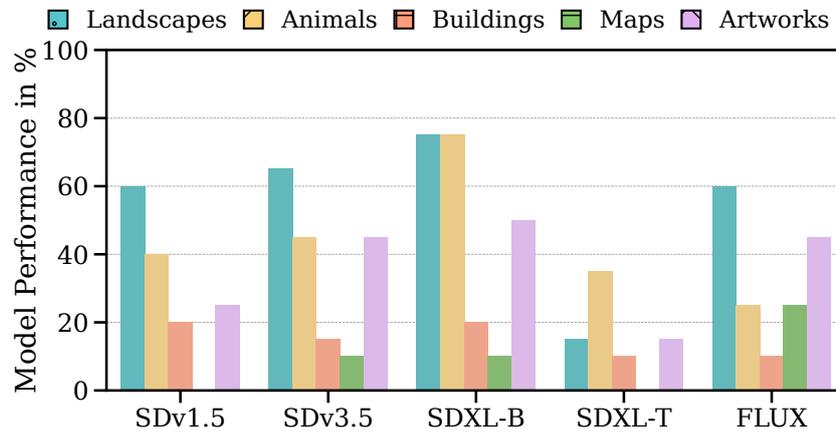


Figure 8: Accuracy per Category in % of images that correctly represent the whole prompt. We present the results for each tested T2I model separately.

D.3 Additional Results for the Direct VLM Judging Approach

Model	Strategy	Prompt	Category	P	R	F1
Baseline			Overall	48.4	48.3	47.7
Qwen-VL	0-Shot	Simple	Landscapes	63.2	56.8	44.0
		Simple	Animals	79.3	51.2	39.3
		Simple	Artworks	74.3	72.6	70.8
		Simple	Buildings	50.3	50.3	50.0
		Simple	Maps	68.4	68.4	68.4
		Instr.	Landscapes	60.9	60.6	56.0
		Instr.	Animals	70.6	63.6	62.6
		Instr.	Artworks	75.5	67.1	61.8
		Instr.	Buildings	50.5	50.5	50.5
		Instr.	Maps	52.3	54.7	52.5
	3-Shot	Simple	Landscapes	66.85	67.03	66.93
		Simple	Animals	58.14	52.18	44.59
		Simple	Artworks	73.81	69.39	66.02
		Simple	Buildings	52.16	52.91	50.91
		Simple	Maps	50.79	53.68	44.77
		Instr.	Landscapes	69.00	70.38	68.47
		Instr.	Animals	78.86	75.78	76.40
		Instr.	Artworks	71.98	68.28	65.11
		Instr.	Buildings	47.61	46.64	42.70
Instr.	Maps	52.40	58.95	50.17		

Table 12: Direct VQA judging per Category: We present macro Precision (P), Recall (R), and F1 against human annotations for the combined judgments of **Q2** and **Q3** (full prompt).

Model	Strategy	Prompt	Category	P	R	F1
Baseline			Overall	48.4	48.3	47.7
GPT-4o mini	0-Shot	Simple	Landscapes	59.15	56.74	47.83
		Simple	Animals	76.27	63.75	61.97
		Simple	Artworks	70.73	70.81	70.76
		Simple	Buildings	57.51	54.15	53.93
		Simple	Maps	52.62	59.47	50.76
		Instr.	Landscapes	60.64	58.32	50.16
		Instr.	Animals	75.76	58.66	53.94
		Instr.	Artworks	74.27	73.33	71.90
		Instr.	Buildings	47.47	48.08	47.44
	Instr.	Maps	51.49	56.32	47.38	
	3-Shot	Simple	Landscapes	63.42	62.53	56.89
		Simple	Animals	69.66	56.61	51.31
		Simple	Artworks	76.29	71.41	68.08
		Simple	Buildings	55.68	52.63	51.64
		Simple	Maps	51.19	55.26	46.32
		Instr.	Landscapes	60.12	58.56	51.52
		Instr.	Animals	62.71	51.52	41.22
		Instr.	Artworks	73.67	73.13	71.96
		Instr.	Buildings	68.36	56.75	56.86
Instr.	Maps	48.09	40.53	39.27		
GPT-5	0-Shot	Simple	Landscapes	67.72	68.08	63.99
		Simple	Animals	54.17	50.66	40.71
		Simple	Artworks	72.11	72.32	71.96
		Simple	Buildings	89.69	56.52	55.79
		Simple	Maps	47.45	48.95	48.19
		Instr.	Landscapes	68.06	67.85	63.00
		Instr.	Animals	71.11	57.80	53.22
		Instr.	Artworks	78.79	78.79	78.00
		Instr.	Buildings	88.89	52.17	47.92
	Instr.	Maps	47.45	48.95	48.19	
	3-Shot	Simple	Landscapes	76.95	78.49	74.88
		Simple	Animals	53.56	52.42	50.18
		Simple	Artworks	76.40	76.57	75.99
		Simple	Buildings	90.10	58.70	59.32
		Simple	Maps	47.45	48.95	48.19
		Instr.	Landscapes	69.23	69.43	65.00
		Instr.	Animals	71.82	63.22	61.90
		Instr.	Artworks	73.94	73.84	73.00
Instr.		Buildings	89.69	56.52	55.79	
Instr.	Maps	47.40	47.89	47.64		

Table 13: Direct VQA judging per Category: We present macro Precision (P), Recall (R), and F1 against human annotations for the combined judgments of **Q2** and **Q3** (full prompt).

D.4 Additional Results on Model Errors

Animals - realistic age	3	3	7	7	2
Animals - unrealistic age	2	4	1	4	2
Animals - youth	1	5	0	0	2
Artworks - period	0	1	0	0	0
Artworks - year	6	0	7	5	5
Buildings - after	1	1	0	1	0
Buildings - before	0	1	2	0	1
Buildings - random after	0	0	1	1	0
Buildings - random before	3	0	1	4	2
Landscapes - Daytime	0	0	0	1	0
Landscapes - Seasonality	3	4	11	4	2
Maps - maps	3	2	0	2	0
	FLUX	SDXL-B	SDXL-T	SDv1.5	SDv3.5

Figure 9: GPT-5 error distribution across temporal categories and T2I models for the best performing prompting strategy.