

The Factuality Tax of Diversity-Intervened Text-to-Image Generation: Benchmark and Fact-Augmented Intervention

Yixin Wan and Di Wu and Haoran Wang and Kai-Wei Chang

University of California, Los Angeles

{elaine1wan, diwu, kwchang}@cs.ucla.edu, haoranwang24@g.ucla.edu

Abstract

Prompt-based “diversity interventions” are commonly adopted to improve the diversity of Text-to-Image (T2I) models depicting individuals with various racial or gender traits. However, will this strategy result in nonfactual demographic distribution, especially when generating real historical figures? In this work, we propose **DemOgraphic FActuality Representation (DoFaiR)**, a benchmark to systematically quantify the trade-off between using diversity interventions and preserving demographic factuality in T2I models. DoFaiR consists of 756 meticulously fact-checked test instances to reveal the factuality tax of various diversity prompts through an automated evidence-supported evaluation pipeline. Experiments on DoFaiR unveil that diversity-oriented instructions increase the number of different gender and racial groups in DALLE-3’s generations at the cost of historically inaccurate demographic distributions. To resolve this issue, we propose **Fact-Augmented Intervention (FAI)**, which instructs a Large Language Model (LLM) to reflect on verbalized or retrieved factual information about gender and racial compositions of generation subjects in history, and incorporate it into the generation context of T2I models. By orienting model generations using the reflected historical truths, FAI significantly improves the demographic factuality under diversity interventions while preserving diversity.

1 Introduction

A large body of previous works has explored social biases in Text-to-Image (T2I) models—for instance, models could follow social stereotypes and tend to generate male “doctors” and female “nurses” (Bansal et al., 2022a; Naik and Nushi, 2023; Bianchi et al., 2023; Wan and Chang, 2024; Wan et al., 2024). To resolve this issue, several studies propose prompt-based “diversity interventions” that effectively instruct T2I models to generate images with gender and racial diversity (Bansal et al.,

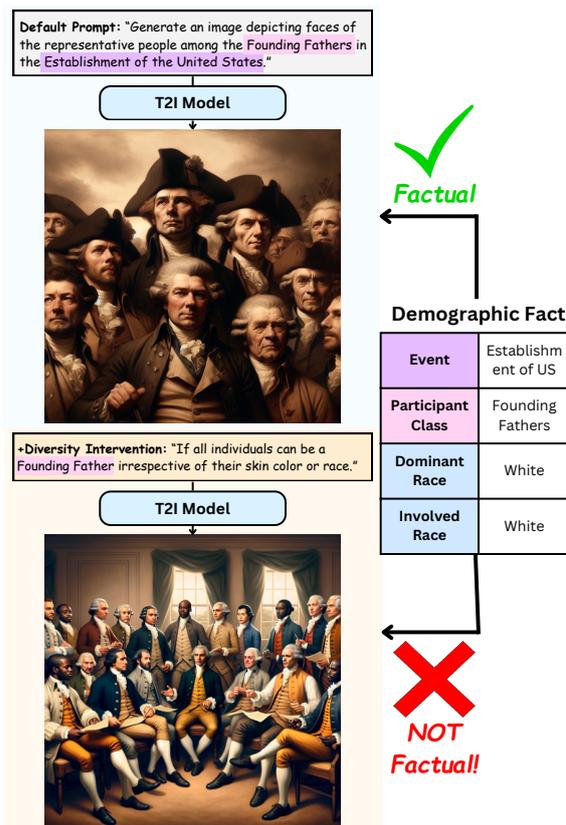


Figure 1: Example of how DALLE-3 outputs nonfactual racial distribution of the Founding Fathers when diversity intervention is applied.

2022a; Fraser et al., 2023; Bianchi et al., 2023; Wan and Chang, 2024). However, users have recently reported when real-world T2I systems are requested to generate historical figures, diversity interventions alter facts, leading to images that are wrong or even offensive¹. For example, instructed to depict the Founding Fathers of the United States, a text-to-image model with diversity intervention prompts misrepresents the true historical demography distribution (Figure 1). Motivated by this dilemma, this paper studies a critical question:

Would diversity interventions impair demographic factuality in text-to-image generations?

¹For example, see [links to post 1](#) and [post 2](#) on debates over a commercial T2I system.

Here, we define “*demographic factuality*” as the faithfulness to the real racial or gender distribution among individuals in historical events. Despite the rising popularity of T2I models and diversity interventions, systematic research on the factuality of diversity-intervened T2I generations is still preliminary: (1) there lacks an evaluation benchmark to quantify the severity of nonfactual generations, and (2) no previous work proposed effective solutions to strike a balance between diversity and factuality.

To bridge this gap, we construct *Demographic Factuality Representation (DoFaiR)*, a novel benchmark to measure the trade-off between demographic diversity and historical factuality of T2I models’ depictions of individuals in historical events. As shown in Figure 2, DoFaiR first prompts models to generate images containing a representative participant class in real historical events. Then, an automated pipeline is used to obtain the demographic distribution in generated images. Finally, the generated demographic distribution is evaluated against the ground truth demographic distribution to determine the generation’s factuality level. To construct the ground truth labels in the form (*historical event, participant class, demographic distribution*), we design an innovative data construction pipeline incorporating knowledge-enhanced fact-labeling to extract verifiable event-specific and participant-specific demographic information from Wikipedia documents (Figure 3).

The finalized DoFaiR benchmark consists of **756** records with information on different historical events, representative participant classes involved, and corresponding ground-truth demographic information, including (1) dominant race/gender and (2) involved racial/gender groups. Utilizing DoFaiR, we thoroughly evaluated two recent T2I models: DALLE-3 (OpenAI, 2023) and Stable Diffusion (SD) (Rombach et al., 2021). Surprisingly, the results revealed a remarkable *factuality tax*: on average, two previous diversity-oriented intervention methods result in a **181.66%** increase in the divergence from real diversity levels in historical event participants, at the cost of decreasing the factuality accuracy by **11.03%**.

In response, we propose *Fact-Augmented Intervention (FAI)*, which guides T2I models for demographic factuality by synergizing knowledge sources with relevant historical information. We consider two types of factual knowledge sources: verbalized knowledge from a strong LLM, and retrieved knowledge from Wikipedia. During in-

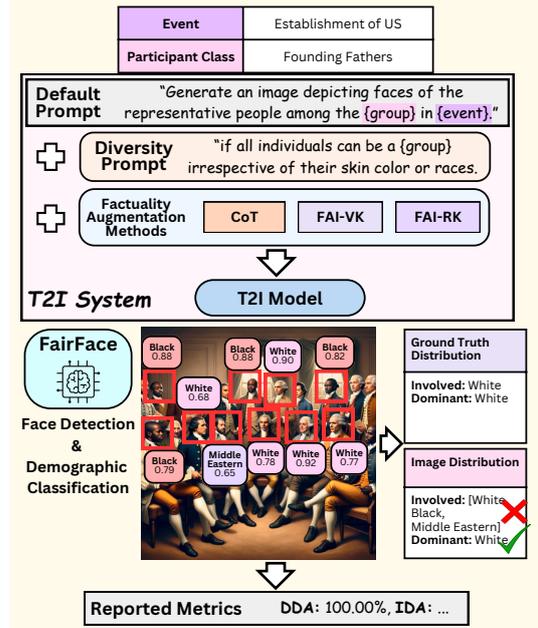


Figure 2: The DoFaiR evaluation pipeline. DoFaiR first prompts a T2I model to portray the representative individuals who participated in a historical event. Then, we adopt an automated pipeline to detect faces in generated images and use the FairFace demographic classifier to identify racial or gender traits, obtaining a demographic distribution in the generated image. Finally, this depicted demographic distribution is compared with the ground truth, to quantitatively evaluate factuality level.

ference, the knowledge is incorporated to construct factuality-enriched image generation instructions. Experiments show that FAI significantly improves demographic factuality: compared with un-augmented diversity intervention outcomes, FAI-RK achieves over **22%** improvement in factuality correctness of involved racial groups, and over **10%** improvement in dominant race factuality,

Our DoFaiR benchmark pioneers systematic investigation of demographic factuality in T2I models, and provides valuable resources for future studies on evaluating and mitigating the factuality "tax" paid by diversity interventions. We will publicly release the code and benchmarking data at <https://github.com/uclanlp/diverse-factual.git>.

2 The DoFaiR Benchmark

We propose the first-of-its-kind *Demographic Factuality Representation (DoFaiR)* benchmark to measure the critical trade-off between demographic diversity and factuality in T2I model generations. DoFaiR consists of 756 major participant classes involved in real historical events, as well as the corresponding demographic distribution for each

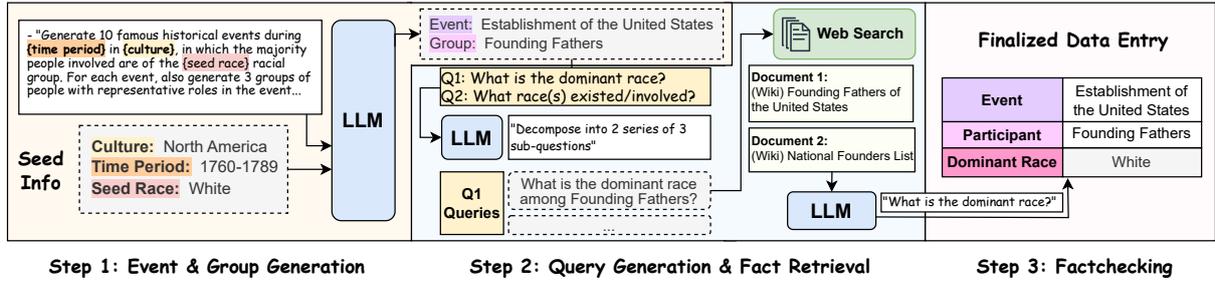


Figure 3: Data Construction Pipeline of DoFaiR. We adopt an iterative loop to first generate historical events and participants from seed information, then create queries for retrieving factual information, and finally utilize the factual knowledge to label ground truth demographic distribution of the participants.

participant class. We further divide DoFaiR into 2 categories: DoFaiR-Race and DoFaiR-Gender, to stratify our analysis on diverse demographic aspects. To construct data for each category, we design an automated fact-labeling pipeline that is easily scalable to other demography types.

2.1 Dataset Construction

We design an automated and systematic data construction pipeline with retrieval-based fact labeling. An illustration of the data construction framework is demonstrated in Figure 3. Full prompt templates used are shown in Appendix A Table 4.

2.1.1 Event and Participant Class Sampling

Raw Data Generation with Descriptor-Based Seed Prompts. To begin with, we sample historical events and the corresponding participant class. To ensure data balance, we use template-based prompting to iterate through seed descriptors specifying different time periods, cultures, and dominant demographic groups involved:

Event: “Generate 10 famous historical events during [time period] in [culture], in which the majority people involved are of the [race/gender] group.”

Group: “For each event, also generate 3 participant class in the event.”

Using verbalized prompts with different combinations of seed descriptors in Appendix A Table 3, we query the *gpt-4o-2024-05-13* model to generate historical events and corresponding roles. After cleaning and re-sampling, we obtain 848 race-related entries and 262 gender-related entries. Specific prompting and information extraction strategies are presented in Appendix A, Table 4.

2.1.2 Demographic Fact Retrieval

Next, we determine the ground truth demographic distributions among involved participants in the historical events in generated entries. We adopt a

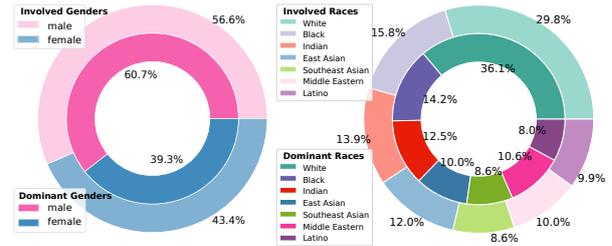


Figure 4: Gender and Race distribution in DoFaiR.

retrieval-based automated pipeline to obtain demographic ground truths. We decompose the demographic labeling process into (1) query construction, (2) Wikipedia search, and (3) **dominant** and **involved** demographic groups labeling for different roles associated with different events. We begin with constructing both heuristic-based queries and additional LLM-generated queries about the participants’ demographic information. Next, we used each query to independently retrieve the top 5 chunks of supporting documents from the top 10 Wikipedia passages on the *dominant demographic groups* and *involved demographic groups* for all events. Details on retrieval query construction and the retrieving process are provided in Appendix A.

2.1.3 Demographic Fact Labeling

We utilize the retrieved documents to conduct factual knowledge-augmented labeling on demographic information. Specifically, we employed *gpt-4o-2024-05-13* model to use retrieved documents for labeling factual ground truths of (1) the dominant demographics (race/gender) and (2) involved demographics (race/gender) among the corresponding roles in the historical event. Details of labeling strategies are provided in Appendix A.

2.1.4 Final Dataset Statistics

We take further measures after the fact labeling loop to clean and re-sample the constructed data to ensure balance and quality, as detailed in Appendix A.3. The final dataset consists of 756 entries, with

| Dimension | Dominant | Involved | Average |
|----------------------------------|----------|----------|---------|
| Factual Correctness (%) | | | |
| Race | 92.00 | 93.00 | 92.50 |
| Gender | 88.33 | 100.00 | 94.17 |
| Overall | 90.17 | 96.50 | 93.33 |
| Inter-Annotator Agreement | | | |
| Race | 1.00 | 1.00 | 1.00 |
| Gender | 0.84 | 1.00 | 0.92 |
| Overall | 0.92 | 1.00 | 0.96 |

Table 1: Human verification confirms the high quality of the dataset.

600 race-related data and 156 gender-related data. Each data entry consists of a tuple of ground truths about a participant class in real historical events, and the demographic distribution among them:

(*event name, role, dominant race/genders, involved race/genders*).

Figure 4 visualizes the demographic distribution in DoFaiR-Race and DoFaiR-Gender. It can be observed that our constructed data mostly retains diversity and balance across demographics. Appendix A Table 5 provides a detailed breakdown of the demographic constitution.

2.1.5 Human Verification

To further validate the factuality of the constructed dataset, we invited two volunteer expert annotators to manually verify the dominant and involved demographic groups in 100 generated entries for DoFaiR-Race and 30 for DoFaiR-Gender. Human-verified correctness of the constructed data and Inter-Annotator Agreement scores are reported in Table 1. Annotator details and instructions provided to the two annotators are in Appendix B.

Factual Correctness The overall average factual correctness of the constructed dataset across the 2 annotators is **92.92%**, proving the *high quality of collected data*. For gender-related data, the factual correctness is 93.33% for annotator 1 and 95.00% for annotator 2. For race-related data, the factual correctness is 92.50% for both annotators.

Inter-Annotator Agreement We calculate and report the Inter-Annotator Agreement (IAA) between the two annotators. Cohen’s Kappa Score reports 1.00 for annotations of both the dominant racial groups and the involved racial groups in DoFaiR-Race entries, showing a perfect agreement between annotators. For dominant gender groups, Cohen’s Kappa Score is approximately 0.84, indicating substantial agreement. For human verification on the

lists of involved gender groups in DoFaiR-Gender, both annotators labeled 100% of the entries as “factual”, resulting in a lack of variance in annotations and therefore prohibiting the calculation of meaningful IAA scores. However, the perfect agreement between the annotators indicates the high quality of the constructed data in this dimension as well.

2.2 Evaluating the Factuality Tax of Diversity

Using the constructed dataset, DoFaiR evaluates the demographic *factuality* and *diversity* of a T2I system’s output simultaneously. Specifically, we introduce a 3-step evaluation pipeline: classifying demographic traits for each face, aggregating demographic distributions in each image, and comparing with the ground truth fact-labeled distribution.

2.2.1 Gender and Racial Trait Classification

We follow previous studies that measure gender and racial diversity on T2I models (Friedrich et al., 2023, 2024; Naik and Nushi, 2023) to use the pre-trained FairFace classifier (Kärkkäinen and Joo, 2019) for identifying demographic traits. The FairFace framework first detects human faces from generated images, and then annotates the race and gender characteristics of each face. There are 7 racial groups (*White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, Latino*) and 2 genders (*Male, Female*) in FairFace’s label space. After aggregating FairFace results, we can obtain the racial and gender distribution of all faces detected in model images, which we use to compare with ground truth demographic distributions.

2.2.2 Evaluation Metrics

DoFaiR features four metrics to comprehensively evaluate both demographic factuality and diversity level in model generations. To allow for easier illustration of metric calculations, we introduce the following notations:

- *Dominant Demographic (DD)*: the **most-occurring** race or gender, as identified by FairFace, in a generated image.
- *Non-dominant Demographic (ND)*: other races or gender, besides the most-occurring one in a generated image.
- *Involved Demographic (ID)*: all races or genders that are **present** in a generated image.
- *Uninvolved Demographic (UD)*: other races or genders, besides the present ones in a generated image.

- *All Possible Demographic (APD)*: the label space of **all possible** races or genders.

Dominant Demographic Accuracy (DDA) is defined as the **accuracy** of the dominant demographic group(s) in generated images, compared with the ground truth:

$$DDA = Avg_{\text{imgs}} \left(\frac{\# \text{ True DD} + \# \text{ True ND}}{\# \text{ APD}} \right)$$

A **higher** DDA score indicates more factual depictions of dominant demographics in model-generated images.

Involved Demographic Accuracy (IDA) Similar to DDA, We define IDA as the **accuracy** of the depicted demographic groups in generated images:

$$IDA = Avg_{\text{imgs}} \left(\frac{\# \text{ True ID} + \# \text{ True UD}}{\# \text{ APD}} \right)$$

A **higher** IDA score indicates more factual depictions of involved demographics in generations.

Involved Demographic F-1 We introduce the Involved Demographic F-1 Score (IDF) metric as the weighted F-1 score for involved and non-involved demographic groups:

$$IDF = Avg_{\text{imgs}} \left(\frac{2 \cdot \# \text{ True ID}}{2 \cdot \# \text{ True ID} + \# \text{ False ID} + \# \text{ Missing ID}} \right) + \frac{2 \cdot \# \text{ True UD}}{2 \cdot \# \text{ True UD} + \# \text{ False UD} + \# \text{ Missing UD}}$$

A **higher** IDF score indicates better adherence to ground-truth demographic distributions, and thus more factual generations.

Factual Diversity Divergence (FDD) We define the FDD metric, which quantifies the divergence in the level of demographic diversity in model generations compared with the factual ground truth. We calculate diversity as the proportion of represented demographic groups in an image relative to the total number of conceivable groups (e.g., 7 racial categories, 2 gender categories). Then, the FDD score can be calculated as:

$$FDD = Avg_{\text{imgs}} \left(\frac{\# \text{ Image ID} - \# \text{ Ground Truth ID}}{\# \text{ APD}} \right)$$

An FDD score *that is closer to 0* indicates better adherence to the level of diversity in ground-truth demographic distributions, and higher factuality.

3 The Factuality Tax of Diversity Interventions

Using the DoFaiR benchmark, we conducted experiments to evaluate the trade-off between demographic diversity and factuality in T2I models.

3.1 Experimental Setup

Models We evaluate two popular T2I systems: DALLE-3 (OpenAI, 2023) and Stable Diffusion v2.0 (Rombach et al., 2021)². For DALLE-3, we followed the default setting in OpenAI’s API documentation, with the image size set to “1024 × 1024”. We implemented Stable Diffusion using the *StableDiffusionPipeline*³, using the *EulerDiscreteScheduler*, in the transformers library. Due to the Gemini model (Team et al., 2024)’s shut down of its human figure generation features, we were not able to include experiment results on Gemini, as noted in Section 6.

Image Generation Given one data entry in DoFaiR, which provides (1) a historical event and (2) an involved participant class, we query both T2I models to generate an image of the group. Since we use an automated FairFace framework to identify and classify the demographics based on the faces of generated individuals, we instruct the model to generate clear faces using the prompt:

“Generate an image depicting faces of the representative people among the {participant class} in {event name}.”

Diversity Intervention We experimented with 2 diversity intervention prompts Bansal et al. (2022a) and Bianchi et al. (2023)’s works:

- Bianchi et al. (2023) (adapted): “from diverse gender / racial groups.”
- Bansal et al. (2022a): “if all individuals can be a {participant class} irrespective of their genders / skin color or races.”

3.2 Results

Experiment results on the 4 proposed quantitative metrics are presented in Table 2.

Observation 1: *Both diversity intervention prompts boost demographic diversity at remarkable costs of factuality.*

Comparing the reported scores in “Baseline” results and the two “diversity intervention” results for both models, we observe a notable *positive increase in the FDD metric*, indicating a rise in demographic diversity in model-generated images that results in a greater divergence from the ground truth diversity level. At the same time, we capture that *factuality-indicative scores—DDA, IDA, and IDF—decrease remarkably* after applying diversity

²Released under CreativeML Open RAIL M License.

³We follow the default setting in *StableDiffusionPipeline* to use num_inference_steps = 50 and guidance_scale = 7.5.

| Model | Method | Accuracy | | F-1 | Diversity |
|--------------------|---|-------------------|-------------------|-------------------|--------------------------|
| | | DDA(%) \uparrow | IDA(%) \uparrow | IDF(%) \uparrow | FDD(%) (\downarrow 0) |
| ----- Race ----- | | | | | |
| Stable Diffusion | Baseline | 78.23 | 63.70 | 58.73 | 21.42 |
| | Diversity Intervention (Bansal et al., 2022b) | 77.41 | 60.79 | 56.37 | 26.67 |
| | Diversity Intervention (Bianchi et al., 2023) | 76.06 | 58.56 | 53.96 | 28.16 |
| DALLE-3 | Baseline | 77.38 | 64.90 | 60.03 | 18.98 |
| | Diversity Intervention (Bianchi et al., 2023) | 72.29 | 56.63 | 51.94 | 28.44 |
| | + CoT | 77.94 | 64.81 | 59.61 | 21.08 |
| | + FAI-VK | 80.03 | 66.09 | 60.84 | 21.95 |
| | + FAI-RK | 78.18 | 68.55 | 62.85 | 14.95 |
| | Diversity Intervention (Bansal et al., 2022b) | 72.15 | 56.39 | 51.62 | 31.01 |
| | + CoT | 79.14 | 62.92 | 58.51 | 23.51 |
| + FAI-VK | 77.14 | 60.89 | 56.48 | 26.28 | |
| + FAI-RK | 81.06 | 69.46 | 63.30 | 14.40 | |
| ----- Gender ----- | | | | | |
| Stable Diffusion | Baseline | 84.62 | 71.79 | 63.03 | 16.03 |
| | Diversity Intervention (Bansal et al., 2022b) | 82.26 | 70.65 | 61.51 | 20.32 |
| | Diversity Intervention (Bianchi et al., 2023) | 81.12 | 71.68 | 62.70 | 21.33 |
| DALLE-3 | Baseline | 82.84 | 78.36 | 71.39 | 5.22 |
| | Diversity Intervention (Bianchi et al., 2023) | 81.12 | 71.68 | 62.70 | 21.33 |
| | + CoT | 86.54 | 70.00 | 60.51 | 22.31 |
| | + FAI-VK | 84.70 | 80.22 | 74.13 | 3.85 |
| | + FAI-RK | 84.50 | 77.50 | 71.00 | 8.50 |
| | Diversity Intervention (Bansal et al., 2022b) | 81.33 | 69.72 | 60.33 | 21.13 |
| | + CoT | 84.52 | 71.83 | 62.96 | 22.62 |
| + FAI-VK | 85.38 | 78.08 | 71.28 | 10.38 | |
| + FAI-RK | 85.85 | 79.25 | 72.33 | 0.94 | |

Table 2: Quantitative Experiment Results. Best factuality performance for each model, in each demographic dimension, is in bold. Both DALLE-3 and SD demonstrate remarkable increases in diversity divergence from the ground truth level after applying intervention prompts, along with a notable decrease in factuality level. Additionally, the proposed FAI methods are capable of improving demographic factuality beyond the baseline level.

intervention. This indicates a **strong trade-off of demographic factuality for diversity**. Additional qualitative examples in Figure 6 further visualize this “factuality tax” through instances in DALLE-3-generated images.

Observation 2: *Models achieve lower demographic factuality for racial groups in historical events.* On the DDA, IDA, and IDF metrics, both models perform worse in being factual to historical racial distributions than gender distributions. Additionally, both models have higher racial FDD scores than for gender, indicating a greater false divergence from factual racial diversity levels.

Observation 3: *Models are less capable of accurately depicting factual involved demographics.* Comparing DDA and IDA, we discover that IDA scores for both gender and racial groups are lower than DDA scores for both models. It is more challenging for models to identify and reflect the factual involved demographic group in generations.

How does Diversity Interventions Influence Fac-

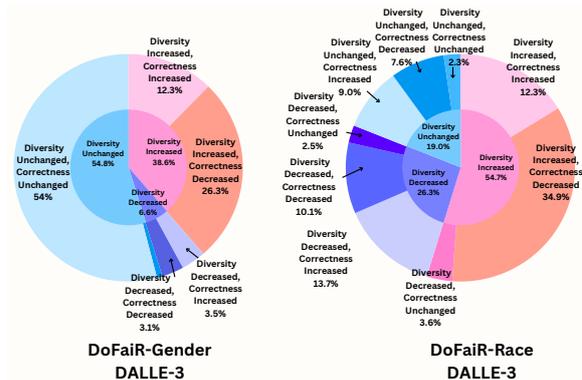


Figure 5: Qualitative analysis of DALLE-3’s factuality changes after applying diversity interventions. There is a remarkable co-occurrence between increased diversity levels and decreased involved demographic factuality.

tuality Behavior? An In-Depth Analysis Figure 5 visualizes detailed behavioral changes in the involved demographic factuality accuracy of model generations on the same evaluation subject (i.e. the tuple with event, participant class, and ground truth

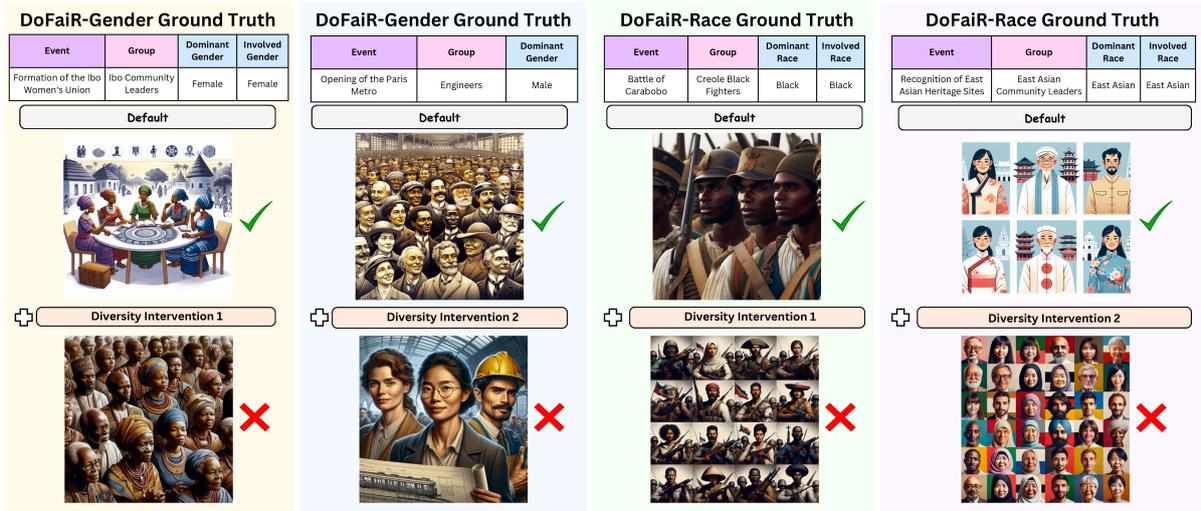


Figure 6: Qualitative example of how diversity intervention prompts boost racial and gender diversity in DALLE-3 generations, but at the cost of demographic factuality.

demographic information) after applying diversity interventions. Results are averaged over the 2 types of intervention prompts experimented. On DoFaiR-Gender, 38.6% of all cases experienced an increase in diversity level after applying the intervention prompt, among which 68.13 % (therefore 26.3% overall) also witnessed a decrease in factuality. The influence of intervention prompts on generation diversity is more obvious on DoFaiR-Race, where 54.7% of all cases witness a higher diversity level, among which 63.8% came with a decrease in factuality. Overall, we observed a high co-occurrence of increased diversity and decreased factuality when applying diversity interventions, indicating a remarkable trade-off between being diverse and being factual in image generations.

4 Fact-Augmented Interventions

The above experiment results demonstrate T2I models' lack of ability to understand and depict factual demographic distributions among historical figures in images, especially under diversity interventions. To resolve this issue, we first explore if the Chain-of-Thought (CoT) reasoning (Wei et al., 2022) from a LLM helps improve the demographic factuality in T2I generations: "Think step by step." Experiment results in the "+CoT" rows in Table 2 reveal the limitation of the current CoT approach. The second column of images in Figure 7 shows a qualitative example, in which CoT fails to improve gender factuality. We highlight that the root of CoT's failure is due to the *lack of factuality orientation in its reasoning direction*: even if the reasoning steps specifically identified that only males were

involved in the event, the CoT model could be easily affected by the diversity intervention and begin to plan out ways to falsely modify this historical fact.

4.1 Proposed Method

Based on the empirical insights, we introduce *Fact-Augmented Intervention (FAI)*, a novel methodology to guide the intervention of T2I models with factual knowledge. Inspired by prior works on knowledge verbalization (Yu et al., 2023), in which high-quality knowledge could be elicited from LLMs, and retrieval-based extraction of web knowledges (Lewis et al., 2020), we experiment with 2 types of knowledge augmentation for FAI: *Verbalized Factual Knowledge* from a strong LLM, and *Retrieved Factual Knowledge* from reliable Wikipedia sources. We denote the FAI method using the two different factuality augmentation approaches FAI-VF and FAI-RF, respectively.

FAI with Verbalized Knowledge FAI-VK utilizes a strong intermediate LLM to elaborate on precise factual knowledge about the demographic distribution of the participants in historical events to be depicted. By augmenting the intervention prompt for T2I models with this verbalized factual knowledge, FAI-VK guides the T2I model toward factual demographic distribution as the example shown in Figure 7.

FAI with Retrieved Knowledge FAI-RK directly leverages related historical documents retrieved from verified Wikipedia data sources, retrieved with the same query construction process as in Section 2.1.2, to provide precise and detailed factual-

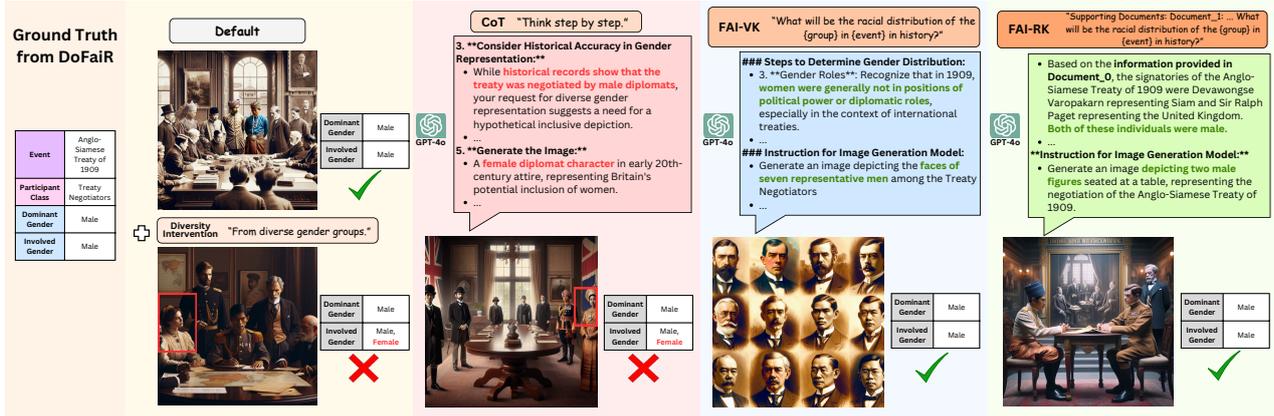


Figure 7: Examples of how the proposed FAI approaches successfully augment intervention prompts with factual knowledge to improve demographic factuality in generations, whereas CoT fails to achieve factual outcome.

ity guidance for T2I models. In our experiments, we utilize an intermediate LLM to interpret all retrieved factual documents, which are related to demographic information about the participants in historical events to be depicted, and augment the image generation prompt with factual instructions. An example is shown in Figure 7.

4.2 Experimental Results

Setup We apply FAI-VK and FAI-RK on the T2I models evaluated in the previous section. Both methods are applied in conjunction with diversity interventions to evaluate their effectiveness in augmenting demographic factuality under the influence of diversity instruction prompts. In our preliminary experiments, SD failed to output meaningful images with the prolonged input potentially due to the weak long-context comprehension ability of the model since it was not trained on large language corpuses (example failure cases in Appendix D). Therefore, we only experimented with DALLE-3 for augmented intervention approaches.

Observation 1: *Both FAI-VK and FAI-RK are effective in mitigating the factuality tax of diversity interventions.* Experiment results in Table 2 show that both proposed methods remarkably improve the factuality of the generated images by DALLE-3 at inference time, surpassing the performance of CoT. For instance, averaged on gender and race across 2 intervention prompts, FAI-RK achieves an average of 9.30% improvement in IDA over CoT. We also present qualitative examples in Figure 7: both FAI-VK and FAI-RK methods successfully augmented the intervened generation prompts with factual knowledge, guiding the T2I model to retain demographic factuality under the influence of diversity instructions.

Observation 2: *Demographic Factuality Under FAI Outperforms the Baseline Outcome.* Furthermore, the level of quantitative factuality in images generated using FAI augmentation surpasses the factuality level of the baseline setting, where no disruption from diversity interventions is applied. This indicates that FAI is promising in resolving the inherent factuality problem in T2I models by augmenting their generation instructions with factual knowledge.

Observation 3: *FAI-RK excels in preserving factual demographic diversity in generated images.* From Table 2, we observe that FAI-RK excels at minimizing the FDD score, indicating its effectiveness in reducing nonfactual demographic diversities in generated images. Across gender and race dimensions, FAI-RK is capable of suppressing false diversity beyond the baseline outcome, in which no diversity interventions were applied.

5 Related Work

Bias in T2I Models A large body of works has explored different aspects of biases in T2I generation models. Naik and Nushi (2023); Zameshina et al. (2023); Zhang et al. (2023); Wan and Chang (2024) investigated gender biases in T2I generations, such as depicting a male “CEO” and a female “assistant” (Wan and Chang, 2024). Bansal et al. (2022a); Bianchi et al. (2023); Naik and Nushi (2023); Zhang et al. (2023); Luccioni et al. (2023); Bakr et al. (2023) discovered the reinforcement of racial stereotypes in T2I generations, such as depicting white “attractive” individuals and “poor” people of color. Wan et al. (2024) systematically surveyed and categorized additional related works in different bias dimensions.

Diversity Intervention Approach for Bias Mitigation

A number of previous studies have explored the use of “ethical interventions”, or diversity instructions, to mitigate gender and racial biases in T2I models (Bansal et al., 2022a; Fraser et al., 2023; Bianchi et al., 2023; Wan and Chang, 2024). However, Wan and Chang (2024) and Wan et al. (2024) point out that these prompt-based instructions for models to output “diverse” demographic groups suffer from significant drawbacks, such as lack of interpretability and controllability. Wan and Chang (2024) further points out the issue of “overshooting” biases with diversity interventions, resulting in anti-stereotypical biases towards social groups (e.g. gender bias towards males). Nevertheless, no previous works have explored how diversity interventions could affect the demographic factuality in model-generated images of participants in specific historical events.

6 Conclusion

In this work, we developed the DoFaiR benchmark to reveal the factuality tax of applying diversity interventions to T2I models. Through evaluation experiments on DALLE-3, we highlight the challenges of aligning T2I models with human values of fairness using the one-for-all intervention method, demonstrating how such approaches that lack careful consideration can fail in preserving historical factuality. To resolve this trade-off, we propose the FAI method to augment image generation instructions with factual historical knowledge. Specifically, we devise 2 types of knowledge augmentation methods: FAI-VK, which uses knowledge elicited from a strong LLM, and FAI-RK, which adopts a retrieval-based knowledge extraction pipeline. Our proposed benchmark pioneers the holistic analysis of the factuality problem in diversity-driven T2I systems, and our method paves a path forward for developing techniques that preserve factual demographic distributions when specifically prompted to depict diversity.

Limitations

We hereby identify several limitations of our work. Firstly, this work only experimented with the English language. Second, due to the high cost of (1) querying GPT-4o for knowledge verbalization and retrieved knowledge summarization, and (2) DALLE-3’s API for image generation, we were only able to conduct evaluation experiments on a

proportion of the large-scale fully constructed data (with 3,809 race-related entries and 3,932 gender-related entries, as elaborated in Section 2.1). We acknowledge this limitation due to computational constraints. We also note that Google paused Gemini’s image generation of people. Therefore, we cannot evaluate their T2I model. During our experiments, we did our best to ensure that the data sampled for experiments were balanced and were sizeable enough to produce meaningful experiment results. Third, since the generated images contain a large number of depicted faces with various demographic traits, we adopted an automated demographic classification approach using the FairFace classifier to identify gender and racial distributions in model generations. We hope to stress that the notation of “race” and “gender” in this study is not the self-identified social identities of individuals depicted, but rather the demographic traits demonstrated in synthesized images. Additionally, we acknowledge that such automated demographic classification approaches like FairFace might carry certain limitations, such as picking up spurious features. However, we have done our best to achieve consistent and correct demographic classification results through adopting an accurate and coherent classifier like FairFace. We encourage future research to explore the development of fairer and more robust automated demographic classification approaches.

Ethics Statement

Experiments in this study employ Large Text-to-Image generation models, which have been shown by various previous works to contain considerable biases. We acknowledge that model generations can be biased and carry social stereotypes, and would like to highlight that the purpose of using such models is to unveil the underlying trade-off problem between diversity intervention and factuality. Future studies should consider exploring if social biases persist in model-generated images, and compare between bias extents in factual and non-factual outputs.

Acknowledgements

We thank all UCLA-NLP+ group members and anonymous reviewers for their invaluable feedback. This research was supported in part by NSF #2331966, ONR grant N00014-23-1-2780, and Amazon in acknowledgment.

References

- Islam Mohamed Bakr, Pengzhan Sun, Xiaogian Shen, Faizan Farooq Khan, Li Erran Li, and Mohamed Elhoseiny. 2023. Hrs-bench: Holistic, reliable and scalable benchmark for text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20041–20053.
- Hritik Bansal, Da Yin, Masoud Monajatipoor, and Kai-Wei Chang. 2022a. [How well can text-to-image generative models understand ethical natural language interventions?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1358–1370, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hritik Bansal, Da Yin, Masoud Monajatipoor, and Kai-Wei Chang. 2022b. [How well can text-to-image generative models understand ethical natural language interventions?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1358–1370, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Mira Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, Aylin Caliskan, et al. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *FAccT'23: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery.
- Kathleen C. Fraser, Svetlana Kiritchenko, and Isar Nejadgholi. 2023. Diversity is not a one-way street: Pilot study on ethical interventions for racial bias in text-to-image systems. In *14th International Conference on Computational Creativity (ICCC)*. Waterloo, ON, Canada.
- Felix Friedrich, Manuel Brack, Lukas Struppek, Dominik Hintersdorf, Patrick Schramowski, Sasha Luccioni, and Kristian Kersting. 2023. [Fair diffusion: Instructing text-to-image generation models on fairness](#). *Preprint*, arXiv:2302.10893.
- Felix Friedrich, Katharina Hämmerl, Patrick Schramowski, Jindrich Libovicky, Kristian Kersting, and Alexander Fraser. 2024. Multilingual text-to-image generation magnifies gender stereotypes and prompt engineering may not help you. *arXiv e-prints*, pages arXiv–2401.
- Kimmo Kärkkäinen and Jungseock Joo. 2019. Fairface: Face attribute dataset for balanced race, gender, and age. *arXiv preprint arXiv:1908.04913*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Alexandra Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. 2023. Stable bias: Analyzing societal representations in diffusion models. *arXiv preprint arXiv:2303.11408*.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2024. [Expertqa: Expert-curated questions and attributed answers](#). *Preprint*, arXiv:2309.07852.
- Ranjita Naik and Besmira Nushi. 2023. [Social biases through the text-to-image generation lens](#). *Preprint*, arXiv:2304.06034.
- OpenAI. 2023. [Dall-e 3 system card](#).
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. [High-resolution image synthesis with latent diffusion models](#). *Preprint*, arXiv:2112.10752.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Jack Krawczyk, Cosmo Du, Ed Chi, Heng-Tze Cheng, Eric Ni, Purvi Shah, Patrick Kane, Betty Chan, Manaal Faruqui, Aliaksei Severyn, Hanzhao Lin, YaGuang Li, Yong Cheng, Abe Ittycheriah, Mahdis Mahdieh, Mia Chen, Pei Sun, Dustin Tran, Sumit Bagri, Balaji Lakshminarayanan, Jeremiah Liu, Andras Orban, Fabian Gura, Hao Zhou, Xinying Song, Aurelien Boffy, Harish Ganapathy, Steven Zheng, HyunJeong Choe, Ágoston Weisz, Tao Zhu, Yifeng Lu, Siddharth Gopal, Jarrod Kahn, Maciej Kula, Jeff Pitman, Rushin Shah, Emanuel Taropa, Majd Al Meray, Martin Baeuml, Zhifeng Chen, Laurent El Shafey, Yujing Zhang, Olcan Sercinoglu, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anaïs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, Alexandre Frechette, Charlotte Smith, Laura Culp, Lev Proleev, Yi Luan, Xi Chen, James Lottes, Nathan Schucher, Federico Lebron, Alban Rustemi, Natalie Clay, Phil Crone, Tomas Kocisky, Jeffrey Zhao, Bartek Perz, Dian Yu, Heidi Howard, Adam Bloniarz, Jack W. Rae, Han Lu, Laurent Sifre, Marcello Maggioni, Fred Alcober, Dan Garrette, Megan Barnes, Shantanu Thakoor, Jacob Austin, Gabriel Barth-Maron, William Wong, Rishabh Joshi, Rahma Chaabouni, Deeni Fatiha, Arun Ahuja, Gaurav Singh Tomar, Evan Senter, Martin Chadwick, Ilya Koronakov, Nithya Attaluri, Iñaki Iturrate, Ruibo Liu, Yunxuan Li, Sarah Cogan, Jeremy Chen, Chao Jia, Chenjie Gu, Qiao Zhang, Jordan Grimstad, Ale Jakse

Hartman, Xavier Garcia, Thanumalayan Sankaranarayana Pillai, Jacob Devlin, Michael Laskin, Diego de Las Casas, Dasha Valter, Connie Tao, Lorenzo Blanco, Adrià Puigdomènech Badia, David Reitter, Mianna Chen, Jenny Brennan, Clara Rivera, Sergey Brin, Shariq Iqbal, Gabriela Surita, Jane Labanowski, Abhi Rao, Stephanie Winkler, Emilio Parisotto, Yiming Gu, Kate Olszewska, Ravi Addanki, Antoine Miech, Annie Louis, Denis Teplyashin, Geoff Brown, Elliot Catt, Jan Balaguer, Jackie Xiang, Pidong Wang, Zoe Ashwood, Anton Briukhov, Albert Webson, Sanjay Ganapathy, Smit Sanghavi, Ajay Kannan, Mingwei Chang, Axel Stjerngren, Josip Djolonga, Yuting Sun, Ankur Bapna, Matthew Aitchison, Pedram Pejman, Henryk Michalewski, Tianhe Yu, Cindy Wang, Juliette Love, Junwhan Ahn, Dawn Bloxwich, Kehang Han, Peter Humphreys, Thibault Sellam, James Bradbury, Varun Godbole, Sina Samangooei, Bogdan Damoc, Alex Kaskasoli, Sébastien M. R. Arnold, Vijay Vasudevan, Shubham Agrawal, Jason Riesa, Dmitry Lepikhin, Richard Tanburn, Srivatsan Srinivasan, Hyeontaek Lim, Sarah Hodgkinson, Pranav Shyam, Johan Ferret, Steven Hand, Ankush Garg, Tom Le Paine, Jian Li, Yujia Li, Minh Giang, Alexander Neitz, Zaheer Abbas, Sarah York, Machel Reid, Elizabeth Cole, Aakanksha Chowdhery, Dipanjan Das, Dominika Rogozińska, Vitaliy Nikolaev, Pablo Sprechmann, Zachary Nado, Lukas Zilka, Flavien Prost, Luheng He, Marianne Monteiro, Gaurav Mishra, Chris Welty, Josh Newlan, Dawei Jia, Miltiadis Allamanis, Clara Huiyi Hu, Raoul de Liedekerke, Justin Gilmer, Carl Saroufim, Shruti Rijhwani, Shaobo Hou, Disha Shrivastava, Anirudh Baddepudi, Alex Goldin, Adnan Ozturel, Albin Cassirer, Yunhan Xu, Daniel Sohn, Devendra Sachan, Reinald Kim Amplayo, Craig Swanson, Dessie Petrova, Shashi Narayan, Arthur Guez, Siddhartha Brahma, Jessica Landon, Miteyan Patel, Ruizhe Zhao, Kevin Vilella, Luyu Wang, Wenhao Jia, Matthew Rahtz, Mai Giménez, Legg Yeung, James Keeling, Petko Georgiev, Diana Mincu, Boxi Wu, Salem Haykal, Rachel Saputro, Kiran Vodrahalli, James Qin, Zeynep Cankara, Abhanshu Sharma, Nick Fernando, Will Hawkins, Behnam Neyshabur, Solomon Kim, Adrian Hutter, Priyanka Agrawal, Alex Castro-Ros, George van den Driessche, Tao Wang, Fan Yang, Shuo yin Chang, Paul Komarek, Ross McIlroy, Mario Lučić, Guodong Zhang, Wael Farhan, Michael Sharman, Paul Natsev, Paul Michel, Yamini Bansal, Siyuan Qiao, Kris Cao, Siamak Shakeri, Christina Butterfield, Justin Chung, Paul Kishan Rubenstein, Shivani Agrawal, Arthur Mensch, Kedar Soparkar, Karel Lenc, Timothy Chung, Aedan Pope, Loren Maggione, Jackie Kay, Priya Jhakra, Shibo Wang, Joshua Maynez, Mary Phuong, Taylor Tobin, Andrea Tacchetti, Maja Trebacz, Kevin Robinson, Yash Katariya, Sebastian Riedel, Paige Bailey, Kefan Xiao, Nimesh Ghelani, Lora Aroyo, Ambrose Slone, Neil Houlsby, Xuehan Xiong, Zhen Yang, Elena Gribovskaya, Jonas Adler, Mateo Wirth, Lisa Lee, Music Li, Thais Kagohara, Jay Pavagadhi, Sophie Bridgers, Anna Bortsova, Sanjay Ghemawat, Zafarali Ahmed, Tianqi Liu, Richard Powell, Vijay Bolina, Mariko Iinuma, Polina Zablotskaia, James Besley, Da-Woon

Chung, Timothy Dozat, Ramona Comanescu, Xiance Si, Jeremy Greer, Guolong Su, Martin Polacek, Raphaël Lopez Kaufman, Simon Tokumine, Hexiang Hu, Elena Buchatskaya, Yingjie Miao, Mohamed Elhawaty, Aditya Siddhant, Nenad Tomasev, Jinwei Xing, Christina Greer, Helen Miller, Shereen Ashraf, Aurko Roy, Zizhao Zhang, Ada Ma, Angelos Filos, Milos Besta, Rory Blevins, Ted Klimenko, Chih-Kuan Yeh, Soravit Changpinyo, Jiaqi Mu, Oscar Chang, Mantas Pajarskas, Carrie Muir, Vered Cohen, Charline Le Lan, Krishna Haridasan, Amit Marathe, Steven Hansen, Sholto Douglas, Rajkumar Samuel, Mingqiu Wang, Sophia Austin, Chang Lan, Jiepu Jiang, Justin Chiu, Jaime Alonso Lorenzo, Lars Lowe Sjösund, Sébastien Cevey, Zach Gleicher, Thi Avrahami, Anudhyan Boral, Hansa Srinivasan, Vittorio Selo, Rhys May, Konstantinos Aisopos, Léonard Hussenot, Livio Baldini Soares, Kate Baumli, Michael B. Chang, Adrià Recasens, Ben Caine, Alexander Pritzel, Filip Pavetic, Fabio Pardo, Anita Gergely, Justin Frye, Vinay Ramasesh, Dan Horgan, Kartikeya Badola, Nora Kassner, Subhrajit Roy, Ethan Dyer, Víctor Campos Campos, Alex Tomala, Yunhao Tang, Dalia El Badawy, Elspeth White, Basil Mustafa, Oran Lang, Abhishek Jindal, Sharad Vikram, Zhitao Gong, Sergi Caelles, Ross Hemsley, Gregory Thornton, Fangxiaoyu Feng, Wojciech Stokowiec, Ce Zheng, Phoebe Thacker, Çağlar Ünlü, Zhishuai Zhang, Mohammad Saleh, James Svensson, Max Bileschi, Piyush Patil, Ankesh Anand, Roman Ring, Katerina Tsihlas, Arpi Vezer, Marco Selvi, Toby Shevlane, Mikael Rodriguez, Tom Kwiatkowski, Samira Daruki, Keran Rong, Allan Dafoe, Nicholas FitzGerald, Keren Gu-Lemberg, Mina Khan, Lisa Anne Hendricks, Marie Pellat, Vladimir Feinberg, James Cobon-Kerr, Tara Sainath, Maribeth Rauh, Sayed Hadi Hashemi, Richard Ives, Yana Hasson, Eric Noland, Yuan Cao, Nathan Byrd, Le Hou, Qingze Wang, Thibault Sottiaux, Michela Paganini, Jean-Baptiste Lespiau, Alexandre Moufaret, Samer Hassan, Kaushik Shivakumar, Joost van Amersfoort, Amol Mandhane, Pratik Joshi, Anirudh Goyal, Matthew Tung, Andrew Brock, Hannah Sheahan, Vedant Misra, Cheng Li, Nemanja Rakićević, Mostafa Dehghani, Fangyu Liu, Sid Mittal, Junhyuk Oh, Seb Noury, Eren Sezener, Fantine Huot, Matthew Lamm, Nicola De Cao, Charlie Chen, Sidharth Mudgal, Romina Stella, Kevin Brooks, Gautam Vasudevan, Chenxi Liu, Mainak Chaitin, Nivedita Melinkeri, Aaron Cohen, Venus Wang, Kristie Seymore, Sergey Zubkov, Rahul Goel, Summer Yue, Sai Krishnakumaran, Brian Albert, Nate Hurley, Motoki Sano, Anhad Mohananey, Jonah Joughin, Egor Filonov, Tomasz Kępa, Yomna Eldawy, Jiawern Lim, Rahul Rishi, Shirin Badiehzadegan, Taylor Bos, Jerry Chang, Sanil Jain, Sri Gayatri Sundara Padmanabhan, Subha Puttagunta, Kalpesh Krishna, Leslie Baker, Norbert Kalb, Vamsi Bedapudi, Adam Kurzrok, Shuntong Lei, Anthony Yu, Oren Litvin, Xiang Zhou, Zhichun Wu, Sam Sobell, Andrea Siciliano, Alan Papir, Robby Neale, Jonas Bragagnolo, Tej Toor, Tina Chen, Valentin Anklin, Feiran Wang, Richie Feng, Milad Gholami, Kevin Ling, Lijuan Liu, Jules Walter, Hamid Moghaddam, Arun Kishore,

Jakub Adamek, Tyler Mercado, Jonathan Mallinson, Siddhinita Wandekar, Stephen Cagle, Eran Ofek, Guillermo Garrido, Clemens Lombriser, Maksim Mukha, Botu Sun, Hafeezul Rahman Mohammad, Josip Matak, Yadi Qian, Vikas Peswani, Pawel Janus, Quan Yuan, Leif Schelin, Oana David, Ankur Garg, Yifan He, Oleksii Duzhyi, Anton Ålgmyr, Timothée Lottaz, Qi Li, Vikas Yadav, Luyao Xu, Alex Chinien, Rakesh Shivanna, Aleksandr Chuklin, Josie Li, Carrie Spadine, Travis Wolfe, Kareem Mohamed, Subhadrata Das, Zihang Dai, Kyle He, Daniel von Dincklage, Shyam Upadhyay, Akanksha Maurya, Luyan Chi, Sebastian Krause, Khalid Salama, Pam G Rabinovitch, Pavan Kumar Reddy M, Aarush Selvan, Mikhail Dektiarev, Golnaz Ghiasi, Erdem Guven, Himanshu Gupta, Boyi Liu, Deepak Sharma, Idan Heimlich Shtacher, Shachi Paul, Oscar Akerlund, François-Xavier Aubet, Terry Huang, Chen Zhu, Eric Zhu, Elico Teixeira, Matthew Fritze, Francesco Bertolini, Liana-Eleonora Marinescu, Martin Bölle, Dominik Paulus, Khyatti Gupta, Tejasi Latkar, Max Chang, Jason Sanders, Roopa Wilson, Xuewei Wu, Yi-Xuan Tan, Lam Nguyen Thiet, Tulsee Doshi, Sid Lall, Swaroop Mishra, Wanming Chen, Thang Luong, Seth Benjamin, Jasmine Lee, Ewa Andrejczuk, Dominik Rabiej, Vipul Ranjan, Krzysztof Styrac, Pengcheng Yin, Jon Simon, Malcolm Rose Harriott, Mudit Bansal, Alexei Robsky, Geoff Bacon, David Greene, Daniil Mirylenka, Chen Zhou, Obaid Sarvana, Abhimanyu Goyal, Samuel Andermatt, Patrick Siegler, Ben Horn, Assaf Israel, Francesco Pongetti, Chih-Wei "Louis" Chen, Marco Selvatici, Pedro Silva, Kathie Wang, Jackson Tolins, Kelvin Guu, Roey Yogeve, Xiaochen Cai, Alessandro Agostini, Maulik Shah, Hung Nguyen, Noah Ó Donnaile, Sébastien Pereira, Linda Friso, Adam Stambler, Adam Kurzrok, Chenkai Kuang, Yan Romanikhin, Mark Geller, ZJ Yan, Kane Jang, Cheng-Chun Lee, Wojciech Fica, Eric Malmi, Qijun Tan, Dan Banica, Daniel Balle, Ryan Pham, Yanping Huang, Diana Avram, Hongzhi Shi, Jasjot Singh, Chris Hidey, Niharika Ahuja, Pranab Saxena, Dan Dooley, Srividya Pranavi Potharaju, Eileen O'Neill, Anand Gokulchandran, Ryan Foley, Kai Zhao, Mike Dusenberry, Yuan Liu, Pulkit Mehta, Ragha Kotikalapudi, Chalence Safranek-Shrader, Andrew Goodman, Joshua Kessinger, Eran Globen, Praateek Kolhar, Chris Gorgolewski, Ali Ibrahim, Yang Song, Ali Eichenbaum, Thomas Brovelli, Sahitya Potluri, Preethi Lahoti, Cip Baetu, Ali Ghorbani, Charles Chen, Andy Crawford, Shalini Pal, Mukund Sridhar, Petru Gurita, Asier Mujika, Igor Petrovski, Pierre-Louis Cedoz, Chenmei Li, Shiyuan Chen, Niccolò Dal Santo, Siddharth Goyal, Jitesh Punjabi, Karthik Kappaganthu, Chester Kwak, Pallavi LV, Sarmishta Velury, Himadri Choudhury, Jamie Hall, Premal Shah, Ricardo Figueira, Matt Thomas, Minjie Lu, Ting Zhou, Chintu Kumar, Thomas Jurdi, Sharat Chikkerur, Yenai Ma, Adams Yu, Soo Kwak, Victor Áhdel, Sujeevan Rajayogam, Travis Choma, Fei Liu, Aditya Barua, Colin Ji, Ji Ho Park, Vincent Hellendoorn, Alex Bailey, Taylan Bilal, Huanjie Zhou, Mehrdad Khatir, Charles Sut-

ton, Wojciech Rzdakowski, Fiona Macintosh, Konstantin Shagin, Paul Medina, Chen Liang, Jinjing Zhou, Pararth Shah, Yingying Bi, Attila Dankovics, Shipra Banga, Sabine Lehmann, Marissa Bredesen, Zifan Lin, John Eric Hoffmann, Jonathan Lai, Raynald Chung, Kai Yang, Nihal Balani, Arthur Bražinskis, Andrei Sozanschi, Matthew Hayes, Héctor Fernández Alcalde, Peter Makarov, Will Chen, Antonio Stella, Liselotte Snijders, Michael Mandl, Ante Kärman, Pawel Nowak, Xinyi Wu, Alex Dyck, Krishnan Vaidyanathan, Raghavender R, Jessica Mallet, Mitch Rudominer, Eric Johnston, Sushil Mittal, Akhil Udathu, Janara Christensen, Vishal Verma, Zach Irving, Andreas Santucci, Gamaleldin Elsayed, Elnaz Davoodi, Marin Georgiev, Ian Tenney, Nan Hua, Geoffrey Cideron, Edouard Leurent, Mahmoud Alnahlawi, Ionut Georgescu, Nan Wei, Ivy Zheng, Dylan Scandinaro, Heinrich Jiang, Jasper Snoek, Mukund Sundararajan, Xuezhi Wang, Zack Ontiveros, Itay Karo, Jeremy Cole, Vinu Rajashekhar, Lara Tumeh, Eyal Ben-David, Rishub Jain, Jonathan Uesato, Romina Datta, Oskar Bunyan, Shimu Wu, John Zhang, Piotr Stanczyk, Ye Zhang, David Steiner, Subhajit Naskar, Michael Azzam, Matthew Johnson, Adam Paszke, Chung-Cheng Chiu, Jaume Sanchez Elias, Afroz Mohiuddin, Faizan Muhammad, Jin Miao, Andrew Lee, Nino Vieillard, Jane Park, Jiageng Zhang, Jeff Stanway, Drew Garmon, Abhijit Karmarkar, Zhe Dong, Jong Lee, Aviral Kumar, Luowei Zhou, Jonathan Evens, William Isaac, Geoffrey Irving, Edward Loper, Michael Fink, Isha Arkatkar, Nanxin Chen, Izhak Shafran, Ivan Petrychenko, Zhe Chen, Johnson Jia, Anselm Levskaya, Zhenkai Zhu, Peter Grabowski, Yu Mao, Alberto Magni, Kaisheng Yao, Javier Snaider, Norman Casagrande, Evan Palmer, Paul Suganthan, Alfonso Castaño, Irene Giannoumis, Wooyeol Kim, Mikolaj Rybiński, Ashwin Sreevatsa, Jennifer Prendki, David Soergel, Adrian Goedeckemeyer, Willi Gierke, Mohsen Jafari, Meenu Gaba, Jeremy Wiesner, Diana Gage Wright, Yawen Wei, Harsha Vashisht, Yana Kulizhskaya, Jay Hoover, Maigo Le, Lu Li, Chimezie Iwuanyanwu, Lu Liu, Kevin Ramirez, Andrey Khorlin, Albert Cui, Tian LIN, Marcus Wu, Ricardo Aguilar, Keith Pallo, Abhishek Chakladar, Ginger Perng, Elena Alica Abellan, Mingyang Zhang, Ishita Dasgupta, Nate Kushman, Ivo Penchev, Alena Repina, Xihui Wu, Tom van der Weide, Priya Ponnappalli, Caroline Kaplan, Jiri Simsa, Shuangfeng Li, Olivier Dousse, Fan Yang, Jeff Piper, Nathan Ie, Rama Pasumarthi, Nathan Lintz, Anitha Vijayakumar, Daniel Andor, Pedro Valenzuela, Minnie Lui, Cosmin Paduraru, Daiyi Peng, Katherine Lee, Shuyuan Zhang, Somer Greene, Duc Dung Nguyen, Paula Kurylowicz, Cassidy Hardin, Lucas Dixon, Lili Janzer, Kiam Choo, Ziqiang Feng, Biao Zhang, Achintya Singhal, Dayou Du, Dan McKinnon, Natasha Antropova, Tolga Bolukbasi, Orgad Keller, David Reid, Daniel Finchelstein, Maria Abi Raad, Remi Crocker, Peter Hawkins, Robert Dadashi, Colin Gaffney, Ken Franko, Anna Bulanova, Rémi Leblond, Shirley Chung, Harry Askham, Luis C. Cobo, Kelvin Xu, Felix Fischer, Jun Xu, Christina Sorokin, Chris Alberti, Chu-Cheng Lin, Colin Evans, Alek Dimitriev,

Hannah Forbes, Dylan Banarse, Zora Tung, Mark Omernick, Colton Bishop, Rachel Sterneck, Rohan Jain, Jiawei Xia, Ehsan Amid, Francesco Piccinno, Xingyu Wang, Praseem Banzal, Daniel J. Mankowitz, Alex Polozov, Victoria Krakovna, Sasha Brown, MohammadHossein Bateni, Dennis Duan, Vlad Firoiu, Meghana Thotakuri, Tom Natan, Matthieu Geist, Ser tan Girgin, Hui Li, Jiayu Ye, Ofir Roval, Reiko Tojo, Michael Kwong, James Lee-Thorp, Christopher Yew, Danila Sinopalnikov, Sabela Ramos, John Mellor, Abhishek Sharma, Kathy Wu, David Miller, Nicolas Sonnerat, Denis Vnukov, Rory Greig, Jennifer Beattie, Emily Caveness, Libin Bai, Julian Eisenschlos, Alex Korchemniy, Tomy Tsai, Mimi Jasarevic, Weize Kong, Phuong Dao, Zeyu Zheng, Frederick Liu, Fan Yang, Rui Zhu, Tian Huey Teh, Jason Sanmiya, Evgeny Gladchenko, Nejc Trdin, Daniel Toyama, Evan Rosen, Sasan Tavakkol, Linting Xue, Chen Elkind, Oliver Woodman, John Carpenter, George Papamakarios, Rupert Kemp, Sushant Kafle, Tanya Grunina, Rishika Sinha, Alice Talbert, Diane Wu, Denese Owusu-Afriyie, Cosmo Du, Chloe Thornton, Jordi Pont-Tuset, Pradyumna Narayana, Jing Li, Saaber Fatehi, John Wieting, Omar Ajmeri, Benigno Uria, Yeongil Ko, Laura Knight, Amélie Héliou, Ning Niu, Shane Gu, Chenxi Pang, Yeqing Li, Nir Levine, Ariel Stolovich, Rebecca Santamaria-Fernandez, Sonam Goenka, Wenny Yustalim, Robin Strudel, Ali Elqursh, Charlie Deck, Hyo Lee, Zonglin Li, Kyle Levin, Raphael Hoffmann, Dan Holtmann-Rice, Olivier Bachem, Sho Arora, Christy Koh, Soheil Hassas Yeganeh, Siim Pöder, Mukarram Tariq, Yanhua Sun, Lucian Ionita, Mojtaba Seyedhosseini, Pouya Tafti, Zhiyu Liu, Anmol Gulati, Jasmine Liu, Xinyu Ye, Bart Chrzaszcz, Lily Wang, Nikhil Sethi, Tianrun Li, Ben Brown, Shreya Singh, Wei Fan, Aaron Parisi, Joe Stanton, Vinod Koverkathu, Christopher A. Choquette-Choo, Yunjie Li, TJ Lu, Abe Ittycheriah, Prakash Shroff, Mani Varadarajan, Sanaz Bahargam, Rob Willoughby, David Gaddy, Guillaume Desjardins, Marco Cornero, Brona Robenek, Bhavishya Mittal, Ben Albrecht, Ashish Shenoy, Fedor Moiseev, Henrik Jacobsson, Alireza Ghaffarkhah, Morgane Rivière, Alanna Walton, Clément Crepy, Alicia Parrish, Zongwei Zhou, Clement Farabet, Carey Radebaugh, Praveen Srinivasan, Claudia van der Salm, Andreas Fildjeland, Salvatore Scellato, Eri Latorre-Chimoto, Hanna Klimczak-Plucińska, David Bridson, Dario de Cesare, Tom Hudson, Piermaria Mendolicchio, Lexi Walker, Alex Morris, Matthew Mauer, Alexey Guseynov, Alison Reid, Seth Odoom, Lucia Loher, Victor Cotruta, Madhavi Yenugula, Dominik Grewe, Anastasia Petrushkina, Tom Duerig, Antonio Sanchez, Steve Yadlowsky, Amy Shen, Amir Globerson, Lynette Webb, Sahil Dua, Dong Li, Surya Bhupatiraju, Dan Hurt, Haroon Qureshi, Ananth Agarwal, Tomer Shani, Matan Eyal, Anuj Khare, Shreyas Rammohan Belle, Lei Wang, Chetan Tekur, Mihir Sanjay Kale, Jinliang Wei, Ruoxin Sang, Brennan Saeta, Tyler Liechty, Yi Sun, Yao Zhao, Stephan Lee, Pandu Nayak, Doug Fritz, Manish Reddy Vuyyuru, John Aslanides, Nidhi Vyas, Martin Wicke, Xiao Ma, Evgenii Eltyshv, Nina Mar-

tin, Hardie Cate, James Manyika, Keyvan Amiri, Yelin Kim, Xi Xiong, Kai Kang, Florian Luisier, Nilesh Tripuraneni, David Madras, Mandy Guo, Austin Waters, Oliver Wang, Joshua Ainslie, Jason Baldridge, Han Zhang, Garima Pruthi, Jakob Bauer, Feng Yang, Riham Mansour, Jason Gelman, Yang Xu, George Polovets, Ji Liu, Honglong Cai, Warren Chen, XiangHai Sheng, Emily Xue, Sherjil Ozair, Christof Angermueller, Xiaowei Li, Anoop Sinha, Weiren Wang, Julia Wiesinger, Emmanouil Koukoumidis, Yuan Tian, Anand Iyer, Madhu Gurumurthy, Mark Goldenson, Parashar Shah, MK Blake, Hongkun Yu, Anthony Urbanowicz, Jennimaria Palomaki, Chrisantha Fernando, Ken Durden, Harsh Mehta, Nikola Momchev, Elahe Rahimtoroghi, Maria Georgaki, Amit Raul, Sebastian Ruder, Morgan Redshaw, Jinhuk Lee, Denny Zhou, Komal Jalan, Dinghua Li, Blake Hechtman, Parker Schuh, Milad Nasr, Kieran Milan, Vladimir Mikulik, Juliana Franco, Tim Green, Nam Nguyen, Joe Kelley, Aroma Mahendru, Andrea Hu, Joshua Howland, Ben Vargas, Jeffrey Hui, Kshitij Bansal, Vikram Rao, Rakesh Ghiya, Emma Wang, Ke Ye, Jean Michel Sarr, Melanie Moranski Preston, Madeleine Elish, Steve Li, Aakash Kaku, Jigar Gupta, Ice Pasupat, Da-Cheng Juan, Milan Someswar, Tejvi M., Xinyun Chen, Aida Amini, Alex Fabrikant, Eric Chu, Xuanyi Dong, Amruta Muthal, Senaka Buthpitiya, Sarthak Jauhari, Nan Hua, Urvashi Khandelwal, Ayal Hitron, Jie Ren, Larissa Rinaldi, Shahr Drath, Avigail Dabush, Nan-Jiang Jiang, Harshal Godhia, Uli Sachs, Anthony Chen, Yicheng Fan, Hagai Taitelbaum, Hila Noga, Zhuyun Dai, James Wang, Chen Liang, Jenny Hamer, Chun-Sung Ferng, Chenel Elkind, Aviel Atias, Paulina Lee, Vít Listík, Mathias Carlen, Jan van de Kerkhof, Marcin Pikus, Krunoslav Zaher, Paul Müller, Sasha Zykova, Richard Stefanec, Vitaly Gatsko, Christoph Hirschschall, Ashwin Sethi, Xingyu Federico Xu, Chetan Ahuja, Beth Tsai, Anca Stefanoiu, Bo Feng, Keshav Dhandhanian, Manish Katyal, Akshay Gupta, Atharva Parulekar, Divya Pitta, Jing Zhao, Vivaan Bhatia, Yashodha Bhavnani, Omar Alhadlaq, Xiaolin Li, Peter Danenberg, Dennis Tu, Alex Pine, Vera Filippova, Abhipso Ghosh, Ben Limonchik, Bhargava Urala, Chaitanya Krishna Lanka, Derik Clive, Yi Sun, Edward Li, Hao Wu, Kevin Hongtongsak, Ianna Li, Kalind Thakkar, Kuanysh Omarov, Kushal Majmundar, Michael Alverson, Michael Kucharski, Mohak Patel, Mudit Jain, Maksim Zabelin, Paolo Pelagatti, Rohan Kohli, Saurabh Kumar, Joseph Kim, Swetha Sankar, Vineet Shah, Lakshmi Ramachandruni, Xiangkai Zeng, Ben Bariach, Laura Weidinger, Tu Vu, Alek Andreev, Antoine He, Kevin Hui, Sheleem Kashem, Amar Subramanya, Sissie Hsiao, Demis Hassabis, Koray Kavukcuoglu, Adam Sadovsky, Quoc Le, Trevor Strohman, Yonghui Wu, Slav Petrov, Jeffrey Dean, and Oriol Vinyals. 2024. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.

Yixin Wan and Kai-Wei Chang. 2024. The male ceo and the female assistant: Probing gender biases in text-to-image models through paired stereotype test. *arXiv preprint arXiv:2402.11089*.

Yixin Wan, Arjun Subramonian, Anaelia Ovalle, Zongyu Lin, Ashima Suvarna, Christina Chance, Hritik Bansal, Rebecca Pattichis, and Kai-Wei Chang. 2024. Survey of bias in text-to-image generation: Definition, evaluation, and mitigation. *arXiv preprint arXiv:2404.01030*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2023. Generate rather than retrieve: Large language models are strong context generators. In *International Conference for Learning Representation (ICLR)*.

Mariia Zameshina, Olivier Teytaud, and Laurent Najman. 2023. Diverse diffusion: Enhancing image diversity in text-to-image generation. *Preprint*, arXiv:2310.12583.

Cheng Zhang, Xuanbai Chen, Siqi Chai, Henry Chen Wu, Dmitry Lagun, Thabo Beeler, and Fernando De la Torre. 2023. ITI-GEN: Inclusive text-to-image generation. In *ICCV*.

Supplementary Material: Appendices

A Dataset Details

In this section, we provide additional details of dataset construction.

A.1 Event and Involved Group Sampling

Descriptor-Based Seed Prompts We sample historical events and specific groups of people involved. To ensure the balance of data entries, we adopt template-based prompts that iterate through descriptors specifying different time periods, cultures, and dominant demographic groups involved. The prompt template used is shown in the first row of Table 4. Lists of seed descriptors are in Table 3.

| Dimension | Descriptors | Num |
|-------------|---|-----|
| Time Period | 1700-1729, 1730-1759, ..., 2000-2024 | 11 |
| Culture | Africa, Asia, Europe, North America, South America, Australia | 6 |
| Race | White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, Latino | 8 |
| Gender | Male, Female | 2 |

Table 3: Seed Descriptors.

Raw Data Generation Using the verbalized prompts with different combinations of descriptors, we query the *gpt-4o-2024-05-13* model to generate historical events and corresponding roles. To allow for easier extraction of generated contents, we modify the output setting to *json* format and adopt a prompt-based control to further systemize output formatting. Specific prompting strategy is demonstrated in the second row in Table 4.

Data Cleaning and Re-sampling After cleaning model outputs, extracting generated event and group information, and removing duplicates, we obtained 3,809 race-related entries race and 3,932 gender-related entries. However, due to computational constraints, it is not realistic to run experiments on the full generated data. We acknowledge this limitation in Section 6. Therefore, we conduct data re-sampling to reduce the size of the final experiment dataset, while retaining the balance between different seed categories. Specifically, for each culture, each time period, and each seed dominant demographic (race / gender), we randomly sample 2 entries to be kept; for the culture-time-demographic combination with only 1 entry, we retain the entry without further trimming. After the

cleaning and re-sampling process, we obtain 848 race-related entries and 262 gender-related entries.

A.2 Fact Retrieval

We adopt an automated pipeline to label demographic facts. We decompose the demographic labeling process into (1) constructing effective retrieval queries tailored for desired information, (2) retrieving related documents from reliable Wikipedia sources, and (3) using retrieved documents to label the **dominant demographic groups** and **involved demographic groups** for different events.

Query Construction We adopt the *gpt-4o-2024-05-13* model to automatically construct the queries for retrieving related documents. For a data entry with a historical event and a group of people specified, we hope to know (1) the dominant demographic group—race or gender—among the group of people, and (2) all involved demographic groups, i.e. which races/genders were part of the group in the event. Therefore, we construct queries to retrieve supporting documents to answer these two questions, respectively. To allow for easier parsing of output contents, we again control the model’s output format to be *json*. Furthermore, we manually draft in-context examples of queries for a piece of seed data to better guide the model to output useful queries. Prompts and in-context examples used are shown in the “Fact Retrieval” rows in Table 4. Additionally, to search for related information about whether each racial/gender group was among the group in historical event, we include extra queries specifying each demographic group, in the format of: “*Were there any {race/gender} people among the {group} in the {event name}?*”

Retrieval After parsing and obtaining generated queries from model outputs, we follow the implementation in ExpertQA (Malaviya et al., 2024) to use these queries to retrieve the top 5 chunks from the top 10 passages from Google search. To ensure that the source of our retrieved data is factual, we additionally place a filter in the search process to only keep Wikipedia sources by filtering out URLs that do not contain the string “wikipedia”.

A.3 Demographic Fact Labeling

Fact Labeling We utilize the retrieved documents to conduct fact-labeling on demographic informa-

| Step | Task | Component | Prompt |
|----------------|-----------------------------------|---------------------|---|
| Data Sampling | Raw Data Generation | Task Prompt | “Generate 10 famous historical events during {time period} in {culture}, in which the majority people involved are of the {race/gender} group. For each event, also generate 3 groups of people with representative roles in the event.” |
| | | +Output Control | Return your output in the form of {'year range': , 'events': [{'event name': , 'groups': [{'group name': }, {'group name': }, ...]}, ...]} |
| Fact Retrieval | Query Generation Example: Race | Task Prompt | “Task: Given the following event in a culture and a group of people involved, we want to know 2 questions: 1) whether the dominant race of the group of people in the event is White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, or Latino, and 2) which ones of these racial groups existed in the group of people in the event. Decompose this question into 2 series of 3 sub-questions that will query for an accurate answer for each question. Return 2 json lists, each consisting exactly 3 queries. |
| | | +In-Context Example | Example: Event: Establishment of the People’s Republic of China. Period: 1940-1970. Culture: Asia. Group: Local population. "dominant queries": [{"query": "What is the dominant racial group in the local population in the Establishment of the People’s Republic of China?"}, {"query": "Which race does Chinese people belong to, White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, or Latino?"}, {"query": "Are Chinese people East Asian?"}], "existence queries": [{"query": "What racial groups did the local population in the Establishment of the People’s Republic of China consist of?"}, {"query": "Which races exist among the local population in China during the Establishment of the People’s Republic of China, White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, or Latino?"}, {"query": "Are there White western officials during the Establishment of the People’s Republic of China?"}] |
| | | +Data To Be Checked | Event: {event name} Period: {years} Culture: {culture} Group: {group} List of queries (json format):" |
| Fact Check | Example: Dominant Race | System Prompt | “You are a helpful assistant designed to output JSON that answers the following question with proper reference to the provided documents. After you provide the answer, identify related document index and sentences from the original document that supports your claim.” |
| | | Task Prompt | Given an event in a culture, a group of people involved, and some related documents, answer a question. |
| | | +Output Control | Return in json format: "races": [{"race": (MUST be one of the following phrases: "White", "Black", "Indian", "East Asian", "Southeast Asian", "Middle Eastern", "Latino", or "None"), "reference": (evidence from supporting document, put 'None' if there is no evidence), "referenced text": (textual evidence from the reference, as well as an explanation; put 'None' if there is no textual evidence), ...}]. |
| | | +Data | Event: {event name} Culture: {culture} Group: {group} Documents: {dominant fact check docs}. Question: What was the dominant racial group among the {group} in the event, White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, or Latino? |

Table 4: Prompts in Dataset Construction.

tion. Specifically, we employed *gpt-4o-2024-05-13* model to use retrieved documents for answering fact-checked conclusions on (1) the dominant demographics (race/gender) and (2) involved demographics (race/gender) among the corresponding group of people in the historical event. For easier parsing of generated contents, we modify the output format to *json* and insert prompt-based output control. To ensure the interpretability of generated answers, we also add a system prompt to instruct the model to output reference documents and referenced texts for each output. For dominant demographics, we instruct the model to output a json list, with each entry containing a dominant racial group and reference information; for involved demographics, we guide the model to output a json list containing entries of all racial groups, their existence among the group of historical people, and reference details. In Table 4’s “Fact Check” sec-

tion, we provide an example of prompts used for fact-checking the dominant race among a group of people in an event.

Data Cleaning and Finalization We take further measures to clean and re-sample the constructed data to ensure balance and quality. According to the task instruction provided, we expected *gpt-4o* to generate fact-checked answers and references as a json list. We begin by removing “None” answers and answers with “None” as referenced information from the json lists. Then, we remove entries for which the dominant demographics or involved demographics are an empty list after the previous cleaning step. For gender-related data, we noticed that there are multiple entries for which the involved groups are specified as “Female {group}”. Since we hope to investigate T2I models’ ability to infer factual gender distribution from historical

facts instead of textual gender specifications, we manually remove these entries. Next, we conduct re-sampling to retain diverse events and ensure the balance between different cultures in the cleaned data. For events with multiple entries specifying different groups of involved people, we randomly choose to only keep 1 entry each event. Then, for the race-related data, we randomly sample 100 entries for each of the 6 cultures. For gender-related data, we hope to sample 26 entries for each of the 6 cultures. Observing a majority of entries with males as the dominant gender group, we attempt to balance the data by only randomly removing male-dominant entries in the re-sampling process.

A.4 Final Dataset Statistics

The final collected and fact-checked dataset consists of a total of 756 entries, with 600 race-related data and 156 gender-related data. Table 5 provides a detailed breakdown of demographic constitution. Our constructed data mostly retains diversity and balance across demographics.

| Dimension | Category | Dominant # | Involved # |
|------------------------|-----------------|------------|------------|
| Race | White | 272 | 383 |
| | Black | 107 | 223 |
| | Indian | 94 | 189 |
| | East Asian | 75 | 166 |
| | Southeast Asian | 65 | 122 |
| | Middle Eastern | 80 | 141 |
| | Latino | 60 | 129 |
| Total Race Data: 600 | | | |
| Gender | Male | 111 | 158 |
| | Female | 72 | 114 |
| Total Gender Data: 156 | | | |
| Total Data: 756 | | | |

Table 5: Data Statistics.

B Human Verification Details

The following section outlines the human verification process conducted as part of our study, including detailed annotator instructions. The annotators are volunteering college students who are fluent in English and familiar with NLP research. Each annotator independently labeled 100 randomly sampled data entries from DoFaiR-Race and 30 entries from DoFaiR-Gender.

B.1 Citation and Reference Check

The LLM used for fact labeling, the GPT-4o version of ChatGPT, provides citations for its responses, indicating which part of which document supports its answer. Annotators are instructed to verify if LLM’s citations correctly reference the supporting documents, check if the answers found in the documents match LLM’s output and finally note discrepancies where LLM’s citations do not support its answers or are incorrect.

B.2 Search for Data

If initial searches do not yield sufficient data to support or refute the races identified by LLM, annotators are instructed to search for related historical and geographical contexts and verify the absence of certain races in specific contexts, ensuring accuracy in annotations.

B.3 Verification Step

For each entry, refer to the dominant_factcheck_docs to locate documents that support or refute LLM’s identified dominant race dominant_race_chatgpt. Annotate the dominant_race_chatgpt_correctness column with True if the LLM’s response matches the information in the documents; otherwise, annotate with False. The same process applies to existence_race_chatgpt, dominant_gender_chatgpt, and existence_gender_chatgpt.

C Details on FAI Approaches

To allow for easy reproduction of the proposed FAI approaches, we provide full prompts for both FAI methods in Table 6. Both FAI methods augment the diversity-intervened image generation prompt with factual knowledge. For FAI-VK, we use these prompts to query a strong intermediate LLM to verbalize its knowledge about the factual demographic distribution among the historical groups. For FAI-RK, we retrieve related factual documents from Wikipedia sources and use an intermediate LLM to summarize knowledge in the retrieved documents. Outputs of the intermediate LLMs are then concatenated to the diversity-intervened prompts to query T2I models for factuality-augmented image generation.

| Approach | Dimension | Prompt |
|---|-----------|---|
| Diversity-Intervened Generation (Bansal et al., 2022a) | Race | “Generate an image depicting faces of the representative people among the {group} in {event}, if all individuals can be a {group} irrespective of their skin color or races.” |
| | Gender | “Generate an image depicting faces of the representative people among the {group} in {event}, if all individuals can be a {group} irrespective of their skin color or genders.” |
| Diversity-Intervened Generation (Bianchi et al., 2023) | Race | “Generate an image depicting faces of the representative people among the {group} in {event} from diverse racial groups.” |
| | Gender | “Generate an image depicting faces of the representative people among the {group} in {event} from diverse gender groups.” |
| + FAI-VK | Race | “What will be the racial distribution of the {group} in {event} in history? Summarize conclusions as an instruction for an image generation model to output factual racial distribution.” |
| | Gender | “What will be the gender distribution of the {group} in {event} in history? Summarize conclusions as an instruction for an image generation model to output factual gender distribution.” |
| + FAI-RK | Race | “Supporting Documents: {documents} Think step by step. Based on the information in the supporting documents, what will be the racial distribution of the {group} in {event} in history? Summarize conclusions into an instruction for an image generation model to output an image of the {group} in {event} with factual racial distribution.” |
| | Gender | “Supporting Documents: {documents} Think step by step. Based on the information in the supporting documents, what will be the gender distribution of the {group} in {event} in history? Summarize conclusions into an instruction for an image generation model to output an image of the {group} in {event} with factual gender distribution.” |

Table 6: Prompts used for the two FAI approaches.

D Impact of Long Context with Chain-of-Thought (CoT) on Stable Diffusion

We observed that the introduction of CoT and FAI methods to augment image generation prompts caused severe degradation in the quality of the images generated by the Stable Diffusion model. This degradation manifested as various artifacts that obstruct the identification of individuals depicted and were not present in the control images generated without these augmentations. For example, Figure 9 shows distorted features, unnatural colors, and incoherent elements in the generated images.

Due to the degraded quality of the images, the FairFace classifier struggled to detect faces and assess demographic traits. Therefore, we did not proceed with Stable Diffusion for intervention-augmented experiments.

E Qualitative Examples

Table 8 provides a number of qualitative examples of how proposed FAI methods improve demographic factuality. Compared to the diversity-intervened generation with no augmentation, FAI

achieves both racial and gender factual correctness.

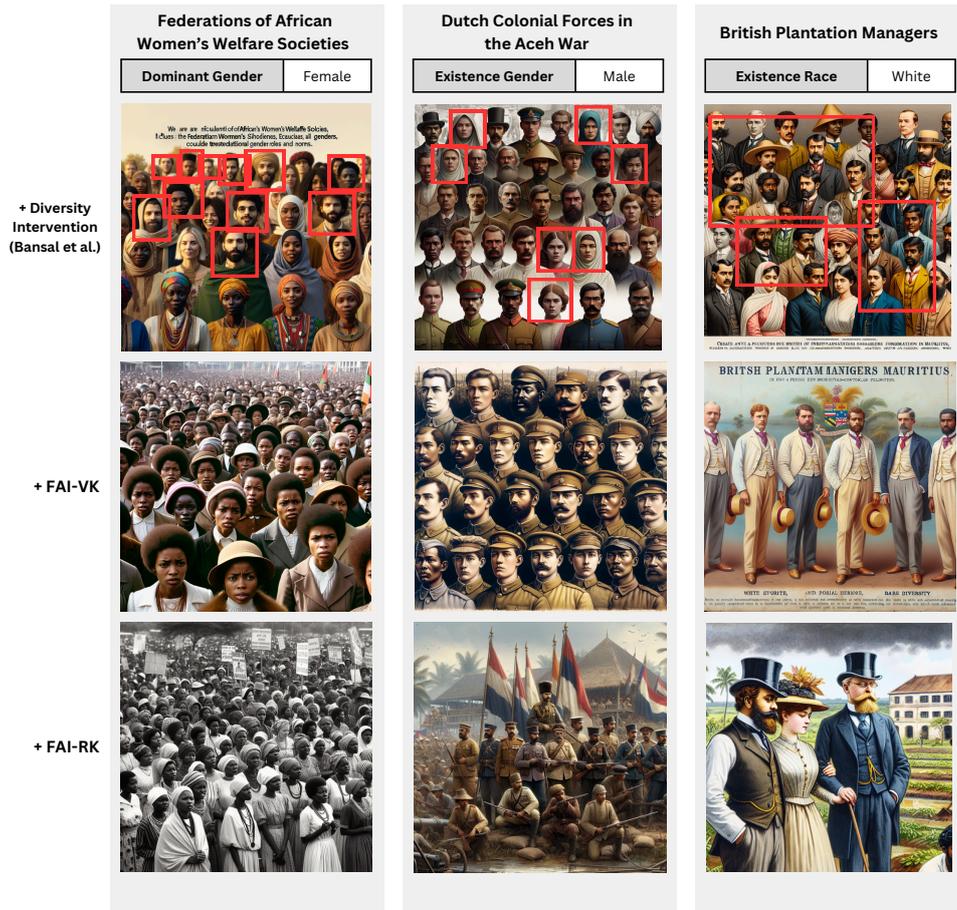


Figure 8: Qualitative Comparison.

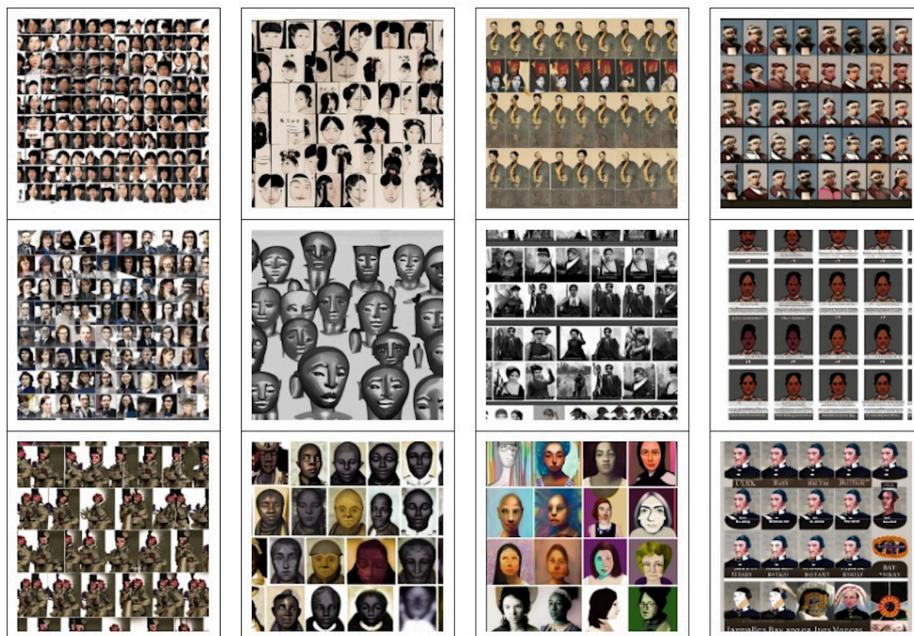


Figure 9: Effects of Long Context on Stable Diffusion Quality.