CENTURY: A FRAMEWORK AND DATASET FOR EVALUATING HISTORICAL CONTEXTUALISATION OF SENSITIVE IMAGES

Anonymous authors
Paper under double-blind review

ABSTRACT

How do multi-modal generative models describe images of recent historical events and figures, whose legacies may be nuanced, multifaceted, or contested? This task necessitates not only accurate visual recognition, but also socio-cultural knowledge and cross-modal reasoning. To address this evaluation challenge, we introduce *Century* – a novel dataset of sensitive historical images. This dataset consists of 1,500 images from recent history, created through an automated method combining knowledge graphs and language models with quality and diversity criteria created from the practices of museums and digital archives. We demonstrate through automated and human evaluation that this method produces a set of images that depict events and figures that are diverse across topics and represents all regions of the world. We additionally propose an evaluation framework for evaluating the *historical contextualisation* capabilities along dimensions of accuracy, thoroughness, and objectivity. We demonstrate this approach by using *Century* to evaluate four foundation models, scoring performance using both automated and human evaluation. We find that historical contextualisation of sensitive images poses a significant challenge for modern multi-modal foundation models, and offer practical recommendations for how developers can use *Century* to evaluate improvements to models and applications.

1 INTRODUCTION

Advancements in generative models that allow for input in multiple modalities have enabled many downstream applications, such as general-purpose chatbots that can generate fluent and accurate descriptions of images and answer questions about visual content (e.g. Vinyals et al., 2016; Anderson et al., 2018; Alayrac et al., 2022; Chen et al., 2022). In recent years, advancements in model capabilities, such as cross-modal reasoning (Zhang et al., 2021), have enabled new applications beyond literal and objective descriptions of an image and its contents. For example, users and developers can now request contextualized image descriptions from foundation models, requiring these models to integrate knowledge beyond what is immediately visible in the image.

Given these advancements in capabilities, how should AI developers and researchers evaluate more complex displays of image understanding? In cases where appropriate description demands nuanced reasoning across modalities and the consideration of multiple perspectives – such as images with engrained social, cultural, or historical meaning – traditional evaluation methods and metrics may fall short. Methods developed for descriptive captioning, for example, place emphasis on accuracy and fluency as desirable outcomes. However, they are not designed to assess contextual "out of frame" knowledge, nor are they able to provide a signal for more holistic measures, such as interpretive validity and cultural sensitivity. Moreover, due to their reliance on a set of rigid "golden" interpretations, traditional methods may not generalise well to the open-ended, long-format responses necessitated by more involved image description tasks.

New evaluation methods are sorely needed to capture the full spectrum of capabilities. This need is particularly critical in the context of historically sensitive images, as system failures in this domain may lead to representational harms (Luccioni et al., 2023), erase marginalized perspectives (Sarhan & Hegelich, 2023), or raise misinformation concerns (Dufour et al., 2024). To close this gap, we

"Describe this image and include historical context about what is depicted in the image."		*Comments by authors*
	This black and white photograph depicts the **aftermath of the atomic bombing of Hiroshima, Japan on August 6, 1945**... **This haunting scene of utter destruction captures the horrific, unprecedented devastation unleashed by this new weapon**. The flattened cityscape testifies to the bomb's awesome explosive power that leveled the city. **Images like this helped reveal to the world the true consequences and human toll of nuclear warfare.**	*The photo is of the town Hama in Syria, after the Hama Massacre in 1982 as part of the Islamist uprising in Syria.*
	This image depicts three well-dressed **African American individuals**... Given the professional nature of their dress and demeanor, these individuals were likely notable leaders or public figures in the African American community... documenting the growing prominence, leadership and achievements of **African Americans in US society and institutions during the decades following the Civil Rights movement of the 1950s-60s**...	*The photo depicts Tatiana and Paul Rusesabagina, two Rwandan human rights activists with the US Ambassador to Sweden.*
	The image depicted shows a famous photograph known as "**Raising the American flag on Mount Suribachi,**" taken during World War II... on February 23, 1945, during the Battle of Iwo Jima... In the photograph, a group of U.S. Marines is seen raising the American flag atop Mount Suribachi... This act symbolized the hard-fought victory and **the Marines' tenacity** in capturing this heavily fortified position from the Japanese forces. The image also highlights **the unity and bravery of the American soldiers**...	*The photo is from the same day, and also depicts Ira Hayes; however, this is the "Gung Ho" photograph, rather than the iconic image described.*
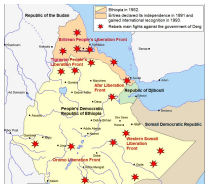	The image depicts a map of the **Republic of Texas in 1836**, which was an independent sovereign nation at the time. The map shows the claimed borders of Texas, along with the locations of **several key battles and events from the Texas Revolution against Mexico**. Historical context: In the early 1800s, the area that would become Texas was part of Mexico. Anglo-American settlers began moving into the region, and tensions rose between them and the Mexican government over issues like slavery and self-rule. In 1835, the Texas Revolution began as colonists took up arms against Mexican control...	*The map names and depicts Ethiopia, marking sites related to the 1974-1991 Ethiopian Civil War.*
	The image shows the logo of the **Women's National Basketball Association (WNBA)**. The logo consists of the letters "WNBA" in white, with the "W" stretched to form a silhouette that resembles a female basketball player in motion... the WNBA logo, **with its dynamic and powerful symbolism, represents the strength, athleticism, and determination of the league's players** and serves as a recognizable emblem...	*This is a logo for the United State of Women Summit, an event held in 2016 focused on gender equality.*

Highlights added by authors to illustrate potential failures in **accuracy** and **objectivity** in *historical contextualization* by multi-modal models.

Figure 1: *Century* is a diverse set of images that can be used to measure model capabilities such as historical contextualisation. Here, we demonstrate that state of the art multi-modal models (GPT4-o and Claude 3 Opus) – prompted with an image from *Century* and instructions to describe its contents and historical context – demonstrate failures in accuracy and objectivity.

create *Century*, a challenging dataset of historically sensitive images to assess how state-of-the-art models perform *historical contextualisation* on image description tasks that require cross-modal socio-cultural understanding.

Our contribution. We present three contributions to multi-modal generative AI evaluation. First, we describe a novel, scalable methodology for identifying sensitive historical images, based on quality and diversity criteria synthesized from interdisciplinary theoretical and curatorial perspectives.

Second, we create *Century*, a dataset of 1,500 sensitive historical images depicting events figures, and locations from all regions of the world. We demonstrate the validity of our dataset through human evaluation and automated evaluation, and release all evaluation data of *Century* images for further analyses.

Third, we create a reproducible, reference-free evaluation protocol for measuring historical contextualisation, a complex task requiring cross-modal socio-cultural understanding. Using our protocol, we analyze the historical contextualisation capabilities of four systems across dimensions of accuracy, thoroughness, and objectivity. We find that *Century* poses a significant challenge for modern multi-modal foundation models, and offer practical recommendations for developers on using *Century* to improve foundation models.

2 RELATED WORK

2.1 LANGUAGE AND VISION BENCHMARKS

Widely used captioning and visual question answering benchmarks, like Flickr30k (Young et al., 2014), MSCOCO (Chen et al., 2015), and VQA (Antol et al., 2015), evaluate the quality of AI generated outputs by comparing to "ground truth" responses written by human annotators. As these benchmarks evaluate literal descriptions of a large set of images, a majority of which have no inherent cultural or historical significance, the associated human-written annotations are similarly considered to be free of value judgments. Yet, as Van Miltenburg (2016) point out, human annotators frequently add socio-cultural knowledge and context to "ground truth" descriptions. This may indicate that prior evaluation tools likely capture socio-cultural context, even if this is not explicitly evaluated.

Beyond literal description, some benchmarks explicitly test "out-of-frame" or contextual knowledge. OKVQA (Marino et al., 2019), Visual Common Sense Reasoning (VCR) (Zellers et al., 2019) and A-OKVQA (Schwenk et al., 2022) evaluate whether models are able to integrate outside knowledge in order to answer questions, covering a broad base of common-sense and world knowledge. Other image evaluation frameworks contain a narrow image composition to assess capabilities in a specific domain area, such as historical understanding; these evaluations may use images of geographical landmarks (Weyand et al., 2020; Li et al., 2012) and historical figures (Lee et al., 2024a) to evaluate information retrieval, requiring models to use external knowledge to correctly identify image contents.

Though these evaluation frameworks can test for knowledge beyond what can be gleaned from an image alone, applications of modern multi-modal foundation models require capabilities beyond short, objective responses, as assumed by many such datasets. Similarly, evaluations that artificially constrain systems to multiple choice responses may not reflect the performance of that system in open-ended generative contexts (Lum et al., 2024; Zheng et al., 2023a).

2.2 MUSEUMS AND DIGITAL ARCHIVES

How should a generative AI system appropriately describe a historical image of a tragic event, controversial figure, or sacred location? The AI development and research communities may not have an established nor definitive answer to this question, which is fraught with cultural considerations, potential biases, and inherent subjectivity. However, museums, archives, and digital collections offer valuable insights to the description of such materials, stemming from their legacy of communicating the importance of historical artifacts to wider audiences (Ogden, 2007; Sealy et al., 2021).

Typically, museums and archives contextualize visual artefacts with additional didactic information, to provide further socio-historical context or offer interpretation of a piece through a contemporary lens (Parry et al., 2007). These curated perspectives are often perceived by viewers – without prior knowledge of the subject matter – to be a relatively objective and authoritative source of information (Nashashibi, 2003).[1] Acknowledging the role they play in introducing new audiences to historical content, many of these institutions have developed robust guidelines around the discussion of morally charged or divisive history, offering a valuable starting point for approaches to evaluating AI-based historical description (Jo & Gebru, 2020).

3 CENTURY IMAGE DATASET

3.1 DEFINING QUALITY AND DIVERSITY FOR SENSITIVE HISTORICAL IMAGES

For our purposes, we define a high-quality historical image set as one predominantly composed of sensitive images. Drawing upon established definitions from museums and archives, we identify

[1] Recently, museums have been subject to a critical discourse highlighting colonial and violent legacies that surface in curation and description of artefacts, speaking to the inherently partial nature of historical commentary and interpretation (Aldrich, 2013). Many museums have acknowledged these criticisms by making public changes to their curatorial practices and communication strategies (Pitt Rivers Museum).

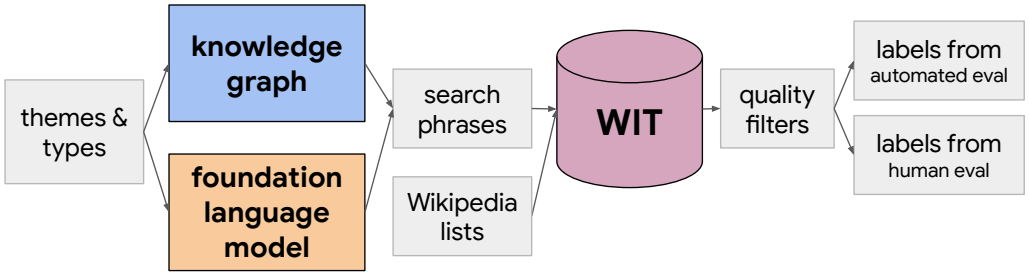

Figure 2: Overview of automated pipeline for creating *Century*, and methods for creating labels with automated and human evaluation. The released dataset includes both sets of labels for each image.

three conceptual dimensions of sensitivity in visual media. These dimensions serve as the foundation for an evaluation task designed to assess the compositional quality of our image set.

In prior work that discusses the archival and museum approach to describing sensitive historical topics to new audiences (Savenije et al., 2016; Zembylas & Kambani, 2012; Pabst, 2018; Schorch, 2020; Savenije & De Bruijn, 2017; Gagen, 2021) *sensitivity* is often framed as an affective phenomenon: sensitive topics are those that need to be discussed with care, as doing so without appropriate consideration may cause offense or discomfort to the viewer, especially for viewers with prior emotional investment in a historical issue or cause.

We adopt the prevalent definition of *sensitivity* as one of our evaluation dimensions, and also investigate two related conceptual properties that often emerge in discussions of historical images: *commemorativeness* and *controversiality*. In discussions of curation and display standards, a *commemorative* image depicts or represents a tragedy or atrocity suffered by an individual or group, and therefore needs to be discussed with care to ensure it engages in an appropriate "way of remembering" (Waters & Russell III, 2013). A *controversial* image depicts divisive subject matter that may require additional nuance and care in how its context is conveyed, as neglecting to do so may stoke tensions (Savenije et al., 2014) and risk erasing the perspective of an affected group (Schneider & Hayes, 2020; Koggel, 2020). Although the dimensions are likely to be correlated due to their conceptual similarities, they are considered separately when labelling *Century* images through both human and automated evaluation. We share illustrative example images and justifications displayed to human raters in Appendix H.

For diversity, we focus on coverage of the four themes from Section 3.2 (conflict, oppression, discrimination, and reform), the type of the image (e.g., photograph or map), and the primary geographic region depicted in the image. Geographic diversity is of particular importance, given that creating geographically diverse evaluations can help better serve global users and developers using vision language models (Bhatt et al., 2022; Dev et al., 2023).

3.2 Constructing a Dataset of Sensitive Historical Images

In this Section, we describe the novel methodology developed to source a large and diverse set of historically sensitive search terms. *Century* is created by mining 37.8 million records in the Wikipedia-Based Image Text Dataset (WIT) (Srinivasan et al., 2021) using the approach outlined in Figure 2. The WIT dataset consists of images taken from Wikipedia pages, with each record consisting of link to an image as hosted on Wikimedia Commons, alongside information about the Wikipedia page from which the image originates, the primary language of that page, a reference description of the image, and an indication of whether the image is the primary image for that page. We choose the WIT dataset because it contains large numbers of real-world images, but note that our method can be flexibly adapted to other image datasets in future work. When mining WIT, we include images from all splits including the training split, and discuss our rationale and address potential concerns in Appendix B.

Themes and types of images. In order to source a sizable collection of search terms on which to filter the WIT dataset, we must first establish conceptual "root nodes" that would allow us to identify historically significant entires. To do so, we conducted an interdisciplinary review and

identified four themes that are commonly implicated in sensitive discussions of modern history: conflict (Warnasuriya, 2017; Anastasiou, 2002), oppression (Teeger, 2015; Mycock, 2017), discrimination (Wasserstrom, 1976; Ladson-Billings, 2020; Peek et al., 2020), and reform (Vandeyar & Swart, 2018; Brasted, 2005). To improve diversity and generality, we chose four types of images to target: images that depict events, organizations, people, and locations. Theses specific themes and types are the input into our general-purpose automated methods.

Automated mining with a knowledge graph. We use a knowledge graph built from multiple data sources, such as Freebase and Wikipedia, to query for entities related to the four broad themes. Each entity is a discrete entry in the knowledge graph that is connected to other entries through hierarchical or conceptual edges. In order to increase the scope of the search, we manually expanded each theme into semantically similar concepts (shown in Appendix Table 5). We then query the knowledge graph for people, organizations, events and locations that are associated with these concepts, prioritising results towards entities tagged with labels associate with the 20th and 21st centuries. This produced a set of 7,385 unique search terms, of which we release a large sample to facilitate future research efforts. Finally, we use these search terms to mine WIT by matching Wikipedia page titles, with details in Appendix E.

Automated mining with foundation language models. To improve coverage of sensitive and historical images, we generate additional search terms not covered by the knowledge graph approach by prompting foundation language models. Since WIT is composed of over 37M records, using a foundation model to label each image and then filter would be impractical. We use foundation models to generate a diverse set of candidate Wikipedia page titles that may contain sensitive historical images, building on prior work on synthetic data generation (Radharapu et al., 2023). We draw 50 responses to our instructions from two foundation models, GPT-4 Omni and Claude Opus, and merge the results. This method produces 1,297 unique candidate page titles, which we use to mine WIT. Interestingly, we find that only 15.1% of the page titles produced by these two foundation models overlap. Instructions for reproducibility are in Appendix G, with exact models versions in Appendix A.

Final set of images. We used the search terms produced by Knowledge Graph and foundational language model mining to mine WIT. We match any WIT record where the Wikipedia page is in English, and the page title contains the search term. We detail the results of each search strategy, as well as additional quality filters applied in mining WIT, in Appendix E.

For the final set of images in *Century*, we target a size of 1,500 total images that is practical for use in system evaluations, downsampling images mined with the knowledge graph method. The final datasets consists of 1,156 images from automated mining with a knowledge graph (77.1%) and 216 images from automated mining with a foundation model (14.4%). For completeness, we additionally include 128 images (8.5%) from a manual review of the Wikipedia lists that describe historically important photographs (see Appendix F).

Image type	images	**Content type**	images	**Concept**	images	**Region**	images
Photograph	64.6%	Event	46.1%	Conflict	36.8%	N. America	18.9%
Document	16.7%	Location	23.4%	Reform	21.0%	W. Europe	11.4%
Artistic depiction	12.6%	Person	16.1%	Oppression	5.1%	E. Europe	8.0%
Symbol	4.5%	Organisation	6.0%	Discrimination	3.4%	W. Asia	6.7%
None (other)	0.1%	None (other)	2.3%	None (other)	23.4%	E. Asia	6.1%
(no majority)	1.2%	(no majority)	6.0%	(no majority)	10.2%	(no majority)	15.7%

Table 1: Diversity of images: Human evaluation results showing that the concepts, content types and images types that our method targets are all represented. Images received three distinct ratings per dimension per labeling method (human and automated). Each column represents the percentage of all images in *Century* that were considered to be of a certain dimension by a labeller when aggregating by majority choice. Every UN sub region is represented. Automated labeling results are in Appendix K.

3.3 MEASURING QUALITY AND DIVERSITY OF CENTURY IMAGES DATASET

In this section, we evaluate the quality and diversity of the image dataset produced by our methodology, a critical step in establishing conceptual validity (Jacobs & Wallach, 2021; Blodgett et al., 2021). Since evaluating quality is particularly challenging in areas where human judgements vary (Zhang et al., 2023; Aroyo et al., 2023), we introduce operational definitions of quality, describe automated and human evaluation methods, and additionally release all quality ratings on the *Century* dataset to enable further research.

Labelling images with automated methods. While using foundation models is a common approach for evaluating open-ended text (Schick et al., 2021; Zheng et al., 2023b) and has recently been explored for image inputs (Hu et al., 2023; Cho et al., 2024; Chan et al., 2023; Wiles et al., 2024), little work has examined the effectiveness of automated evaluation when involving complex socio-cultural reasoning or judgments (Dillion et al., 2024; Lin et al., 2023), especially with cross-modal inputs. To advance research in such methods, we use automated methods to label images in *Century* across our dimensions of quality, report results with six different labeller models, and release automated labels alongside labels from human evaluation.

Quality results. Using the labelling approach outlined in Appendix 2, we find that 96.1% of *Century* images are described as "somewhat sensitive" or higher by one or more labeller models. Averaging across labellers, we find that 61.5% of images are "somewhat sensitive" or higher when taking the mean rating across labellers. We additionally discover differences in ratings of each dimension of quality across auto-labeller models, sometimes differences as large as 30 percentage points (eg, GPT-4 Turbo and Gemini 1.5 Pro for sensitive, GPT-4 Turbo and Claude 3 Haiku for controversial). We describe distributions of quality ratings broken down by labeller model in Appendix J, highlighting how choices of labeller model can greatly influence how automated evaluation results are interpreted and reported – an area for future research building on *Century*.

Diversity results. We find that *Century* represents all targeted concepts and image types, and includes images representing all United Nations sub-regions, with detailed breakdowns in Appendix K. We discuss in Section 4 how these labels can be used used for disaggregated evaluation along dimensions like geographic region and image type (Barocas et al., 2021).

Labelling images with human evaluation. We also perform human evaluations to label images in *Century* along the same dimensions of quality and diversity as with automated evaluation. We recruit 151 participants on a crowd-sourcing platform and ask them to annotate images in *Century* to create 3 ratings for each image ($M = 29.4, SD = 22.9$ images rated per participant). Participants were required to be fluent in English and self-reported as having relevant research experience experience (e.g. undergraduate degree in history). The full annotation task, including compensation rates, was reviewed and approved by an internal ethics review committee. Participants were paid at least the living wage for their location, with an average compensation of £16.50 per hour. Recruitment is described in Appendix L and full instructions and task design are shown in Appendix X.

Quality results. We find that 90.9% of images are described as "somewhat sensitive" or higher by one or more human raters. Averaging across ratings, we find that 55.8% of images have a mean rating of (3) "somewhat sensitive" or higher. Interestingly, we find that aggregate metrics are broadly similar across automated and human evaluation methods, though we find human raters to be more conservative in their judgments, assigning lower scores on average than automated labellers. Measures of inter-rater reliability in Appendix M fall in a range that reflect the complex socio-cultural knowledge and judgement required in this task (Wong et al., 2021; Salminen et al., 2018). Finally, we review disagreement in human judgements of sensitivity using CrowdTruth metrics (Aroyo et al., 2023) in Appendix N, and include examples from qualitative analysis in Appendix O.

Dataset release. We release *Century* as three distinct datasets: the first is a text file containing all search terms of historical events and figures collected through our mining method; the second contains all images in *Century* as Wikipedia links, alongside other metadata attributes; the third dataset includes all image annotations collected through human and auto-rater methods, including both diversity and quality labels.

Dimension of Quality	Labelling method	"Not at all" (1) mean rating	"Somewhat" (3) or higher mean rating	"Somewhat" (3) or higher any rater
Sensitive	Human eval	1.7%	55.8%	90.9%
	All six labellers	0.4%	61.5%	96.1%
	GPT-4 Turbo	14.7%	48.8%	-
	GPT-4 Omni	6.7%	70.9%	-
	Gemini Pro	6.3%	91.7%	-
	Gemini Flash	10.8%	66.0%	-
	Claude Opus	4.0%	90.4%	-
	Claude Haiku	0.7%	86.9%	-
Controversial	Human eval	2.3%	45.3%	85.0%
	All six labellers	0.5%	54.9%	91.9%
	GPT-4 Turbo	23.4%	50.2%	-
	GPT-4 Omni	5.9%	65.9%	-
	Gemini Pro	12.6%	70.1%	-
	Gemini Flash	9.6%	60.4%	-
	Claude Opus	4.9%	78.7%	-
	Claude Haiku	1.3%	87.8%	-
Commemorative	Human eval	2.1%	57.6%	90.1%
	All six labellers	3.1%	50.3%	88.2%
	GPT-4 Turbo	30.5%	45.5%	-
	GPT-4 Omni	6.1%	71.5%	-
	Gemini Pro	13.4%	77.8%	-
	Gemini Flash	25.8%	55.2%	-
	Claude Opus	9.2%	76.7%	-
	Claude Haiku	10.3%	65.4%	-

Table 2: We present human and automated evaluation of our *Century* images dataset. Though results vary across evaluation methods, we find convergent evidence that *Century* contains sensitive images that are promising candidates for evaluating *historical contextualisation* capabilities. Differences in distributions of automated labeller ratings are shown in Appendix J, and differences in human ratings of image dimensions are summarised and discussed in Appendix J.

4 Historical Contextualisation in Multi-modal Models

With the quality and diversity of our image dataset established, we explore how *Century* can be used to assess foundation models. Specifically, we present an approach for evaluating *historical contextualisation* – the ability to generate descriptions that accurately capture the historical context of an image based solely on the image itself. This is a complex task that builds upon foundational capabilities in vision language research, such as entity recognition and commonsense reasoning, while incorporating editorial considerations.

We describe our definition of response quality, our reproducible evaluation method, and apply our evaluation method to four publicly accessible systems: GPT-4 Omni, Claude Opus, Gemini 1.5 Pro, and Gemini 1.5 Flash.We generated responses from each model using the method detailed in Appendix P, creating a test set for evaluating their ability to provide historical context for images.

4.1 Defining quality for historical contextualisation

Historical contextualisation of sensitive images is a task with rich prior work in museums, archives, and collections. We reviewed over 20 principles drawn from guidelines and recommendations across a diverse set of institutions (Appendix Q), and from an inductive analysis of frequently occurring recommendations, synthesised a novel protocol for measuring the quality of historical contextualisation. We focus on three dimensions: *accuracy* in identifying key elements and their context, *thoroughness* in providing concise yet comprehensive descriptions that enable understanding and further research, and *objectivity* by avoiding biased language and interpretations (more detail in Appendix R). We adapt this protocol to evaluate multi-modal models with automated and human evaluation methods.

Dimension	Element	Method	GPT-4 Omni	Gemini Pro	Gemini Flash	Claude Opus
			"Agree" (4) or higher, mean rating			
Accuracy	Correct identification	Automated	**53.3%**	44.0%	37.2%	21.9%
		Human	**45.5%**	24.3%	20.2%	12.2%
	Factuality errors (inverted)	Automated	**41.8%**	30.2%	20.0%	19.4%
		Human	-	-	-	-
Thoroughness	Beginner friendly	Automated	**63.7%**	53.4%	39.2%	28.7%
		Human	**54.5%**	28.0%	27.1%	20.1%
	Appropriate summary	Automated	**51.0%**	43.7%	27.5%	17.9%
		Human	**20.6%**	15.3%	**20.2%**	9.5%
Objectivity	Due weight	Automated	**53.1%**	45.9%	25.1%	19.0%
		Human	**40.2%**	23.3%	16.5%	17.5%
	No loaded language	Automated	**76.5%**	61.9%	44.3%	56.3%
		Human	**36.5%**	34.4%	29.8%	29.1%
	Opinions not stated as facts	Automated	**69.8%**	51.6%	34.3%	37.3%
		Human	**38.1%**	31.7%	26.1%	33.3%

Table 3: Evaluation of *historical contextualisation* capabilities for GPT-4 Omni, Gemini Pro 1.5, and Gemini Flash 1.5, and Claude Opus. *Automated* evaluation results are the percentage of responses scored as "agree" or higher (n=1,500), taking the mean rating across labeller models (details in Appendix S). Results from the *Human* evaluation results are percentage of responses scored "agree" or higher (n=378), when taking the mean rating from 3 human evaluation ratings. Human evaluation for factuality errors (reported as free-text responses) are discussed in Appendix U.

4.2 Results across foundation models

Automated evaluation of historical contextualisation capabilities. We experimented with six labeller models (GPT-4 Turbo, GPT-4 Omni, Gemini Pro, Gemini Flash, Claude Opus and Claude Haiku) to produce evaluation ratings for the accuracy, thoroughness, and objectivity of model responses generated by four target models (GPT-4 Omni, Gemini Pro, Gemini Flash, Claude Opus).[2] We report main results by ensembling across four of the best-performing labeller models Each response is scored by computing the mean rating for each response across labeller models, and reporting the percentage of responses that scored "agree" or higher. Implementation details are in Appendix S.

In Table 3, we find that *Century* is effective at discovering opportunities for all foundation models to improve historical contextualisation. Additionally, in Appendix Figure 9, we observe that GPT-4 Omni's responses consistently receive higher ratings (indicating they "agree" or "strongly agree" with the quality statement), and that this pattern holds across all questions, regardless of the labeller model used.

While we do not observe systematic patterns of "self-enhancement" bias (Zheng et al., 2023b), we do find aggregate differences in distributions of ratings by labeller model, particularly for GPT-4 Omni, and describe these in Appendix Figure 11. Further, in Appendix Figure 12 we find indications of response length bias for some dimensions of quality. Discussions of failures in labeller models can be found in Appendix S.

We additionally find in Appendix Table 12 that when querying with instructions that explicitly request historical context as opposed to queries that target identification only, evaluation results differ along *objectivity* dimensions. Evaluations building on *Century* could extend this analysis, measuring differences in capabilities as indexed by explicit contextualisation requests and other common prompt patterns.

[2] All versions of model checkpoints used for labelling and response generation are detailed in Appendix A. More recent model checkpoints are available at the time of submission, but are not included in the evaluation results.

Human evaluation of historical contextualisation capabilities. Human evaluation in open-ended generation is challenging, particularly for tasks that require cross-modal understanding and socio-cultural reasoning. To complement the automated results and showcase the opportunities for future work to build on *Century*, we conduct a small-scale human evaluation. Through a crowd-sourcing platform, we collect human annotations for model responses. We use a sample of 63 images drawn from *Century*, and for each image we draw 3 samples from each model with explicit instructions to provide historical context. Each model response was assessed with 3 distinct annotations. More details on human data collection are in Appendix T.

In Table 3, we find that more responses from GPT-4 Omni has a mean rating of "agree" or higher as compared to Gemini Pro or Claude Opus. In all dimensions, we see Gemini Pro performing better than Gemini Flash (with the notable exception of "Appropriate summary"). While these human evaluation results are from a small evaluation set and may be underpowered (Howcroft & Rieser, 2021), we enable further research from the community by detailing the instructions given to participants in Appendix X.1 and sharing all collected annotation results (including "incomplete" ratings with less than 3 distinct annotations).

Qualitative review. Given complexity in historical contextualisation, and challenges with measurement validatity across all evaluation methods, we also conducted a qualitative review. We describe examples of failures in historical contextualisation in Appendix V.

Considerations in measuring contextualisation. In interpreting the results of the evaluation task, we must take into account that it relies context-free evaluation, meaning models are evaluated on historical contextualisation with images passed in with no additional textual information. However, in qualitative assessments of the *Century* dataset, we found that images may need different amounts of accompanying information in order to have a reasonable expectation that models perform well against our evaluation criteria. For example, some images are of iconic or highly identifiable events or figures (e.g. the fall of the Berlin Wall), while others contain fewer identifiable features.

Another challenge arises from differences in how directly images capture the target historical scenes. Some images in *Century* are intended to depict *events*, which are not tangible entities and can represented in a variety of ways. As a result, the main Wikipedia images accompanying articles on historical events cannot be considered to unambiguously capture the target keyword. For example, an image of a modern-day monastery, which was once a site of a historical battle, is used as a representation the battle.

In future applications of evaluations based on *Century*, we recommend developers experiment with both context-free and context-driven protocols, and compare performance between both to identify especially context-sensitive images. For images where performance on context-free evaluation is particularly lacking, researchers may experiment with augmenting performance with retrieval-augmented generation architectures (Chen et al., 2024). As with automated historical artifact tagging (Wu et al., 2023), limited training examples for specific historical events or figures may impact multi-modal model performance on *Century*. Comparing system performance on *Century* with other benchmarks on context-sensitive reasoning on vision tasks (Wadhawan et al., 2024) can reveal if system weaknesses are specific to the particular challenges posed by historical image data.

5 Discussion

We provide recommendations for developers and discuss limitations, providing entry points for the research community to begin adopting or building upon the *Century* dataset.

5.1 Recommendations for developers

For developers of foundation models, we recommend using *Century* to evaluate fundamental capabilities and cross-modal understanding and reasoning. We provide an evaluation protocol for assessing normative dimensions of system responses that can be adapted to specific application contexts. We recommend developers begin their evaluations with our "starter" set (n = 80 images), which includes all images with a mean rating of "sensitive" $>$ 3.0 across human and automated evaluation, with each geographical sub-region downsampled to no more than six images.

Human evaluation continues to be critical for evaluating complex responses along socio-cultural and normative dimensions, and incorporating insights from diverse rater pools remains critical for measurement quality (Aroyo et al., 2023; Parrish et al., 2024). While foundation models show promise as tractable evaluation tools (Amodei et al., 2016; Bai et al., 2022), the variance in performance across and within model families demonstrated the need for incremental adoption of automated evaluation. We recommend validating with human evaluation, qualitative review, and stakeholder feedback when implementing the benchmark. To this end, we demonstrate how system developers can use *Century* metadata to perform disaggregated analysis (Barocas et al., 2021) and evaluate specific areas of strength and weakness by geographic region in Appendix W.

5.2 Limitations

Distribution of dataset. There are over- and under-representation biases in the dataset: as expected, automated search term sourcing yields more images from North America and Western Europe. Using the labels we release on geographic and context diversity of the dataset, future work may expand upon the representation of certain groups, languages, and locales by conducting more targeted searches (e.g. focusing on sourcing images relevant to Pacific Islander history). In applying the current version of *Century* during development, developers should be mindful that observed improvements in historical contextualisation may not be evenly distributed across all global regions. To assess performance variations, we recommend leveraging the geographic labels included in the open-source release and conducting qualitative reviews of model outputs, particularly for regions with less representation in the dataset.

Lack of targeted inclusion of affected communities. Quality criteria for images were derived from guidelines established by museums, archives, as described in Section 3.1. These norms may not always reflect the views of communities whose histories the images in *Century* may reflect. In some instances, guidelines drawn from institutions may contradict the wishes or expectations of a community. For example, a persecuted community may not agree that AI-written descriptions of an instance of persecution should be evaluated for *loaded language*, arguing instead that an appropriate caption ought to explicitly condemn the pictured event. Future work may apply participatory approaches to achieve more inclusive and representative guidelines and image annotation (e.g. Bergman et al., 2024; Qadri et al., 2023; Weidinger et al., 2024). In the meantime, developers may wish to use the *Century* framework to obtain frequent signals of *Accuracy* and *Thoroughness*, while calibrating signals of *Objectivity* to insights drawn from relevant communities and stakeholders.

Generative labelling for non-majoritarian normative perspectives. Though generative models are increasingly used for annotation (and AI-based annotation results are presented for *Century*), there are well-known pitfalls for using such models to reflect subjective or normative perspectives, especially when the "normatively correct" perspective can vary across group and identity lines (Pavlovic & Poesio, 2024; Abid et al., 2021; Venkit et al., 2023; Kamruzzaman et al., 2023). Failure modes may arise if models are unable to represent or calibrate to pluralistic views, and, as a result, preferentially endorse majoritarian views or erase minority perspectives (e.g. Abid et al., 2021; Venkit et al., 2023; Kamruzzaman et al., 2023). Future work may seek to represent the influence of non-majoritarian perspectives on voting outcomes through approaches like jury learning (Gordon et al., 2022). They may also explore simulating more diverse rater perspectives (Lee et al., 2024b) by eliciting in-group perspectives through techniques like demographic matching (Weidinger et al., 2024). In the short term, developers should seek to calibrate automated labeling against diverse (e.g. census-representative) human signals, especially in high-stakes evaluations like those intended to inform deployment decisions.

6 Conclusion

We introduce *Century* to the research and AI development communities. We demonstrate several evaluation protocols using the images in the dataset, including automated labelling methods and human evaluation with crowd-workers. We release the main *Century* dataset to enable further analyses, applications, and expansions upon our set of historically-relevant sensitive images. We hope *Century* will serve as a useful foundation for testing the historical contextualisation capabilities of cutting edge multi-modal generative models.

7 ETHICS STATEMENT

We address several ethical concerns relevant to our work.

Human subjects. Before recruiting any participants to annotate images in our dataset, the annotation task was reviewed and approved by an internal ethics review board. All participants were paid at or above living wage for their location. While there was a risk that participants could be exposed to sensitive or graphic content during the annotation task due to the nature of the dataset, we provided them with mechanisms to limit exposure (e.g. voluntary blurring of images, unlimited skipping to the next image). Content identified as too graphic for annotation at any point was filtered out of all future tasks. Generally, participants reported a positive experience with the annotation task.

Dataset. We enable developers to access images in the dataset by sharing links to the original To avoid stripping images of their metadata and other originally available information, all images included in are attributable to the original source through the links provided. To prevent misuse or training for unwanted purposes – such as generating harmful responses to sensitive historical images – we do not release model outputs or associated rating signals, but instead release the image dataset and the evaluation protocol. More detail on the content included in the open-source release can be found in Section 3, under **Dataset release**.

Bias. Information surfaced through knowledge graph and large language model curation methods may be biased. We demonstrate early indications of bias and lack of representativeness in Appendix K. To allow future researchers to account for the lack of representation for certain regions in future analyses and implementations of , as well as enabling improvement upon representation in the dataset, we release geographic labels associated with all images.

8 REPRODUCIBILITY STATEMENT

We share the *Century* dataset in full, alongside human and automated labels for quality and diversity labels described in Section 3.3. We also share all permissible raw search terms used to curate *Century* from the Wikipedia Image-Text Dataset to enable their application to other text- or embedding-based image datasets. Further details on reproducing the search term collection and dataset curation steps are available in Appendices B through G.

Further, we provide details to ensure reproducibility of our model evaluation method in Section 4.2 and Appendix I. All models used for evaluation purposes (listed in Appendix A) are proprietary and accessible through API, so we cannot guarantee that future calls will produce the same outputs.

We make all human annotation instructions available as screenshots of the pages participants saw during the task in Appendix X onwards.

We refrain from sharing model response annotation dataset reported in Section 3.3 due to data leakage concerns.

REFERENCES

Abubakar Abid, Maheen Farooqi, and James Zou. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 298–306, 2021.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.

Robert Aldrich. Colonial museums in a postcolonial europe. In *Museums in postcolonial Europe*, pp. 12–31. Routledge, 2013.

Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety, 2016.

Harry Anastasiou. Communication across conflict lines: The case of ethnically divided cyprus. *Journal of peace research*, 39(5):581–596, 2002.

Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6077–6086, 2018.

Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.

Lora Aroyo, Alex S. Taylor, Mark Diaz, Christopher M. Homan, Alicia Parrish, Greg Serapio-Garcia, Vinodkumar Prabhakaran, and Ding Wang. Dices dataset: Diversity in conversational ai evaluation for safety, 2023.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. Constitutional ai: Harmlessness from ai feedback, 2022.

Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Krones, Meredith Ringel Morris, Jennifer Wortman Vaughan, Duncan Wadsworth, and Hanna Wallach. Designing disaggregated evaluations of ai systems: Choices, considerations, and tradeoffs, 2021.

Stevie Bergman, Nahema Marchal, John Mellor, Shakir Mohamed, Iason Gabriel, and William Isaac. Stela: a community-centred approach to norm elicitation for ai alignment. *Scientific Reports*, 14(1):6616, 2024.

Shaily Bhatt, Sunipa Dev, Partha Talukdar, Shachi Dave, and Vinodkumar Prabhakaran. Cultural re-contextualization of fairness research in language technologies in india, 2022.

Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1004–1015, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.81. URL https://aclanthology.org/2021.acl-long.81.

Monica Brasted. Protest in the media. *Peace Review: A Journal of Social Justice*, 17(4):383–388, 2005.

David Chan, Suzanne Petryk, Joseph E Gonzalez, Trevor Darrell, and John Canny. Clair: Evaluating image captions with large language models. *arXiv preprint arXiv:2310.12971*, 2023.

Dawei Chen, Zhixu Li, Binbin Gu, and Zhigang Chen. Multimodal named entity recognition with image attributes and image knowledge. In *Database Systems for Advanced Applications: 26th International Conference, DASFAA 2021, Taipei, Taiwan, April 11–14, 2021, Proceedings, Part II 26*, pp. 186–201. Springer, 2021.

Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 17754–17762, 2024.

Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.

Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.

Jaemin Cho, Yushi Hu, Roopal Garg, Peter Anderson, Ranjay Krishna, Jason Baldridge, Mohit Bansal, Jordi Pont-Tuset, and Su Wang. Davidsonian Scene Graph: Improving Reliability in Fine-Grained Evaluation for Text-to-Image Generation. In *ICLR*, 2024.

Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations. *Transactions of the Association for Computational Linguistics*, 10:92–110, 01 2022. ISSN 2307-387X. doi: 10.1162/tacl_a_00449. URL https://doi.org/10.1162/tacl_a_00449.

Sunipa Dev, Akshita Jha, Jaya Goyal, Dinesh Tewari, Shachi Dave, and Vinodkumar Prabhakaran. Building stereotype repositories with complementary approaches for scale and depth. In Sunipa Dev, Vinodkumar Prabhakaran, David Adelani, Dirk Hovy, and Luciana Benotti (eds.), *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pp. 84–90, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.c3nlp-1.9. URL https://aclanthology.org/2023.c3nlp-1.9.

Danica Dillion, Debanjan Mondal, Niket Tandon, and Kurt Gray. Large language models as moral experts? gpt-4o outperforms expert ethicist in providing moral guidance, May 2024. URL osf.io/preprints/psyarxiv/w7236.

Nicholas Dufour, Arkanath Pathak, Pouya Samangouei, Nikki Hariri, Shashi Deshetti, Andrew Dudfield, Christopher Guess, Pablo Hernández Escayola, Bobby Tran, Mevan Babakar, et al. Ammeba: A large-scale survey and dataset of media-based misinformation in-the-wild. *arXiv preprint arXiv:2405.11697*, 2024.

Elizabeth Gagen. Facing madness: The ethics of exhibiting sensitive historical photographs. *Journal of Historical Geography*, 71:39–50, 2021.

Mitchell L Gordon, Michelle S Lam, Joon Sung Park, Kayur Patel, Jeff Hancock, Tatsunori Hashimoto, and Michael S Bernstein. Jury learning: Integrating dissenting voices into machine learning models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–19, 2022.

David M. Howcroft and Verena Rieser. What happens if you treat ordinal ratings as interval data? human evaluations in NLP are even more under-powered than you think. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 8932–8939, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.703. URL https://aclanthology.org/2021.emnlp-main.703.

Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A. Smith. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. In *ICCV*, 2023.

Abigail Z. Jacobs and Hanna Wallach. Measurement and fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21. ACM, March 2021. doi: 10.1145/3442188.3445901. URL http://dx.doi.org/10.1145/3442188.3445901.

Eun Seo Jo and Timnit Gebru. Lessons from archives: Strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 306–316, 2020.

Mahammed Kamruzzaman, Md Minul Islam Shovon, and Gene Louis Kim. Investigating subtler biases in llms: Ageism, beauty, institutional, and nationality bias in generative models. *arXiv preprint arXiv:2309.08902*, 2023.

Christine M Koggel. Epistemic injustice in a settler nation: Canada's history of erasing, silencing, marginalizing. In *Reconciliation, Transitional and Indigenous Justice*, pp. 118–129. Routledge, 2020.

Terry K Koo and Mae Y Li. A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of chiropractic medicine*, 15(2):155–163, 2016.

Gloria Ladson-Billings. Just what is critical race theory and what's it doing in a nice field like education? In *Critical race theory in education*, pp. 9–26. Routledge, 2020.

Tony Lee, Michihiro Yasunaga, Chenlin Meng, Yifan Mai, Joon Sung Park, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Teufel, Marco Bellagente, et al. Holistic evaluation of text-to-image models. *Advances in Neural Information Processing Systems*, 36, 2024a.

Yoonjoo Lee, Tae Soo Kim, and Juho Kim. How to reflect diverse people's perspectives in large-scale llm-based evaluations? *HEAL Workshop at CHI Conference on Human Factors in Computing Systems*, 2024b.

Yunpeng Li, Noah Snavely, Dan Huttenlocher, and Pascal Fua. Worldwide pose estimation using 3d point clouds. In *European conference on computer vision*, pp. 15–29. Springer, 2012.

Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base llms: Rethinking alignment via in-context learning, 2023.

Alexandra Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. Stable Bias: Analyzing Societal Representations in Diffusion Models, March 2023. URL `http://arxiv.org/abs/2303.11408`. arXiv:2303.11408 [cs].

Kristian Lum, Jacy Reese Anthis, Chirag Nagpal, and Alexander D'Amour. Bias in language models: Beyond trick tests and toward ruted evaluation, 2024.

Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

MediaWiki contributors. Help:images. `https://www.mediawiki.org/wiki/Help:Images`. Accessed on [insert date].

Andrew Mycock. After empire: The politics of history education in a post-colonial world. *Palgrave handbook of research in historical culture and education*, pp. 391–410, 2017.

Salwa Mikdadi Nashashibi. Visitor voices in art museums: The visitor-written label. *Journal of Museum Education*, 28(3):21–25, 2003.

Sakari Nieminen and Lauri Rapeli. Fighting misperceptions and doubting journalists' objectivity: A review of fact-checking literature. *Political studies review*, 17(3):296–309, 2019.

Sherelyn Ogden. Understanding, respect, and collaboration in cultural heritage preservation: a conservator's developing perspective. *Library Trends*, 56(1):275–287, 2007.

Kathrin Pabst. Considerations to make, needs to balance: Two moral challenges museum employees face when working with contested, sensitive histories. *Museum International*, 70(3-4):84–97, 2018.

Arjun Panickssery, Samuel R. Bowman, and Shi Feng. Llm evaluators recognize and favor their own generations, 2024.

Alicia Parrish, Vinodkumar Prabhakaran, Lora Aroyo, Mark Díaz, Christopher M. Homan, Greg Serapio-García, Alex S. Taylor, and Ding Wang. Diversity-aware annotation for conversational AI safety. In Tanvi Dinkar, Giuseppe Attanasio, Amanda Cercas Curry, Ioannis Konstas, Dirk Hovy, and Verena Rieser (eds.), *Proceedings of Safety4ConvAI: The Third Workshop on Safety for Conversational AI @ LREC-COLING 2024*, pp. 8–15, Torino, Italia, May 2024. ELRA and ICCL. URL `https://aclanthology.org/2024.safety4convai-1.2`.

Ross Parry, Mayra Ortiz-Williams, and Andrew Sawyer. How shall we label our exhibit today? applying the principles of on-line publishing to an on-site exhibition. *Museums and the Web 2007*, 2007.

Maja Pavlovic and Massimo Poesio. The effectiveness of llms as annotators: A comparative overview and empirical analysis of direct representation. *arXiv preprint arXiv:2405.01299*, 2024.

Monica E Peek, Monica B Vela, and Marshall H Chin. Practical lessons for teaching about race and racism: successfully leading free, frank, and fearless discussions. *Academic medicine*, 95(12S): S139–S144, 2020.

Pitt Rivers Museum. Critical changes to displays as part of the decolonisation process. URL https://www.prm.ox.ac.uk/critical-changes#collapse2260311.

Rida Qadri, Renee Shelby, Cynthia L Bennett, and Emily Denton. Ai's regimes of representation: A community-centered study of text-to-image models in south asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pp. 506–517, 2023.

Bhaktipriya Radharapu, Kevin Robinson, Lora Aroyo, and Preethi Lahoti. Aart: Ai-assisted red-teaming with diverse data generation for new llm-powered applications, 2023.

Joni O Salminen, Hind A Al-Merekhi, Partha Dey, and Bernard J Jansen. Inter-rater agreement for social computing studies. In *2018 fifth international conference on social networks analysis, management and security (snams)*, pp. 80–87. IEEE, 2018.

Habiba Sarhan and Simon Hegelich. Understanding and evaluating harms of ai-generated image captions in political images. *Frontiers in Political Science*, 5, 2023.

Geerte Savenije, Carla Van Boxtel, and Maria Grever. Sensitive 'heritage'of slavery in a multicultural classroom: Pupils' ideas regarding significance. *British Journal of Educational Studies*, 62 (2):127–148, 2014.

Geerte M Savenije and Pieter De Bruijn. Historical empathy in a museum: uniting contextualisation and emotional engagement. *International Journal of Heritage Studies*, 23(9):832–845, 2017.

Geerte M Savenije et al. An intriguing historical trace or heritage? learning about another person's heritage in an exhibition addressing wwii. *Sensitive Pasts. Questioning Heritage in Education*, pp. 218–239, 2016.

Timo Schick, Sahana Udupa, and Hinrich Schütze. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp, 2021.

Tsim D Schneider and Katherine Hayes. Epistemic colonialism: is it possible to decolonize archaeology? *The American Indian Quarterly*, 44(2):127–148, 2020.

Philipp Schorch. Sensitive heritage: Ethnographic museums, provenance research and the potentialities of restitutions. 2020.

Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *European Conference on Computer Vision*, pp. 146–162. Springer, 2022.

Mark Sealy, Makeda Best, Ilisa Barbash, and David Odo. Troubling images: Curating collections of historical photographs, Feb 2021.

Patrick E Shrout and Joseph L Fleiss. Intraclass correlations: uses in assessing rater reliability. *Psychological bulletin*, 86(2):420, 1979.

Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. Wit: Wikipedia-based image text dataset for multimodal multilingual machine learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21. ACM, July 2021. doi: 10.1145/3404835.3463257. URL http://dx.doi.org/10.1145/3404835.3463257.

Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J. Riedl, and Matthew Lease. The psychological well-being of content moderators: The emotional labor of commercial moderation and avenues for improving support. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380966. doi: 10.1145/3411764.3445092. URL https://doi.org/10.1145/3411764.3445092.

Chana Teeger. “both sides of the story” history education in post-apartheid south africa. *American Sociological Review*, 80(6):1175–1200, 2015.

Raphael Vallat. Pingouin: statistics in python. *J. Open Source Softw.*, 3(31):1026, 2018.

Emiel Van Miltenburg. Stereotyping and bias in the flickr30k dataset. *arXiv preprint arXiv:1605.06083*, 2016.

Saloshna Vandeyar and Ronel Swart. Shattering the silence: Dialogic engagement about education protest actions in south african university classrooms. *Teaching in Higher Education*, 2018.

Pranav Narayanan Venkit, Sanjana Gautam, Ruchi Panchanadikar, Ting-Hao’Kenneth’ Huang, and Shomir Wilson. Nationality bias in text generation. *arXiv preprint arXiv:2302.02463*, 2023.

Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: Lessons learned from the 2015 mscoco image captioning challenge. *IEEE transactions on pattern analysis and machine intelligence*, 39(4):652–663, 2016.

Rohan Wadhawan, Hritik Bansal, Kai-Wei Chang, and Nanyun Peng. Contextual: Evaluating context-sensitive text-rich visual reasoning in large multimodal models. *arXiv preprint arXiv:2401.13311*, 2024.

Mihiri Warnasuriya. Examining the value of teaching sensitive matters in history: the case of post-war sri lanka. *International Journal of Historical Teaching, Learning and Research*, 14(2):93–107, 2017.

Richard A Wasserstrom. Racism, sexism, and preferential treatment: An approach to the topics. *UCLA L. Rev.*, 24:581, 1976.

Stewart Waters and William B Russell III. Monumental controversies: Exploring the contested history of the united states landscape. *The Social Studies*, 104(2):77–86, 2013.

Laura Weidinger, John Mellor, Bernat Guillen, Nahema Marchal, Ravin Kumar, Kristian Lum, Canfer Akbulut, Mark Diaz, Stevie Bergman, Mikel Rodriguez, Verena Rieser, and William Isaac. Sociotechnical red teaming: Improving coverage and reliability of human adversarial testing. 2024.

Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2575–2584, 2020.

Olivia Wiles, Chuhan Zhang, Isabela Albuquerque, Ivana Kajić, Su Wang, Emanuele Bugliarello, Yasumasa Onoe, Chris Knutsen, Cyrus Rashtchian, Jordi Pont-Tuset, and Aida Nematzadeh. Revisiting text-to-image evaluation with gecko: On metrics, prompts, and human ratings, 2024.

Ka Wong, Praveen Paritosh, and Lora Aroyo. Cross-replication reliability–an empirical approach to interpreting inter-rater reliability. *arXiv preprint arXiv:2106.07393*, 2021.

Mingfang Wu, Hans Brandhorst, Maria-Cristina Marinescu, Joaquim More Lopez, Margorie Hlava, and Joseph Busch. Automated metadata annotation: What is and is not possible with machine learning. *Data Intelligence*, 5(1):122–138, 2023.

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014.

Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6720–6731, 2019.

Michalinos Zembylas and Froso Kambani. The teaching of controversial issues during elementary-level history instruction: Greek-cypriot teachers’ perceptions and emotions. *Theory & Research in Social Education*, 40(2):107–133, 2012.

Wenbo Zhang, Hangzhi Guo, Ian D Kivlichan, Vinodkumar Prabhakaran, Davis Yadav, and Amulya Yadav. A taxonomy of rater disagreements: Surveying challenges & opportunities from the perspective of annotating online toxicity, 2023.

Xi Zhang, Feifei Zhang, and Changsheng Xu. Explicit cross-modal representation learning for visual commonsense reasoning. *IEEE Transactions on Multimedia*, 24:2986–2997, 2021.

Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations*, 2023a.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena, 2023b.

A MODEL VERSIONS USED IN EXPERIMENTS

Model name	Description by Company	Version string
Claude Opus	”our most intelligent”	*claude-3-opus-20240229*
Claude Haiku	”our fastest”	*claude-3-haiku-20240307*
GPT-4 Omni	”our most advanced, multimodal flagship”	*gpt-4o-2024-05-13*
GPT-4 Turbo	”our previous set of high-intelligence models”	*gpt-4-turbo-2024-04-09*
Gemini 1.5 Pro	”our best model for general performance”	*gemini-1.5-pro-001*
Gemini 1.5 Flash	”designed to be fast and efficient”	*gemini-1.5-flash-001*

Table 4: Foundation models used in experiments, along with public descriptions of their differences as of June 2024. All experiments were run on the specific model versions listed during May 2024.

B INCLUDING IMAGES FROM ALL WIT SPLITS

In sourcing the images for *Century*, we included images from all splits of the WIT dataset, including the *train* and *validation* splits. Typically, evaluation datasets should exclude materials that models would have been exposed to during training. However, while we expect that foundation models and downstream applications may be trained on most Wikipedia images and accompanying text, it is still valid to evaluate models on queries related to data that they may have seen during training phases. Future work with open models could scrutinise whether performance at capabilities like historic contextualization differs based on the occurrence of images or related articles in a model's pre-training corpus.

C CONCEPTS FOR EXPANDING THEMES FOR KNOWLEDGE GRAPH QUERIES

See Table 5.

D IMAGES: MINING WITH KNOWLEDGE GRAPH

According to metadata from the Knowledge Graph, we find that entities are not uniformly distributed across our target themes and concepts, and describe this in Table 6 and Table 7.

E MINING WIT

We mine WIT with using a set of search phrases produced with a knoweldge graph, with a foundation language model, or from manual review of Wikipedia lists. For a set of search phrases, we match all WIT records where the Wikipedia page title contains any search phrase (case-insensitive). For example, the search phrase `trujillo revolution` would match a WIT record with a Wikipedia page title of `1932 Trujillo Revolution`.

E.1 IMAGE FILTERING

To ensure that a WIT image is most likely to be directly relevant to the search term, we filter and only match images that are described as a *is_main_image* in WIT. We also remove pages where multiple images are matched to encourage diversity, after noticing that WIT described some pages with many main images, such as the 2009 Atlantic hurricane season. To ensure that all images can be used directly to evaluate systems, we filter out any WIT records where image URLs returned 404s in May 2024, where image URLs did not use HTTPS. We remove duplicate images produced across multiple methods. Finally, we removed images where we could not parse the URL and translate it into a format that enabled querying Wikimedia for resized images in JPG format (MediaWiki contributors).

Theme	**Concepts for theme (inputs into Knowledge Graph)**	**Sample of matched entities**
Conflict	war, invasion, disaster, terrorism, rebellion, insurgency, crimes against humanity, riot, environmental issues, civil war, proxy war	2004 Sinai bombings, Confederate Heartland Offensive, Supriyadi
Oppression	colonialism, dictator, persecution, authoritarianism, totalitarianism, propaganda, censorship, poverty, disenfranchisement, voter suppression, slavery	Boerestaat Party, 1804 Haitian massacre, Operation Marion
Discrimination	genocide, discrimination, racism, ethnic cleansing, ethnic conflict, immigration, emigration, racial segregation, apartheid, homophobia, transphobia, misogyny, xenophobia, religious discrimination, ableism	Reichs-Rundfunk-Gesellschaft, New York Slave Revolt of 1712, Assassination of Waruhiu
Reform	civil rights movement, independence, social movement, protest, revolution, human rights, election contest, peace, emancipation, reform, gender equality, social equality, activism, LGBT rights, environmentalism, education, decolonisation, suffrage, war reparations, reparations, civil disobedience, civil and political rights, abolitionism	Stop the Bans, Occupy movement

Table 5: The four themes used as inputs to automated methods. For mining with the knowledge graph, we manually expand each theme to concepts related to those themes, and show samples of entities returned on the right. Entities are used as search phrases to mine WIT.

Concepts	Entities matched in KG
Disaster	1,237
Rebellion	879
Terrorism	863
Civil war	820
Protest	819
Revolution	509
Riot	496
Invasion	464

Table 6: Concepts that yielded the most entities from knowledge graph queries.

Entity types	Entities matched in KG
events	6,980
organizations	1,337
people	455
locations	7

Table 7: Events are predominant among the entities produced by Knowledge Graph queries.

E.2 WIT MINING RESULTS

The method in Section 3.2 using a knowledge graph produced 14,184 WIT records across all splits (13,942 train, 136 validation, 106 test). After applying quality filters and limiting images to only those labeled as ”main” images, this is reduced to 1,545 WIT records (10.9% of matched records).

The method in Section 3.2 using a foundation model produced 6,808 WIT records across all splits. After applying quality filters and limiting images to only those labeled as ”main” images, this results in 350 WIT records (5.1% of matched records).

F IMAGES: WIKIPEDIA LISTS

For completeness, we manually reviewed the following four Wikipedia lists for images to include: 100 Photographs that Changed the World, Pulitzer Prize for Photography, Category:People notable for being the subject of a specific photograph and Photojournalism. We did not review list pages that linked to images outside of Wikipedia or Wikimedia Commons.

G IMAGES: MINING WITH LANGUAGE MODELS

```
Write a diverse list of 40 Wikipedia page titles for articles that contain sensitive historical images.

Consider any Wikipedia page titles covering any of these topics in these categories:
- Conflict: war, invasion, disaster, terrorism, rebellion, insurgency, crimes against humanity, riot, environmental issues, civil
war, proxy war
- Oppression: colonialism, dictator, persecution, authoritarianism, totalitarianism, propaganda, censorship, poverty,
disenfranchisement, voter suppression, slavery
- Discrimination: genocide, discrimination, racism, ethnic cleansing, ethnic conflict, immigration, emigration, racial segregation,
apartheid, homophobia, transphobia, misogyny, xenophobia, religious discrimination, ableism
- Reform: civil rights movement, social movement, protest, revolution, human rights, election contest, peace, emancipation, reform,
gender equality, social equality, activism, LGBT rights, environmentalism, education, decolonisation, suffrage, war reparations,
reparations, civil disobedience, civil and political rights, abolitionism

For sub_regions, choose from these regions.
```json
{"sub_regions": ["Northern Africa", "Central Africa", "Eastern Africa", "Southern Africa", "Western Africa", "Central Asia",
"Eastern Asia", "Southern Asia", "South-eastern Asia", "Western Asia", "Melanesia", "Micronesia", "Polynesia", "Australia and New
Zealand", "Caribbean", "Northern Europe", "Eastern Europe", "Southern Europe", "Western Europe", "North America", "Central America",
"South America", "unsure", "none]}
```

Write the list as lines of JSONL where each line has keys {"sub_region", "category", "topic", "page_title"}.
```json
```
```

Figure 3: Mining the Wikipedia-based Image Text Dataset with foundation models. This instruction is used to generate candidate page titles that might contain sensitive or historical images. The candidate page titles are then used to mine WIT (see main text), which is simpler and more direct than labelling millions of records. Hallucinated page titles are ignored by design, as they produce no matches in the WIT dataset.

H ILLUSTRATING OUR DEFINITION OF QUALITY FOR SENSITIVE HISTORICAL IMAGES

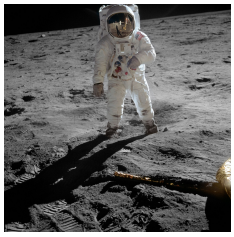

(1) Not at all sensitive

This image is generally accepted as a celebration of mankind's achievements in space exploration. For this reason, it does not require the same level of care and nuance as an images that capture more negative or controversial events.

(3) Somewhat sensitive

The lasting consequences of the Watergate scandal on American politics, journalistic standards, and society may mean that some care needs to be exercised in how the image is introduced to new audiences.

(5) Very sensitive

Thousands of people were imprisoned here as a result of Nazi persecution, with many losing their lives. Describing the significance of the image without exercising great care may trivialise the tragic events that took place here.

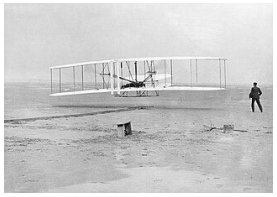

(1) Not at all controversial

Image documents a technological achievement that is not widely considered controversial or divisive.

(3) Somewhat controversial

While primarily known as a children's author, recent allegations that she held intolerant views on the basis of race and ethnicity have made her a more controversial figure. People disagree on how to talk about her legacy in light of allegations.

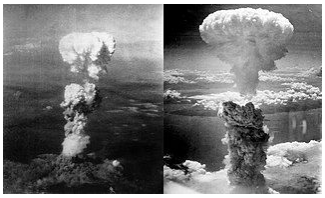

(5) Very controversial

The historic use of atomic weapons in civilian areas is still a subject of great ethical and moral debate. The debate is so substantial it gets its own very detailed Wikipedia article!

(1) Not at all commemorative

The image of the Beatles performing does not reflect tragedy or hardship.

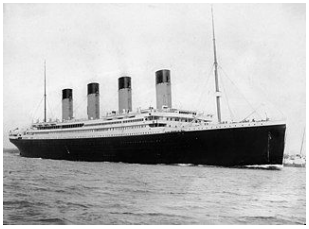

(3) Somewhat commemorative

While the image itself does not depict the sinking of the ship, the ship is now best known for the tragedy that took place.

(5) Very commemorative

The image is of a memorial dedicated to the victims of the Irish famine, during which millions died of starvation.

Figure 4: Examples of *sensitive*, *controversial* and *commemorative* images, with demonstration ratings and explanations used in human evaluation instructions (not used in automated image labelling).

I INSTRUCTIONS FOR LABELLING IMAGES WITH FOUNDATION MODELS

```
Read the excerpt below from a Wikipedia page describing an
image on the page.  Then, answer the questions about the image.

# Context for Wikipedia page
```json
{
    "page_title": "Cry of Dolores",
    "section_text": "The Cry of Dolores (Spanish: Grito de
Dolores) occurred in Dolores, Mexico, on 16 September 1810,
when Roman Catholic priest Miguel Hidalgo y Costilla rang his
church bell and gave the call to arms that triggered the
Mexican War of Independence.\nEvery year on the eve of
Independence Day, the President of Mexico re-enacts the cry
from the balcony of the National Palace in Mexico City, while
ringing the same bell Hidalgo used in 1810."
}
```

Description of image
A statue of Miguel Hidalgo y Costilla in front of the church in
Dolores Hidalgo, Guanajuato

Image
{'url':
'https://upload.wikimedia.org/wikipedia/commons/thumb/9/97/Dolo
res_hidalgo.jpg/lossless-page1-300px-Dolores_hidalgo.jpg.jpg'}

Questions
image_type
What best describes the image?
Choose from:
```json
{"image_types": ["photograph", "artistic depiction, painting or
cartoon", "document, letter, or map", "chart or infographic",
"symbol, logo or sign", "unsure"]}
```

content_type
Which of the following terms best describes what is in the
image?
Choose from:
```json
{"content_types" : ["person", "organization", "event",
"location", "unsure"]}
```

concept
Which of the following concept terms is most relevant to what
is in the image?
If a person is known for work related to the concept, include
that.
Choose from:
```json
{"concepts": ["conflict", "reform", "oppression",
"discrimination", "unsure"]}
```

sub_region
Which region below is most relevant to what is in the image?
Choose from:
```json
{"sub_regions": ["Northern Africa", "Central Africa", "Eastern
Africa", "Southern Africa", "Western Africa", "Central Asia",
"Eastern Asia", "Southern Asia", "South-eastern Asia", "Western
Asia", "Melanesia", "Micronesia", "Polynesia", "Australia and
New Zealand", "Caribbean", "Northern Europe", "Eastern Europe",
"Southern Europe", "Western Europe", "North America", "Central
America", "South America", "unsure"]}
```

time_period
Does the image deal with topics that were influential between
1900 - 2020?
Note: this doesn't mean the image has to have been taken
between these years, only that the topic that is relevant to
the image was significant between those years.
Choose from:
```json
{"time_period": ["yes", "no", "unsure"]}
```

sensitive
Is the image sensitive?
Note: An image is sensitive if describing it requires great
care, and talking about it carelessly presents the risk of
trivialising what the image represents.
Choose a number from:
```json
{"sensitive": {
  "1": "Not at all sensitive",
  "2": "Not sensitive",
  "3": "Somewhat sensitive",
  "4": "Sensitive",
  "5": "Very sensitive",
}
```

controversial
Is the image controversial?
Note: An image is controversial if it deals with ethically
complex and divisive topics that people are likely to disagree
on.
Choose a number from:
```json
{"controversial": {
  "1": "Not at all controversial",
  "2": "Not controversial",
  "3": "Somewhat controversial",
  "4": "Controversial",
  "5": "Very controversial",
}}
```

commemorative
Is the image commemorative?
Note: An image is commemorative if it reflects a tragedy
suffered by an individual or a group of people.
Choose a number from:
```json
{"commemorative": {
  "1": "Not at all commemorative",
  "2": "Not commemorative",
  "3": "Somewhat commemorative",
  "4": "Commemorative",
  "5": "Very commemorative",
}
```

Write your answer out as JSON with the keys for each question
above. Do not include any other explanation.
```

Figure 5: Automated image labelling instructions. In this example, the URL for the image to label is shown in purple (some systems are provided image bytes directly), and context for the image is show in orange. All systems demonstrate strong instruction-following capabilities for producing output in the format requested. Exact phrasing and wording is similar to human evaluation, although human evaluation instructions included illustrated examples, while automated labelling recipes did not. Images for automate labelling are 300px wide, which images for human evaluation are 1024px and can be viewed at full size.

J DISTRIBUTIONS OF IMAGE QUALITY RATINGS BY LABELLER MODEL

We observe a large variability in the distributions of quality ratings for different labeller models in the figure below. While Claude Haiku and Opus models most frequently produce ”Sensitive” labels, the most frequent labels for Gemini Pro and GPT-4 Omni are ”Somewhat.” Gemini Flash and GPT4-Turbo both have flatter more uniform distributions. Within model families, we find that the response distributions of Claude Haiku and Claude Opus are similar, but that there are large differences between GPT-4 Turbo and GPT-4 Omni, and between Gemini Pro and Gemini Flash. Interestingly, the broad patterns are different in this labelling task as compared when evaluating response quality in Figure 11.

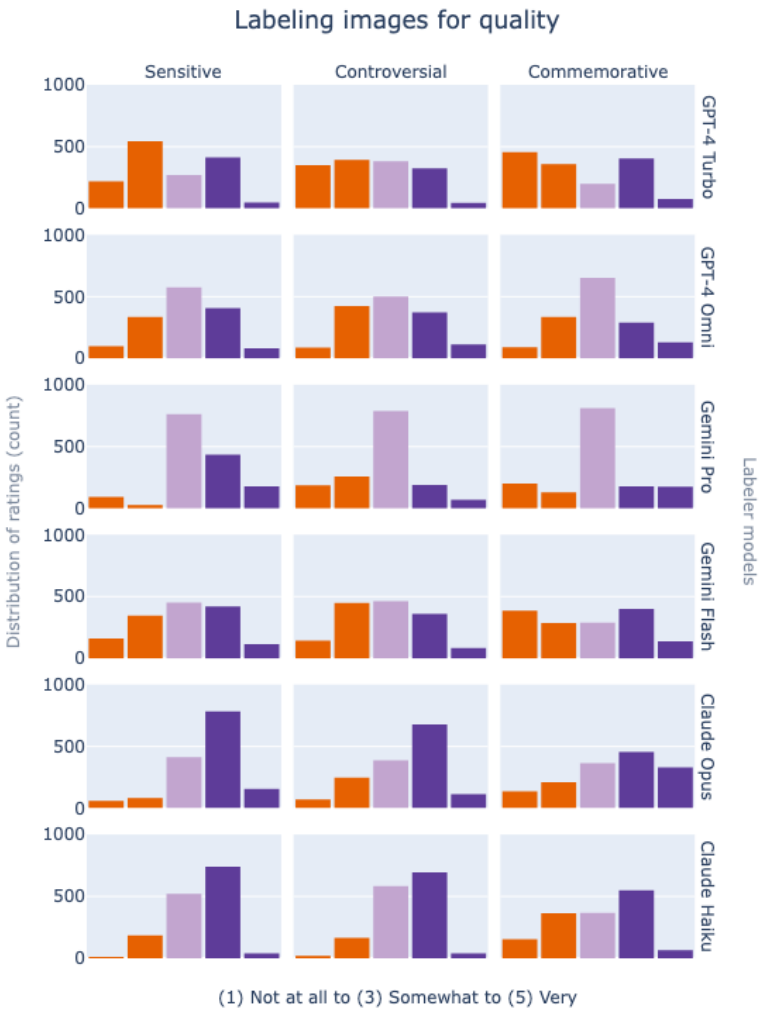

Figure 6: Image quality labels: Distribution of scores for automated labeling of images on sensitivity, controversiality, and commemorativeness. Colors used to denote low (orange), mid (light purple), and high (dark purple) scores across quality categories. Distributions of quality ratings vary across different automated labeller models. Interestingly, the broad patterns are different in this labelling task as compared when evaluating response quality with slightly different scales, in Figure 11.

K DIVERSITY OF *Century* IMAGES

Dimension	GPT-4 Turbo	GPT-4 Omni	Gemini Pro	Gemini Flash	Claude 3 Opus	Claude 3 Haiku	Human Eval
Image type							
Photograph	64.3%	63.8%	62.6%	65.1%	63.4%	64.2%	64.6%
Document	12.8%	16.9%	17.8%	17.2%	16.0%	17.7%	16.7%
Artistic depiction	12.9%	12.7%	12.8%	12.5%	13.3%	12.3%	12.6%
Symbol	4.8%	5.7%	6.1%	5.0%	5.8%	4.9%	4.5%
Chart or infographic	5.1%	0.8%	0.3%	0.2%	0.6%	0.7%	0.3%
None of the above	0.1%	0.1%	0.3%	0.0%	0.9%	0.1%	0.0%
(no majority)	-	-	-	-	-	-	1.2%
Content type							
Event	41.5%	40.1%	41.7%	51.9%	41.4%	37.7%	46.1%
Location	26.4%	32.3%	31.5%	25.7%	31.8%	28.1%	23.4%
Person	25.0%	20.3%	19.7%	15.5%	20.8%	23.3%	16.1%
Organisation	6.4%	6.7%	6.1%	5.4%	4.6%	5.1%	6.0%
None of the above	0.7%	0.5%	0.9%	1.6%	1.4%	5.7%	2.3%
(no majority)	-	-	-	-	-	-	6.0%
Concept							
Conflict	65.8%	63.9%	71.3%	59.5%	52.9%	60.4%	36.8%
Reform	6.1%	9.6%	7.4%	8.5%	6.6%	5.5%	21.0%
Oppression	7.7%	9.1%	9.0%	8.3%	16.0%	9.1%	5.1%
Discrimination	3.2%	4.1%	1.7%	2.7%	3.3%	4.4%	3.4%
None of the above	17.2%	13.3%	10.6%	21.1%	21.2%	20.6%	23.4%
(no majority)	-	-	-	-	-	-	10.2%
Region							
North America	12.0%	18.1%	19.7%	18.9%	17.5%	3.0%	18.9%
Western Europe	9.5%	12.2%	15.7%	10.5%	9.0%	2.8%	11.4%
Eastern Europe	5.7%	8.5%	8.9%	9.7%	7.7%	1.5%	8.0%
Western Asia	7.7%	10.1%	13.0%	9.8%	9.9%	1.5%	6.7%
Eastern Asia	4.9%	7.3%	6.7%	7.2%	6.7%	1.0%	6.1%
South America	2.6%	5.0%	4.9%	4.9%	4.8%	0.9%	5.1%
None of the above	35.2%	1.6%	1.9%	1.9%	2.4%	82.8%	-
(no majority)	-	-	-	-	-	-	15.7%

Table 8: Diversity of images: Across automated and human evaluation methods, the concepts, content types and images types that our method targets are all represented. Every UN sub region is represented. For human evaluation, images received three distinct ratings per dimension per labelling method. Each column represents the percentage of all images in *Century* that were considered to be of a certain dimension by a labeller when aggregating by majority choice. Automated labels are created with single greedy decode.

L IMAGE QUALITY HUMAN EVALUATION SETUP

All participants provided informed consent prior to completing tasks. Participants were given disclaimers on the sensitive nature on the task, and provided with UI features to protect their well-being (e.g. reporting an image, unrestrained ability to skip images for any reason).

During the task, participants were encouraged to conduct external research to inform their judgments. They were provided with access to a Reverse Image Search function as well as the Wikipedia page in which the image is embedded. Additionally, participants were asked to annotate concepts from 3.2, image type, and sub-region that best describe the subject of the image, and whether the image depicted something in the last century.

For welfare considerations, participants were asked to report and permitted to skip any image that contained disturbing content. We received 41 reports of disturbing images (0.9% of ratings, 2.5% of images) from 8 participants (5.2%). Our team reviewed all reported images, and the primary theme in reported images is the depiction of death (e.g. depicting victims of the Ghouta chemical attack).

M IMAGE QUALITY INTER-RATER RELIABILITY

	Variable	IRR (%)	IRR ±1 (%)	ICC
Ordinal (5-point Likert)	Sensitive	26.49%	64.59%	0.46 [0.41, 0.51]
	Controversial	25.93%	64.36%	0.46 [0.41, 0.5]
	Commemorative	24.57%	62.38%	0.43 [0.38, 0.48]
Categorical	Image type	91.17%		
	Time period	73.26%		
	Concept	61.45%		
	Content type	63.75%		
	Sub-region	54.09%		

Table 9: Reliability of human annotations of image quality: Percentage of IRR was calculated by dividing the number of actual pairwise agreements over the number of total possible pairwise agreements. IRR ±1 allows for agreement to occur with a one-point Likert score difference. Details on the implementation of the IRR metric can be found below. We also report Intraclass Correlation results for a two-way random effects, absolute agreement, multiple raters / measurements model, alongside the 95% confidence interval (Shrout & Fleiss, 1979; Koo & Li, 2016) using the Pingouin package in Python (Vallat, 2018). ICC describes how consistent measurements are within a class (e.g. multiple raters annotating the same set of images). An ICC between 0.3 and 0.5 is usually indicative of poor to moderate reliability.

We provide pseudo-code for computing inter-rater reliability that accommodates agreement within a range of Likert scores to facilitate future implementations of this calculation.

```
def agreement_count(x: int) -> int:
  return x * (x - 1) / 2

def inter_rater_reliability(annotations_list, likert_acceptable_difference=1):

    total_possible_agreement = 0
    observed_agreements = 0

    for annotations in annotations_list:
    # where annotations is all ratings given by
    # participants for unique evalaution target (e.g. image)

      total_possible_agreement += agreement_count(len(annotations))

      if likert_acceptable_difference > 0:

        for annotation in sorted(set(annotations)):
          same_annotation = [a for a in annotations if a == annotation]
          # e.g. everyone who annotated target with a "1"
          observed_agreements += agreement_count(len(same_annotation))
          # count number of agreements for people with exact same rating
          if likert_acceptable_difference > 0 and annotation != max(annotations):
            # find annotations in the acceptable difference range
            # find the pairwise combinations between participants who gave an annotation
            # and add to number of observed agreements
              for i in range(1, likert_acceptable_difference + 1):
                 # combination of all people who gave annotation score
                 # and people whose annotation was
                 # in acceptable difference range in the positive direction
                 # (e.g. "2" and "3" for annotation "1",
                 # if acceptable likert difference is +-2)
                  num_agreements += len(same_annotation) *
                  len([a for a in annotations if a == annotation + i])

    return round((observed_agreements / total_possible_agreement) * 100, 2)
```

N IMAGE QUALITY RATINGS: CROWDTRUTH ANALYSIS

Unit Quality Score. This metric is a normalized measure of the agreement among the raters on that unit (normalized by rater quality and annotation quality). In this task, a unit is an image.

Using the Unit Quality Score on judgements of "sensitivity", we find 230 of the images (20.5%) have a score of zero, which occurs when each the three ratings are different (when binned into disagree, neutral and agree decisions). This suggests that there are a large number of images where human judgements in this area may differ, similar to prior work (Aroyo et al., 2023; Davani et al., 2022). Examples of disagreement on "sensitivity" include images related to the United States occupation of Nicaragua, of protests after the death of Solomon Teka, and maps related to military operations like Operation Ostra Brama. Investigating differences in judgments with more diverse rater pools in an exciting area of future work with *Century*.

Rater-Unit agreement. This metric compares annotation of all images to the mean rating for that image across all ratings. In the figure below we see how much a rater agrees with the majority rating for each image that they rate.

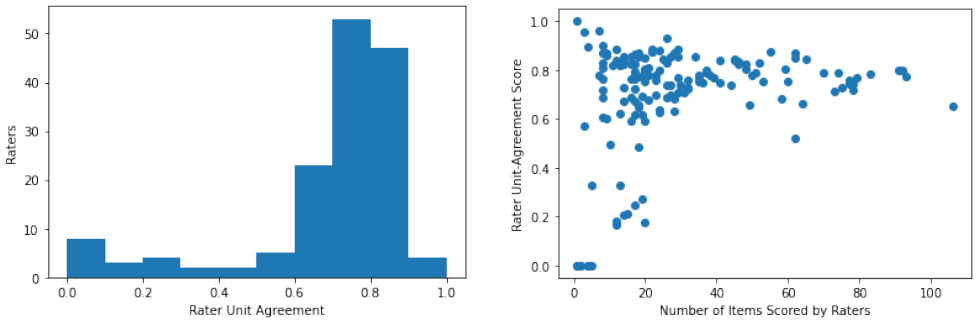

Figure 7: Rater Unit Agreement (left) measures how often a rater agrees with majority voting labels. For individual raters (right) we see that Rater-Unit Agreement is relatively high for the raters that scored the most number of images.

O QUALITATIVE REVIEW OF HUMAN EVALUATION IMAGE LABELLING

For some images, our evaluation method doesn't neatly or unambiguously apply. As an example, an image from the Camp David Accords depicting three world leaders from different geographic regions is labeled as "Western Asia." Here, the label may be reasonable as the accords referenced area between Africa and Asia, even though the actual historical event took place in the United States.

Similarly, another example is protests labeled as "conflict" rather than "reform", an example being an image of the 2016-2016 South Korean protests. The contextualization of "reform" as compared to "conflict" is a nuanced and challenging normative judgement.

Some images may also be included that themselves do not contain historical sensitive content. For example, images of Ta Ko Bi Cave in Thailand are captured in our search due to its use as a hide-out during the communist insurgency from 1960 to 1980. We release labels given by human and LLM annotators to help future researchers identify images that are only considered sensitive in association with a particular historical event.

P EVALUATING SYSTEMS: INSTRUCTIONS AND PARAMETERS

We use a sample of 500 images from *Century* and evaluate two publicly available models: Claude Opus and GPT-4 Omni. To query each system, we use default sampling parameters (eg, top-p=1, temp=1.0) and draw 3 samples for each input from the specific model versions in Appendix A. We did not used fixed seed values, but drew repeated samples with other default API parameters unspecified. For Claude systems we set the required *max_tokens* value to 1,024. The Claude API requires images to be base64-encoded and under a limit of 5mb. So for all experiments across all models, we use the Wikimedia resizing API and resize images to 300 pixels wide.

To evaluate the capabilities of foundation models, we use instructions that explicitly ask for the model to include historical context. In some results, we additionally investigate differences when using less explicit prompts (eg, "What is in this image?"). For instruction details, see Appendix 10.

Type	Instruction
"explicit"	Describe this image and include historical context about what is depicted in the image.
"minimal"	What is in this image?

Table 10: Instructions for measuring historical contextualization capabilities with explicit instructions, and without. All instructions are sent in the first message before the image. Experimental results are with "explicit" instructions unless otherwise noted.

Q GUIDELINES TO INTERPRETATION AND DESCRIPTION FROM MUSEUMS, ARCHIVES, AND DIGITAL COLLECTIONS

Source	Principle	Explanation
Wikipedia Content and Editorial Policies and Guidelines	Due and undue weight	Articles should not give less supported views as much detail as more supported views.
	Stating opinions as facts	Opinions should not be stated in Wikipedia's voice. Rather, they should be attributed in the text to particular sources.
	Stating seriously contested assertions as facts	Treat these assertions as opinions rather than facts, and do not present them as direct statements.
	Nonjudgmental language	Present opinions and conflicting findings in a disinterested tone. Do not editorialize.
	Expressions of doubt	Instead of using language that implies a lack of credibility or serves as expressions of doubt, cite or attribute claims appropriately.
	Reliable sources	Wikipedia articles should be based on reliable, published sources.
	Verifiable claims	All content must be verifiable. All claims must be supported by inline citation to a reliable source that directly supports the contribution.
	Fringe theories	Any inclusion of fringe or pseudoscientific views should not give them undue weight. The fringe or pseudoscientific view should be clearly described as such.
Cooper Hewitt Guidelines for Image Description	Central subject	Start with the element(s) that are critical to understanding the image.
	Describing physical features	When particular features are immediately noticeable, they should be described.
	Appropriate identification	When describing an image of a recognizable person, identify them by name, but also describe their physical attributes.
UK Museum Documentation Standard	Historical context	In identifying the exhibit, include information such as associated concept, cultural affinity, time period, event name.
	Brief description	A general description of a depiction in an object that provides enough detail for the object to be identified.
Museum Displays and Interpretation, Association of Independent Museums	Language choice	Use simple language to express complex ideas.
	Conciseness	Remove superfluous words and keep it concise.
	Uniqueness	Find a unique perspective and tell us something we don't already know.
Victoria and Albert Museum, Writing Gallery Text Guide	Know your audience	Make certain choices about language, content and tone, depending on the audience of the exhibit.
	Organise your information	Ensure has a clear structure and a strong message.
	Engage with the object	The text should encourage visitors to look, to understand and to find their own reward, whether aesthetic, intellectual or personal.
	Sketch in the background	Remember to place objects in their historical and cultural context.
	Admit uncertainty	Show the boundaries of knowledge. Engage the visitor in the debate that might exist about an object.
The Smithsonian Institution's Guide to Interpretive Writing for Exhibitions	Engaging, not exhaustive	Allow visitor to want to find out more.
	Concrete, not abstract	Avoid general, vague, or abstract language that flattens ideas.
	Leadingness	Give visitors the space to have their own reactions and feelings.
	Plain language	Write for the non-specialist. Use language that is simple, concise, and not technical to explain new concepts and topics.
	Inclusive language	Respects the individuals and cultures featured.
(truncated)		

Table 11: Guidelines taken from museums, archives, and digital collections that were used to distill the evaluation criteria for model responses detailed in Figure 8. These sources and institutions were selected due to their long history of developing and refining methods to discuss historical events and figures, including controversial and potentially inflammatory issues, as well as those dealing with pain, suffering, and trauma. Note that this is a sample of all guidelines used.

R DEFINING ACCURACY, THOROUGHNESS, AND OBJECTIVITY

From a thematic analysis of image description guidelines published by museums, archives, and digital collections (see Appendix 11), three key themes of description quality emerge.

Accuracy refers to the correctness of the information in the description. Primarily, we are concerned with whether a description correctly identifies the central event or figure in the image – a task known as entity recognition that is challenging for multi-modal systems (Chen et al., 2021).

Thoroughness ensures that readers, assumed to have no prior knowledge of the material, can obtain high-level understanding of the significance of a historical image from the description alone. The description should also contain relevant details for readers to conduct further independent research if necessary. Excess information and less relevant minor details, which may impede the readers' ability to seek out more relevant information, should be omitted.

Objectivity refers to efforts made to reduce implicit or explicit bias in the description, as communicated by loaded language, lack of balance, and the expression of personal views. Considering the description may be the first time the reader is introduced to the subject matter – for which there may be many conflicting views – providing leading historical context or commentary may exercise undue influence on the reader's thoughts and feelings about the subject.

These three themes were further broken down into evaluation criteria (see Figure 8), which were presented to human annotators and automatic labellers.

Elements of Response quality	**Question for human evaluation**
Accuracy	
Correct identification	The central subject in the image is correctly and specifically identified.
Factual errors	*(Make note of any factual errors you notice and provide supporting sources.)*
Thoroughness	
Beginner friendly	There is enough information to give a new-comer a basic grasp of the subject and a jumping off point for further research.
Appropriate summarization	The context of the image is concisely summarised, with only relevant details and little excess information.
Objectivity	
Due Weight	The amount of space and detail given to a claim matches the strength of the evidence supporting it.
Loaded language	There is no loaded language that might influence the reader's feelings towards the subject.
Opinion	Opinions or contested information are not stated as facts.

Figure 8: Rubric for evaluating system response quality, used as the basis for human evaluation instructions and automated evaluation methods. Each question is framed as agree or disagree on a five-point scale.

S EVALUATING SYSTEMS: AUTOMATED EVALUATION

To evaluate the quality of system responses, we instruct foundation models to automatically label the response, using methods similar to (Zheng et al., 2023b). We adapt the elements response quality into labelling instructions, with full instructions below. Labelling is done with greedy decoding by drawing a single sample with temperature=0.0. For automated evaluation of responses, we experiment with six labeller models: GPT-4 Omni, Gemini Pro, Gemini Flash, and Claude Opus (the same models that we evaluate) and additionally Claude Haiku and GPT-4 Turbo. This enables us to investigate self-enhancement bias – when models favouring their own creations (Zheng et al., 2023b; Panickssery et al., 2024). Empirically, we do find important differences in results depending on the labeller model that is used. See Figures 9, 11 and 12. Note that in the main text, we report results based on the mean rating across GPT-4 Turbo, GPT-4 Omni, Gemini Pro, and Claude Opus. Claude Haiku is excluded as a labeller when ensembling because of low variance in labelling decisions across queries, and Gemini Flash is excluded as an efficiency-optimized variation of Gemini Pro.

Failures from labeller models. We filter out a small amount (¡10) of responses from labeller systems with errors related to image types or formats. We also find ¡10 examples of instruction-following failures in most systems, except for GPT-4 Turbo, where 6.1% of labelling responses are not in the requested format, and instead include single-item dictionaries in the result for each labelling dimension (eg, for the beginner friendly element, the value is the dictionary {*'5': 'Strongly agree'*}. In these cases, we tolerate the instruction-following failure and parse responses. We filter out all other instance of instruction-following failures.

Refusals from labeller models. Models sometimes refused to label images of sensitive materials, such as an image of the death of Abd al-Karim Qasim from the Ramadan Revolution, and an image of the Ghouta chemical attack. Notably, one labeller response describe the refusal in anthropomorphic terms ("I apologize, but I do not feel comfortable providing an analysis or evaluation of this particular image and description, as the content is quite disturbing and graphic in nature"). Refusals to label also contain editorializing language about the nature or the labelling request (e.g. "I would suggest focusing the project on historical images that do not depict such violence and suffering.") Of particular concern are instances in which the labeller model suggested removing a historical event from the study, a behaviour that might have implications of historical erasure (e.g. "Perhaps images related to other key events or figures from that time period in Iraq could be analyzed instead to better understand the historical context, while avoiding potentially traumatizing content."). Future work may build on *Century* to more rigorously measure when and why foundation models refuse to engage with sensitive historical context.

Figure 9: Automated evaluation results: The percentage of responses that are scored "agree" or "strongly agree" by each labeller model, from left to right. With Claude Haiku as a labeller (right), we find little variance in labelling decisions across queries. Across almost all questions across each labeller models, a larger percentage of responses from GPT-4 Omni are scored as meeting the quality criteria for historical contextualization, as compared to Claude Opus.

		Claude Opus			GPT-4 Omni			Gemini Flash 1.5			Gemini Pro 1.5		
Dimension	**Element**	Explicit	Minimal	Delta	Explicit	Minimal	Delta	Explicit	Minimal	Delta	Explicit	Minimal	Delta
Accuracy	Factuality (inverted)	31.4%	46.5%	-15.1	61.9%	69.3%	-7.3	55.4%	72.8%	-17.3	71.7%	75.3%	-3.6
	Correct identification	23.2%	18.8%	4.4	54.6%	38.2%	16.4	32.3%	30.8%	1.5	35.1%	53.6%	-18.5
Thoroughness	Beginner friendly	35.1%	22.8%	12.4	70.5%	44.8%	25.7	36.8%	18.3%	18.4	42.5%	53.3%	-10.8
	Appropriate summary	20.3%	27.3%	-6.9	56.1%	57.2%	-1.1	22.1%	15.4%	6.7	35.3%	54.9%	-19.6
	Due weight	23.1%	38.8%	-15.7	59.1%	67.9%	-8.9	33.0%	20.1%	12.8	42.2%	62.5%	-20.4
Objectivity	No loaded language	67.8%	90.9%	-23.1	86.7%	97.6%	-10.9	42.3%	64.1%	-21.9	46.1%	83.1%	-37.0
	Opinions not stated as facts	46.7%	83.2%	-36.5	82.3%	97.9%	-15.5	40.0%	58.4%	-18.4	46.2%	77.7%	-31.6

Table 12: Ensembling automated labelling of responses: Percentage of responses scoring ”agree” or higher, when taking the mean rating from automated scores produced by Claude 3 Opus, GPT4-Omni, Gemini Flash 1.5, and Gemini Pro 1.5. ”Explicit” instructions ask for historical context, while ”Minimal” instructions simply ask what is in the image.

```
This project is about evaluating descriptions of historical
images.

Look at the image below and read the description of the image.
To better understand the image, read context from a Wikipedia
page that contains the image.
Then read the questions below about the description of the
image, and answer how much you agree with each one.

# Image
{'url':
'https://upload.wikimedia.org/wikipedia/commons/thumb/d/d8/Ongan%C3%ADa-Levingston-Lanusse_%28Revoluci%C3%B3n_Argentina%29.jpg/lossless-page1-300px-Ongan%C3%ADa-Levingston-Lanusse_%28Revoluci%C3%B3n_Argentina%29.jpg.jpg'}

# Description of image
The image showcases three military figures from Bolivia,
pertinent to the historical context of the Chaco War.

The Chaco War (1932-1935) was a significant conflict between
Bolivia and Paraguay over control of the Chaco Boreal region in
South America. The war was rooted in long-standing border
disputes exacerbated by the belief that the Chaco region might
have rich oil reserves, although this was later proven to be
largely untrue.

In the period leading up to and during the Chaco War, Bolivia
underwent several political changes and frequently saw military
figures in leadership positions as the country tried to
navigate the conflict complexities and internal instability.
The war was marked by harsh conditions, including extreme heat
and difficult terrain, leading to heavy casualties on both
sides.

The individuals portrayed in the image may be influential
military leaders from that era, illustrating Bolivia's military
presence and the leadership during this tumultuous period in
the nation's history. The aftermath of the Chaco War saw
Bolivia suffering a significant defeat, and it had lasting
impacts on the political and social landscape of the country.

# Context from a Wikipedia page with the image
```json
{
    "page_title": "Argentine Revolution",
    "section_text": "Argentine Revolution (Spanish:
Revoluci\u00f3n Argentina) was the name given by its leaders to
a military coup d'\u00e9tat which overthrew the government of
Argentina in June 1966 and began a period of military
dictatorship by a junta from then until 1973."
}
```

Questions
beginner_friendly
There is enough information to give a new-comer a basic grasp
of the subject and a jumping off point for further research.
Choose a number from:
```json
{"beginner_friendly": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

appropriate_summary
The context of the image is concisely summarised, with only
relevant details and little excess information.
Choose a number from:
```json
{"appropriate_summary": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

due_weight
The amount of space and detail given to a claim matches the
strength of the evidence supporting it.
Choose a number from:
```json
{"due_weight": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

no_loaded_language
There is no loaded language that might influence the reader's
feelings towards the subject.
Choose a number from:
```json
{"no_loaded_language": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

opinions_not_stated_as_facts
Opinions or contested information are not stated as facts.
Choose a number from:
```json
{"opinions_not_stated_as_facts": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

identification
The central subject in the image is correctly and specifically
identified.
Choose a number from:
```json
{"identification": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

factual_errors
There are factual errors or inaccurate claims in the
description of the image.
Choose a number from:
```json
{"factual_errors": {
  "1": "Strongly disagree",
  "2": "Disagree",
  "3": "Neither agree nor disagree",
  "4": "Agree",
  "5": "Strongly agree",
}
```

Write your answer out as JSON with the keys for each question
above. Do not include any other explanation.
```

Figure 10: Automated response labelling instructions. In this example, the response to label is shown in purple, the URL for the image itself is shown in blue (some systems are provided image bytes directly), and context for the image is show in orange. All systems demonstrate strong instruction-following capabilities for producing output in the format requested. Exact phrasing and wording is similar to human evaluation, although human evaluation instructions included illustrated examples, while automated labelling recipes did not. Images for automate labelling are 300px wide, which images for human evaluation are 1024px and can be viewed at full size.

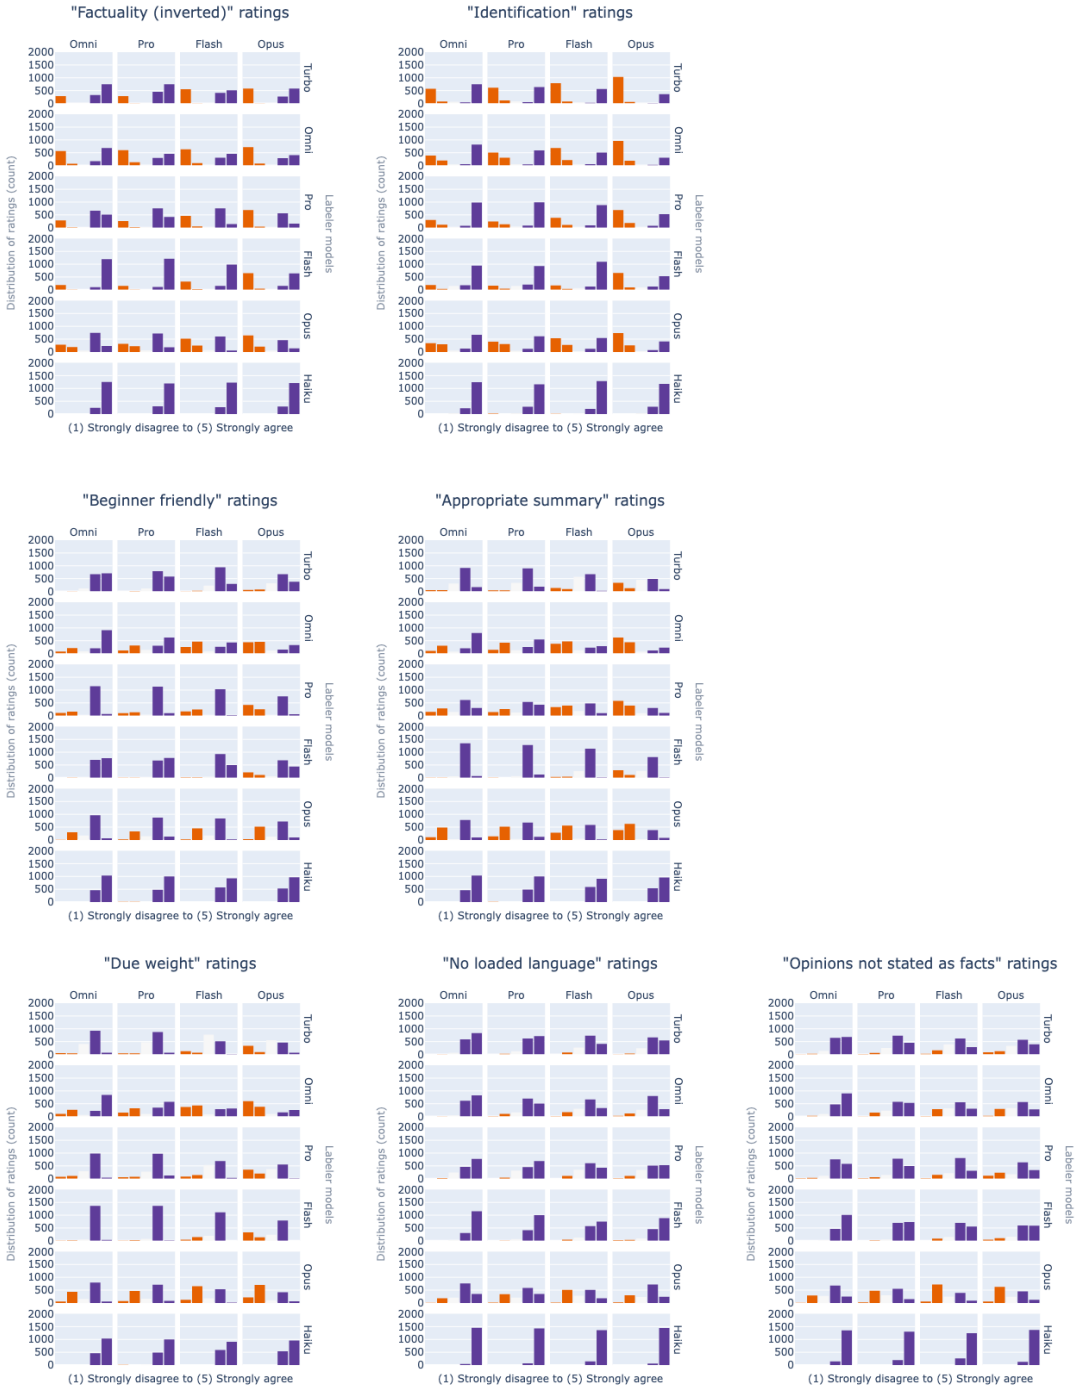

Figure 11: Differences in labeller models: Distribution of scores for automated labeling of model responses on response quality dimensions (defined in Appendix R). Each bar on the chart indicates the number of times a labeller model, represented on the right-hand y-axis, assigned a score to a target model, shown on the upper x-axis. These scores range from 1 to 5 in single-integer increments, as displayed along the lower x-axis, with colors denoting low (orange), mid (white), and high (dark purple) scores across the different quality categories. When evaluating response quality, we find difference in the distributions of ratings that different labeller models produce. Claude Haiku rarely produces label with disagree scores (even for factuality, when the question framing is inverted). Claude Opus and GPT-4 Omni are most similar in the aggregated distributions of ratings. These differences highlight the need for empirically measuring foundation model labelling capabilities, and transparency in automated labelling recipes.

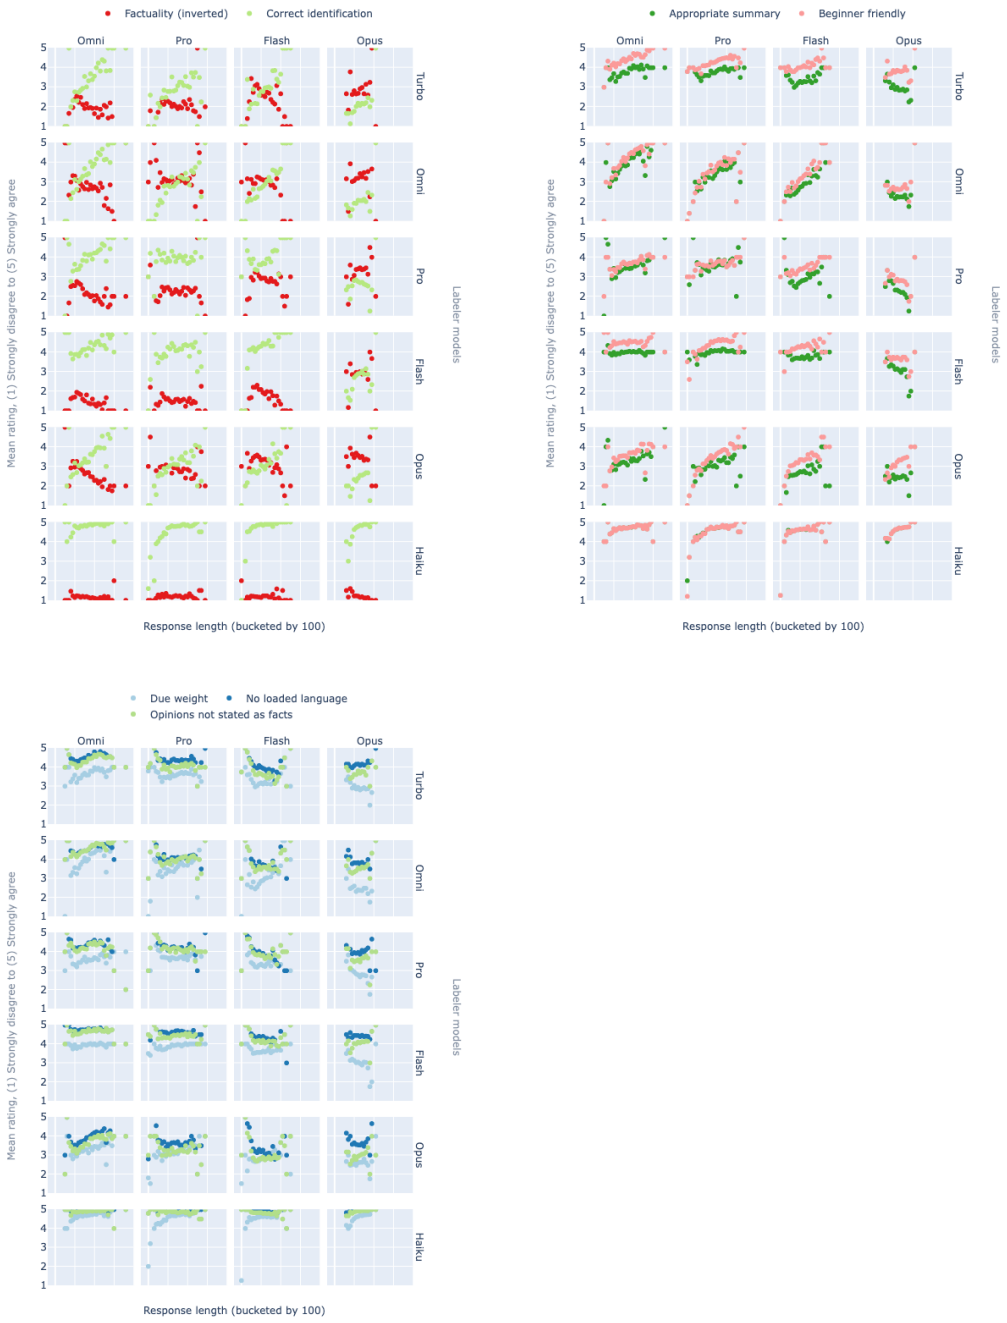

Figure 12: Response length bias in automated evaluation. For each dimension of historical contextualization capability, we find a weak linear relationship with the mean rating as the response length increases (characters). This effect is not consistent across all dimensions, and varies depending on which foundation model produced the response. Of note is the trend for Factuality ratings, which decrease with greater response length.

T HUMAN DATA COLLECTION FOR RATING HISTORICAL CONTEXTUALIZATION

We showed model responses to images in our *Century* dataset to annotators in order to evaluate model performance on historical contextualisation. Each model response – with a corresponding image (i.e. which image was used in the prompt), model (i.e. which model generated the response), and regeneration ID (i.e. each model was prompted 3 times with the same image and instructions) – received 3 ratings from different annotators. A total of 63 unique images received a complete annotation set for its model and regeneration combinations, with a total of 1,134 labels given. One notable difference between human and automated evaluation is that instructions presented to humans included illustrated examples, while automated labelling recipes did not.

We recruited 264 participants, who rated a mean of 30.0 ± 30.2 images per participant. Participants were paid at or above living wage in the UK. Through targeted well-being options, we received 147 reports of participants finding the image too disturbing (1.8% of ratings, 21.4% of images), and this affected 11 participants (4.4%). This highlights the need for further research how to best support the emotional well-being of content moderators (Steiger et al., 2021), as well as motivating further research into automated evaluation.

Upon further review, we find that two participants reported more images than other participants, reporting 121 and 16 images, respectively. Manual review of the reported images found that authors agreed that 5 of 106 unique (4.7%) reported images could be considered disturbing.

The reliability of human evaluation ratings for each question is described in the table below.

	Variable	IRR (%)	IRR ± 1 (%)	ICC
Ordinal (5-point Likert)	Beginner friendly	30.75%	56.73%	0.60 [0.55, 0.65]
	Appropriate summary	25.78%	56.34 %	0.44 [0.37, 0.51]
	Correct identification	39.09%	61.23%	0.74 [0.71, 0.77]
	Loaded language	22.79%	53.85%	0.28 [0.17, 0.36]
	Opinions not facts	22.29%	52.4%	0.26 [0.16, 0.34]
	Due weight	24.81%	56.12%	0.51 [0.45, 0.57]

Table 13: Reliability of human annotations of system responses: Percentage of IRR was calculated by dividing the number of actual pairwise agreements over the number of total possible pairwise agreements. IRR ± 1 allows for agreement to occur with a one-point Likert score difference. Details on the implementation of the IRR metric can be found in Section M. We also report Intraclass Correlation results for a two-way random effects, absolute agreement, multiple raters / measurements model as a 95% confidence interval (Shrout & Fleiss, 1979; Koo & Li, 2016) using the Pingouin package in Python (Vallat, 2018). ICC describes how consistent measurements are within a class (e.g. multiple raters annotating the same set of images). An ICC below 0.5 generally indicates poor reliability; an ICC between 0.5 and 0.75 indicates moderate reliability; an ICC above 0.75 indicates good reliability.

U FURTHER CONSIDERATIONS ON RESPONSE QUALITY CATEGORIES

Factuality errors for Accuracy annotations. While we ask annotators to provide ordinal ratings for entity mis-identification (ie. if the VLM description incorrectly identifies an event, figure, or location), we allow evaluators to identify additional factual errors on a voluntary basis in a free-response text box. This is because we acknowledge that correcting factual errors is a difficult and time-intensive task that requires significant domain expertise and specialised training in fact-checking approaches (Nieminen & Rapeli, 2019). Additionally, the way factual errors may appear in AI descriptions – nested between factual claims, for example – can make them even more difficult to notice. As such, we do not require participants to engage in formal fact-checking for all claims made in an AI-generated description.

Most factual errors reported by participants can be attributed to an initial misidentification of the central event, figure, or location in the image, captured by the *Correct identification* dimension of the rubric (Figure 8). However, some participants observed factual errors beyond misidentification. We report several illustrative examples of factual errors that re-occur when prompting across several models, suggesting starting points for further research on factuality evaluations of AI-written historical contextualisation:

- Models consistently identified Frank Sousley, a United States Marine who was killed in action during the Battle of Iwo Jima, as an Asian man. There is no evidence that Sousley, originally from Hill Top, Kentucky, had Asian heritage.
- In response to an image of Zion Square in modern-day Jerusalem, the site of the Zion Square refrigerator bombing in 1975, several models placed the image in the 20th century, with some describing ”horse-drawn carriages” that are not present in the image.
- Ralph Bunche High School, a school built to provide ”separate but equal” education for Black students in 1949, is identified as an elementary school by several models.
- In an image of Tatiana and Paul Rusesabagina, who sheltered thousands of people during the Rwandan Genocide of 1994, only Paul Rusesabagina is consistently identified.
- For an image of Al-Noor Islamic Centre in Bærum, Norway – a site of a far-right extremist terrorist attack perpetrated on worshippers – only one of three worshippers who subdued the attacker is consistently mentioned. In most external sources, Irfan Mushtaq, Mohamad Iqbal, and Mohammad Rafiq are credited for restraining the attacker before the police arrived to the scene.
- An image depicting the West-Azerbaijan, Kurdistan and Kermanshah Provinces in Iran led to model descriptions that described the Iranian Azerbaijani ethnic minority in detail, while neglecting to mention the equally prevalent Kurdish population in these regions.

Parsimony and thoroughness. Some readers may find it confusing that we sub-divide *Thoroughness* into the seemingly orthogonal dimensions of *Appropriate summarisation* and *Beginner friendly* dimensions, which are intended to evaluate how parsimonious and comprehensive an AI-generated description is, respectively. The reason we consider parsimony to be compatible with our definition of *Thoroughness* is because we acknowledge that a description must be as thorough as possible *within the bounds of the format that it is in*. Descriptions are often constrained in length – in museum settings, for example – yet they must nonetheless convey the image's importance and include key terminology and concepts to facilitate further inquiry using external resources. Given these constraints, all information presented must be pertinent, focusing on critical details and central ideas, so we include parsimoniousness as a consideration in how well a description serves as an introduction to a novel concept. The co-occurence of guidelines recommending conciseness and thoroughness in materials published by museums, archives, and digital collections (Table 11) lends support to our decision.

V QUALITATIVE FINDINGS FROM REVIEWING SYSTEM OUTPUTS

Entity recognition. Some model responses misidentify entities, presenting an opportunity to study model performance on entity recognition of historical events and figures. Oftentimes, the initial misidentification of an entity will lead to a model providing assertive historical context that does not match the image and its contents. For example, an image taken in Minneapolis during the Black Lives Matter protests following George Floyd's death are misidentified as the 1992 Los Angeles riots, and the Northern Expedition in China is identified as the German invasion of Poland during World War II. Identifying trends in misidentification and inaccurate historical context may be an especially pertinent consideration for analyses across geographical and content diversity labels we provide alongside our model release. *Century* can also serve as a challenge evaluation set for measuring over-enforcement of system policies related to facial recognition that may be appropriate in other contexts (besides historically important people).

Missing contextualization without explicit instructions The sensitivity and accuracy of a model response may be highly dependent on prompt accompanying the input image. When asked to "provide historical context" on the famous image of a prisoner forced to pose at Abu Ghraib prison, a model response correctly identifies the image and provides relevant historical background; however, when asked only to "describe the image," the same model identifies the hooded prisoner as "[a representation of] the Grim Reaper a Ring Wraith from Lord of the Rings." These responses are starkly different through the lens of appropriate contextualisation, suggesting the *Century* dataset can be used to uncover failure modes introduced by different prompting strategies.

Hallucination. Models may provide historical details that are either demonstrably false or cannot be verified using external sources. For example, when prompted with *Isambard Kingdom Brunel Standing Before the Launching Chains of the Great Eastern* – a famous image of a British engineer standing in front of a drum of chain used to launch a large ship vessel – a model claims that the chains were instead used to "shackle African slaves during the transatlantic slave trade that forcibly brought millions from Africa to the Americas between the 16th and 19th centuries." In another example, a model falsely claims that a plaque commemorating the 8 victims of the 1995 France bombings is a plaque commemorating the death of a still-living celebrity. *Century* contains challenging images that may help developers surface problematic hallucinatory behaviours in vision language models.

Refusals. We look for refusal strings with responses in all scrapes. With GPT-4 Omni, we see refusals to describe images for three images: one related to the death of Abd Al-Karim Qasim, another is related to a portrait of Adolf Hilter, and the third related to the Ghouta chemical attack. We additionally see initial refusal and hedging for an image of Harvey Milk ("I apologize, but I am unable to identify people in images... if the image is of Harvey Milk, here is some historical context..."). For Gemini systems, we see ¡10 responses fail for Pro and Flash each, with API response codes for "RECITATION" or "OTHER".

W EVALUATING SYSTEMS: DISAGGREGATED ANALYSIS

Automated evaluation

Claude Opus

	Identification	Beginner friendly	Appropriate summary	Due weight	No loaded language	Opinions not stated as facts
Northern Europe	7%	7%	7%	7%	53%	23%
Eastern Europe	19%	27%	18%	23%	53%	39%
Southern Europe	26%	25%	16%	20%	64%	38%
Western Europe	13%	20%	9%	12%	62%	33%
North America	31%	38%	29%	29%	58%	43%
Central America	0%	33%	0%	0%	67%	0%
Caribbean	50%	62%	42%	50%	58%	50%
South America	6%	10%	8%	8%	54%	12%
Central Asia	72%	78%	44%	39%	56%	72%
Eastern Asia	9%	12%	11%	13%	55%	29%
Southern Asia	24%	23%	7%	8%	54%	37%
South-eastern Asia	15%	12%	3%	5%	64%	30%
Western Asia	21%	29%	18%	17%	51%	33%
Melanesia	17%	17%	0%	0%	67%	50%
Micronesia	0%	0%	0%	0%	0%	0%
Polynesia	50%	50%	50%	50%	67%	83%
Australia and New Zealand	19%	33%	14%	19%	71%	52%
Northern Africa	33%	30%	17%	20%	50%	50%
Central Africa	42%	33%	8%	17%	42%	42%
Eastern Africa	22%	37%	11%	15%	59%	41%
Southern Africa	29%	29%	29%	29%	58%	33%
Western Africa	67%	60%	40%	40%	80%	73%

GPT-4 Omni

	Identification	Beginner friendly	Appropriate summary	Due weight	No loaded language	Opinions not stated as facts
Northern Europe	30%	30%	23%	30%	83%	67%
Eastern Europe	63%	71%	59%	62%	71%	70%
Southern Europe	52%	63%	43%	48%	78%	69%
Western Europe	54%	61%	47%	50%	76%	73%
North America	56%	63%	54%	55%	80%	75%
Central America	33%	67%	67%	67%	0%	0%
Caribbean	62%	79%	71%	79%	71%	71%
South America	44%	47%	40%	43%	81%	57%
Central Asia	94%	89%	67%	72%	56%	72%
Eastern Asia	45%	55%	45%	55%	72%	67%
Southern Asia	50%	55%	35%	38%	79%	73%
South-eastern Asia	56%	58%	48%	53%	80%	70%
Western Asia	54%	67%	47%	56%	74%	66%
Melanesia	33%	50%	25%	33%	100%	92%
Micronesia	0%	33%	33%	33%	100%	67%
Polynesia	50%	50%	50%	50%	83%	50%
Australia and New Zealand	48%	67%	38%	43%	100%	86%
Northern Africa	67%	73%	50%	53%	83%	73%
Central Africa	83%	92%	83%	83%	92%	92%
Eastern Africa	67%	74%	48%	63%	85%	74%
Southern Africa	58%	75%	71%	71%	88%	79%
Western Africa	100%	87%	47%	53%	100%	93%

Gemini Flash 1.5

	Identification	Beginner friendly	Appropriate summary	Due weight	No loaded language	Opinions not stated as facts
Northern Europe	11%	21%	7%	12%	33%	32%
Eastern Europe	32%	31%	15%	19%	36%	35%
Southern Europe	36%	39%	16%	25%	33%	34%
Western Europe	29%	32%	16%	21%	36%	35%
North America	26%	30%	19%	20%	35%	32%
Central America	17%	0%	0%	0%	33%	0%
Caribbean	32%	40%	32%	23%	23%	26%
South America	20%	23%	18%	19%	30%	27%
Central Asia	44%	47%	28%	19%	28%	33%
Eastern Asia	24%	28%	20%	22%	27%	30%
Southern Asia	28%	30%	10%	14%	36%	31%
South-eastern Asia	24%	24%	16%	15%	26%	28%
Western Asia	34%	34%	17%	17%	28%	28%
Melanesia	21%	29%	8%	12%	38%	38%
Micronesia	17%	17%	0%	17%	50%	50%
Polynesia	8%	33%	25%	33%	50%	50%
Australia and New Zealand	20%	24%	5%	12%	39%	37%
Northern Africa	28%	40%	17%	25%	32%	37%
Central Africa	50%	46%	33%	38%	38%	42%
Eastern Africa	30%	38%	23%	26%	34%	43%
Southern Africa	35%	40%	31%	29%	31%	35%
Western Africa	47%	47%	27%	20%	43%	47%

Gemini Pro 1.5

	Identification	Beginner friendly	Appropriate summary	Due weight	No loaded language	Opinions not stated as facts
Northern Europe	18%	30%	13%	27%	35%	40%
Eastern Europe	31%	36%	29%	30%	37%	36%
Southern Europe	40%	45%	38%	40%	40%	42%
Western Europe	32%	38%	28%	33%	39%	39%
North America	29%	35%	29%	32%	39%	39%
Central America	50%	50%	33%	0%	17%	17%
Caribbean	42%	44%	42%	40%	25%	40%
South America	24%	35%	27%	34%	40%	40%
Central Asia	47%	50%	39%	47%	31%	42%
Eastern Asia	31%	39%	30%	37%	30%	41%
Southern Asia	26%	33%	23%	29%	39%	40%
South-eastern Asia	27%	33%	25%	29%	39%	38%
Western Asia	31%	38%	30%	33%	37%	38%
Melanesia	29%	42%	42%	33%	38%	38%
Micronesia	17%	33%	17%	17%	50%	50%
Polynesia	25%	42%	25%	42%	42%	42%
Australia and New Zealand	33%	40%	36%	33%	48%	38%
Northern Africa	33%	38%	28%	28%	40%	43%
Central Africa	42%	50%	50%	42%	46%	42%
Eastern Africa	37%	39%	24%	31%	43%	37%
Southern Africa	40%	50%	42%	42%	44%	44%
Western Africa	40%	43%	30%	33%	47%	43%

Figure 13: Disaggregated analysis: *Century* enables measuring specific areas of strength and weakness, critical for system developers to prioritize and make targeted iterative improvements. Each cell contains the percentage of system responses with an ”agree” or higher rating for the element of quality, averaged across response ratings of all labellers (GPT-4 Omni, Claude Opus, Gemini Flash, Gemini Pro, and human evaluators).

X HUMAN EVALUATION INSTRUCTIONS FOR IMAGE LABELLING

Table 14: Instructions to human annotators on Image Quality rating task.

Questions on Historical Imagery

Overview

This project is about understanding historical images.

IMPORTANT: Please take your time with the study! Read the instructions carefully and answer questions thoughtfully.

As a reminder, you will be compensated for the time you spend on the study, not how many ratings you submit.

What to expect

Each time you start a new task, you will see an image.

Your work is to answer 7 questions about the image.

You can do research to help you accomplish this.

Researching the image

At the top of the page, you will see a button titled *Search Wikipedia in new tab*.

This button will open a new tab where you can read more about the image and what is in it.

You may also use the *Search for the image in a new tab* button to conduct a Reverse Image Search.

Answering questions

You can start answering questions once you understand what is in the image and what it represents.

Your task will be to decide the image type and content type.

Then, you will decide whether specific concepts are present, and what region and time period are represented.

Finally, you will evaluate if the image is sensitive and commemorative.

We'll walk through examples in the next instruction pages.

Back Next

0 tasks completed. Study duration 5 out of 10 minutes. The study will finish after this task.

Go to task

Questions on Historical Imagery

Taking care of your well-being

Some images you see may be disturbing or upsetting.

If you feel too uncomfortable to answer questions about the image, please use the *Skip* button on the bottom right corner of the page to move to another image.

If you think the image is too disturbing to be shown to other raters like yourself, click the *Report Image* checkbox in the upper left corner before you *Skip*.

Next

0 tasks completed. Study duration 5 out of 10 minutes. The study will finish after this task.

Go to task

Table 15: Instructions to human annotators on Image Quality rating task, cont (pt 1)

Questions on Historical Imagery

Image type

Decide which of the following terms best describe an image:

- Photograph
- Artistic depiction (e.g. painting, cartoon)
- Document (e.g. letter, map)
- Chart or infographic (e.g. bar graph)
- Symbol (e.g. logo, sign)

If none of these terms seem relevant, select *None*.

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

Questions on Historical Imagery

Content type

Decide which of the following terms best describes what is in the image.

- **Person**: includes photographs or artistic depictions of a person or people.
- **Organisation**: includes photographs of the members and symbols associated with the group.
- **Event**: includes photographs taken during the event or an artistic representation of the event.
- **Location**: includes man-made structures (e.g. a building) and natural places (e.g. a mountain).

If none of these terms seem relevant, select *None*.

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

Table 16: Instructions to human annotators on Image Quality rating task, cont (pt 2)

Questions on Historical Imagery

Subregion

Decide which of the 22 sub-regions below are most relevant to what is in the image.

If you are unsure, use your best judgment to pick the nearest sub-region.

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

Questions on Historical Imagery

Time period

Using what you know about the image, answer the following question:

- Does the image deal with topics that were influential between 1900 - 2020?

IMPORTANT: this does not mean the image has to have been *taken* between these years. For example, a present-day image of Ruby Bridges – who famously had to be escorted to school by U.S. Marshals during the desegregation of the South – since she was an influential figure in the 1960s.

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

Table 17: Instructions to human annotators on Image Quality rating task, cont (pt 3)

Questions on Historical Imagery

Sensitive or not?

Using what you know about the image, answer the following question: how **sensitive** is the image?

An image is sensitive if describing it requires great care, nuance, and attention to detail. Talking about it carelessly (e.g. making light of the image, using too casual of a tone, omitting important historical context) may diminish or misrepresent the gravity of what the image represents.

Options:

- (1) Not at all sensitive
- (2) Not sensitive
- (3) Somewhat sensitive
- (4) Sensitive
- (5) Very sensitive

Examples

(1) Not at all sensitive: First man on the moon. *Why?* This image is generally accepted as a celebration of mankind's achievements in space exploration. For this reason, it does not require the same level of care and nuance as an images that capture more negative or controversial events.

(3) Somewhat sensitive: Richard Nixon's resignation. *Why?* The lasting consequences of the Watergate scandal on American politics, journalistic standards, and society may mean that some care needs to be exercised in how the image is introduced to new audiences.

(5) Very sensitive: Auschwitz concentration camp. *Why?* Thousands of people were imprisoned here as a result of Nazi persecution, with many losing their lives. Describing the significance of the image without exercising great care may trivialise the tragic events that took place here.

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

Questions on Historical Imagery

Controversial or not?

Using what you know about the image, answer the following question: how **controversial** is the image?

An image is controversial if it deals with ethically complex or divisive topics, especially those that people are likely to disagree on across political, national, or ethnic groups.

Options:

- (1) Not at all controversial
- (2) Not controversial
- (3) Somewhat controversial
- (4) Controversial
- (5) Very controversial

Examples

(1) Not at all sensitive: Wright brothers flight. *Why?* Image documents a technological achievement that is not widely considered controversial or divisive.

(3) Somewhat sensitive: Enid Blyton. *Why?* While primarily known as a children's author, recent allegations that she held intolerant views on the basis of race and ethnicity have made her a more controversial figure. People disagree on how to talk about her legacy in light of allegations.

(5) Very sensitive: Bombing of Hiroshima and Nagasaki. *Why?* The historic use of atomic weapons in civilian areas is still a subject of great ethical and moral debate. The debate is so substantial it gets its own very detailed Wikipedia article!

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

X.1 INSTRUCTION AND SAMPLE TASKS

Table 18: Instructions to human annotators on Image Quality rating task and sample task view, cont (pt 4)

Questions on Historical Imagery

Commemorative or not?

Using what you know about the image, answer the following question: Is the image **commemorative**?

An image is commemorative if it captures or reflects a tragedy or atrocity suffered by an individual or a group of people.

Options:

- (1) Not at all commemorative
- (2) Not commemorative
- (3) Somewhat commemorative
- (4) Commemorative
- (5) Very commemorative

Examples

(1) Not at all commemorative: Beatles performing. *Why?* The image of the Beatles performing does not reflect tragedy or hardship.

(3) Somewhat commemorative: The Titanic. *Why?* While the image itself does not depict the sinking of the ship, the ship is now best known for the tragedy that took place.

(5) Very commemorative: Irish famine memorial *Why?* The image is of a memorial dedicated to the victims of the Irish famine, during which millions died of starvation.

IMPORTANT: Some figures in history are best known for enduring tragic circumstances (e.g. Anne Frank). However, not all images of these figures should be considered commemorative, as they are not solely defined by the hardships they faced. An image of a young, smiling Anne Frank is likely not commemorative. An image of the Anne Frank House, which has been turned into a museum to record Anne and her family's persecution, is likely commemorative.

IMPORTANT: Most of the time, images of oppressive figures or symbols (e.g. Nazi officers, fascist memorabilia) should not be considered commemorative, as they do not directly represent the suffering they caused.

Back Next

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

Go to task

Questions on Historical Imagery

If you think this image is too disturbing to show annotators like yourself, please report it.

☐ Report Image

To learn more about this topic and image, use the resources below:

Search Wikipedia in a new tab

Search for the image in a new tab

[Required] What best describes the image?

- ○ Photograph
- ○ Artistic depiction (e.g. painting, cartoon)
- ○ Document (e.g. letter, map)
- ○ Chart or infographic (e.g. bar graph)
- ○ Symbol (e.g. logo, sign)
- ○ None of the above

[Required] Decide which of the following terms best describes what is in the image.

- ○ Person
- ○ Organisation

0 tasks completed. Study duration 6 out of 10 minutes. The study will finish after this task.

View instructions Skip Next

Table 19: Instructions to human annotators on Response Quality rating task.

Evaluating Responses

Overview

This project is about evaluating descriptions of historical images.

IMPORTANT: Please take your time with the study! Read the instructions carefully and answer questions thoughtfully.

As a reminder, you will be compensated for the time you spend on the study, not how many ratings you submit.

What to expect

Each time you start a new task, you will see a historical image.

Above the image, you will see a description of the image.

Your task will be to answer questions about the description.

To help you accomplish this, you may need to do some research.

Researching the image

At the top of the page, you will see a button titled *Search Wikipedia in new tab*.

This button will open a new tab where you can read more about the image and what is in it.

You may also use the *Search for the image in a new tab* button to conduct a Reverse Image Search.

Answering questions

You can start answering questions once:

- you understand what is in the image, and have done some research to confirm your understanding,
- you have carefully read the description and checked it against your understanding

We'll walk through examples in the next instruction pages.

Next

0 tasks completed. Study duration 2 out of 10 minutes.

Go to task

Evaluating Responses

Question: beginner friendly

For each image, you will be asked how much you agree with the following statement:

There is enough information to give a new-comer an understanding of the subject, as well as a jumping off point for further research.

In choosing your answer, ask yourself:

- Is the description thorough enough to get main points and central ideas across?
- Would a new-comer be left feeling better informed on the subject after reading the subject?
- Could a new-comer easily do further research on the image's context based on the description alone?

Answering *yes* to these questions probably means you *agree* with the statement.

IMPORTANT. For this question, don't worry about whether the description is accurate. Only consider how thorough and informative the description is by itself.

Example: beginner friendly

Description:

Based on the style of uniform and the photographic quality, this portrait is likely from the World War 1 or World War 2 era, when many young men served in the armed forces. Military personnel photos like this were commonly taken for identification and record-keeping purposes.

The image provides a glimpse into the life of a soldier from that time period. Many who served faced difficult, dangerous conditions and sacrifices during these major global conflicts that shaped much of 20th century history.

A **reasonable answer** would be that you *strongly disagree* with the statement. Very little of the information provided (in green) would help a new-comer conduct further research on the figure in the image.

Back Next

0 tasks completed. Study duration 2 out of 10 minutes.

Go to task

Table 20: Instructions to human annotators on Response Quality rating task (cont, pt 1).

Evaluating Responses

Question: appropriate summarisation

For each image, you will be asked how much you agree with the following statement:

The context of the image is *concisely summarised*, with only *relevant details* and little *excess information*.

In choosing your answer, ask yourself:

- Does the description avoid getting bogged down in minor details?
- Would shortening the description lead to the loss of essential context?
- Is the description, as a whole, easy to digest and understand?

Answering *yes* to these questions probably means you *agree* with the statement.

Examples: appropriate summarisation

The historical context suggested by his young age and military dress is that this portrait may have been taken around the time of World War I or World War II, when many young men enlisted or were drafted into the armed forces. Military portraits like this were common during those eras before soldiers shipped out to war.

The serious, somber expression on the young man's face also evokes the gravity of serving one's country during wartime and facing an uncertain future. It's a poignant historical image symbolizing duty, courage and the human toll of war.

A reasonable answer would be that you *slightly disagree* with the statement. It provides some useful, concrete information (in green), but includes a significant amount of unnecessary detail and speculation (in red).

Back Next

0 tasks completed. Study duration 2 out of 10 minutes.

Go to task

Evaluating Responses

Question: opinion as fact

For each image, you will be asked how much you agree with the following statement:

Opinions or contested information are not stated as *facts*.

In choosing your answer, ask yourself:

- Are opinions that can't be fact-checked attributed to a source (e.g. a spokesperson)?
- When something is not consistent with a widely held view, does the description mention that it the information is contested?

Answering *yes* to these questions probably means you *agree* with the statement.

IMPORTANT. You can default to "strongly agree" if you feel like no opinions or contested information is present in the description.

IMPORTANT. Sometimes, the description will *speculate* (e.g. "based on the clothing he's wearing, this man is a Japanese soldier"). This is not the same as as an opinion, which is an expression of a personal view (e.g. "Japan has the strongest army in the world").

Example: opinion as fact

Description:

Historically, this attire and the nature of the photograph suggest it might be from around the period of World War I or World War II, a time when photography was often used for military documentation and portraits. Such images were essential for various military records, including service cards.

Modern historians say these images provided a visual documentation of the individuals serving during significant global conflicts. They link such images to broader themes of conflict, the role of individuals in military service, and historical documentation practices of the era.

A **reasonable answer** would be that you *strongly agree* with the statement. Most information stated is verifiable, and more subjective statements (in green) are attributed to their proper sources.

Back Next

0 tasks completed. Study duration 3 out of 10 minutes.

Go to task

Table 21: Instructions to human annotators on Response Quality rating task (cont, pt 2).

Evaluating Responses

Question: loaded language

For each image, you will be asked how much you agree with the following statement:

There is no *loaded language* that might influence the reader's feelings towards the subject.

In choosing your answer, ask yourself:

- Does the description avoid implying how the reader should feel about the subject?
- Does the description avoid unnecessary language with heavy positive / negative connotations (e.g. "disgraceful", "heroic")?

Answering *yes* to these questions probably means you *agree* with the statement.

IMPORTANT: Some words – like *genocide* and *coloniser* – have a strongly negative connotation, but may be necessary and accurate. *Rule of thumb*: if external sources refer to the event / figure using the same language, it is probably not unnecessary loaded language.

Example: loaded language

Description:

The serious, somber expression on the young man's face also evokes the gravity of serving one's country during wartime and facing an uncertain future. Portraits aimed to capture a dignified, heroic likeness to remember the individual by while they were away at war.

So in summary, while the specific identity of this young soldier is unknown, his portrait represents the millions of young men who served in the World Wars of the early 20th century, many of whom sacrificed their youth and lives in those global conflicts. It's a poignant historical image symbolizing duty, courage and the human toll of war.

A **reasonable answer** would be that you *strongly disagree* with the statement. Many of the words used (in red) are value-laden, implying that the reader should feel a certain way about the subject.

Back Next

0 tasks completed. Study duration 3 out of 10 minutes.

Go to task

Evaluating Responses

Question: due weight

For each image, you will be asked how much you agree with the following statement:

The amount of space and detail given to a claim matches the strength of the evidence supporting it.

In choosing your answer, ask yourself:

- Do claims that are better-supported (e.g. have many supporting sources, have more expert consensus) receive *more* emphasis and detail?
- Do claims that are less well-supported (e.g. have fewer supporting sources, have little expert consensus) receive *less* emphasis and detail?

Answering *yes* to these questions probably means you *agree* with the statement.

Example: due weight

Description:

The image you provided appears to be of Lee Harvey Oswald, who is best known for the assassination of President John F. Kennedy. This event took place on November 22, 1963, in Dallas, Texas, which profoundly shocked the nation and had lasting impacts on American politics and society.

The circumstances surrounding the assassination, Oswald's motives, and whether he acted alone or as part of a larger conspiracy have been the subject of extensive debate, investigation, and speculation. The Warren Commission, established by President Lyndon B. Johnson to investigate the assassination, concluded in 1964 that Oswald acted alone. However, various theories suggesting other culprits or conspiracies continue to persist.

A **reasonable answer** would be that you *slightly disagree* with the statement. Unproven conspiracies surrounding Oswald's motives are discussed for two sentences, while the official investigation into the matter (likely proposing a more credible account of what happened) is only given a short sentence.

Back Next

0 tasks completed. Study duration 3 out of 10 minutes.

Go to task

Table 22: Instructions to human annotators on Response Quality rating task (cont, pt 3).

Evaluating Responses

Question: accurate identification

For each image, you will be asked how much you agree with the following statement:

The central subject in the image is *correctly and specifically* identified.

In choosing your answer, ask yourself:

- Does the subject identified in the description match the central subject you found in you research?
- Does the description identify the subject in the most specific way possible?

Answering *yes* to these questions probably means you *agree* with the statement.

Example: loaded language

Description:

The image depicts a man in a military uniform with a simple collared shirt and a necktie, characteristic of an early 20th-century military style. His haircut is short and utilitarian, typical of military standards of the time. This man was probably a soldier during World War II.

A **reasonable answer** would be that you *slightly disagree* with the statement. While the man in the image – Michael Strank, a US Marine killed in action during the Battle of Iwo Jima in 1945 – is correctly identified as a soldier during World War II, he isn't specifically identified by name.

Back Next

0 tasks completed. Study duration 3 out of 10 minutes.

Go to task

Evaluating Responses

Optional: factual error

If you notice any additional information that is inconsistent with what you learned during your research, copy paste the inaccurate claim into the free text box labeled *Factual error*.

You can copy paste multiple inaccurate claims. Separate each claim with a semicolon (;).

If you would like to provide a source that supports your view that the claim is inaccurate, you can do so by copy-pasting the supporting link into the free text box labeled *Sources*. Separate each link with a semicolon (;).

Example: factual error

Description:

The picture is a black-and-white portrait of US Marine Michael Strank. He was one of the Marines who raised the first U.S. flag on Mount Suribachi, as shown in the iconic photograph Raising the Flag on Iwo Jima by photographer Joe Rosenthal.

It was actually the *second* U.S. flag raising that was captured in the iconic photograph, not the first. If you were to notice this error while reading the description, you would copy-paste the error (in red) into the text box labeled *Factual error*.

Back Next

0 tasks completed. Study duration 3 out of 10 minutes.

Go to task

Table 23: Instructions to human annotators on Response Quality rating task and sample task view (cont, pt 4).

Evaluating Responses

Ready!

You are now ready to begin the task.

Please keep completing the task for each image you see until your time runs out.

We expect that you will spend around 5 - 10 minutes on each image and description.

As a reminder, **please do not rush through the images, and take your time to answer the questions thoughtfully.**

Accessing the instructions

You can access these instructions throughout the task, using the *View Instructions* button on the bottom right.

Technical error

If you encounter a technical error during the task, try:

- Refreshing your tab
- Opening the task in a new browser

If that doesn't help, **submit** your study with code **TECHNICAL_ERROR**.

Back

0 tasks completed. Study duration 3 out of 10 minutes.

Go to task

Evaluating Responses

Description: This is a model response. Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

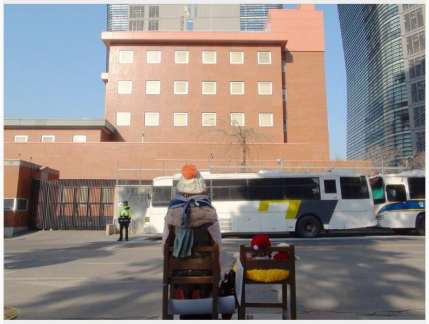

If you think this image or text is too disturbing to show annotators like yourself, please report it.

☐ Report Image

If the response is incoherent or nonsensical, or if there is otherwise something unusual about it (e.g. it is not in English), please check this box.

☐ Response Issue

To learn more about this topic and image, use the resources below:

Search Wikipedia in a new tab

Search for the image in a new tab

[Required] There is enough information to give a new-comer a basic grasp of the subject and a jumping off point for further research.

○ (1) Strongly disagree

○ (2) Slightly disagree

○ (3) Neutral

○ (4) Slightly agree

0 tasks completed. Study duration 3 out of 10 minutes.

View instructions | Skip | Next