



Socratis: Are large multimodal models emotionally aware?

^{1*}Katherine Deng, ^{1*}Arijit Ray, ^{1*}Reuben Tan, ³Saadia Gabriel, ¹Bryan A. Plummer, ^{1,2}Kate Saenko
¹Boston University, ²Meta AI (FAIR), ³MIT

*Equal Contribution

Abstract

Existing emotion prediction benchmarks contain coarse emotion labels which do not consider the diversity of emotions that an image and text can elicit in humans due to various reasons. Learning diverse reactions to multimodal content is important as intelligent machines take a central role in generating and delivering content to society. To address this gap, we propose *Socratis*, a societal reactions benchmark, where each image-caption (IC) pair is annotated with multiple emotions and the reasons for feeling them. *Socratis* contains 18K free-form reactions for 980 emotions on 2075 image-caption pairs from 5 widely-read news and image-caption (IC) datasets. We benchmark the capability of state-of-the-art multimodal large language models to generate the reasons for feeling an emotion given an IC pair. Based on a preliminary human study, we observe that humans prefer human-written reasons over 2 times more often than machine-generated ones. This shows our task is harder than standard generation tasks because it starkly contrasts recent findings where humans cannot tell apart machine vs human-written news articles, for instance. We further see that current captioning metrics based on large vision-language models also fail to correlate with human preferences. We hope that these findings and our benchmark will inspire further research on training emotionally aware models.

1. Introduction

A crucial prerequisite for effective communication and collaboration is the ability to possess emotional awareness [7]. As intelligent machines become increasingly prevalent ranging from generative content creation [19] to collaborative [13] and embodied AI [15], they need to possess emotional awareness for effective communication, and greater trust and acceptance [7]. Emotionally unaware messaging undermine efforts to inform people of global crises [4], spread political division [23], and fail to engage people for the correct social causes [5]. Existing work on emo-




Figure 1. **Socratis benchmark.** The same image and caption pair can evoke different emotions and reactions. We release a benchmark dataset of diverse human-annotated emotions and reactions to images and captions. We show that current state-of-the-art language models and metrics fail to capture the nuance of this task.

tion prediction oversimplifies the problem by categorizing emotions into coarse buckets [14, 25, 3], ignoring the nuance that the same content can elicit various shades of emotional reactions for various reasons. For instance, as shown in Figure 1, an image and caption can generate two conflicting emotions with valid reasons. Learning this diversity of reactions is crucial to tailor a machine’s interactions to individual emotional states.

To encourage further research on emotionally aware AI, we propose the *SOCRATIS* benchmark - a dataset of detailed diverse reactions written by humans on images and captions. *SOCRATIS* includes 980 shades of emotions and 18K free-form reactions written by humans on 2075 image and caption pairs collected from 5 existing news and image-caption datasets. Given an image, a news caption and an emotion word, our task is to generate a “reaction” caption that explains why the image and caption may elicit the specified emotion in a person.

Unlike related benchmarks that focus on reactions to artistic images [12, 1], niche topics such as gun violence [18], emotionally stylized captions [2], or on morality [9], we focus on reasons why humans might feel various emotions on real-life images and captions from widely used news and image-caption datasets. Our task has higher practical relevance since models that fare well on this benchmark can be used to create more effective and inclusive messaging by news agencies and social workers. For in-

stance, as shown in Figure 1, a writer can look at why someone may be “scared” and update the content to reflect that transporting the rhinoceros by helicopter is safe in this context to effectively mitigate the doubt. While a similar benchmark [6] focuses on reactions to reading a text caption, we focus reactions to both images and text since content on the web is largely multimodal.

Using Socratis as a testbed, we evaluate a state-of-the-art multimodal language model [10] and check if commonly used metrics in language-generation evaluation can distinguish between good “reactions” and poor ones. We generate “reactions” given an image, caption, and emotion and ask human raters to blindly choose between the machine generation and the human annotation.

Our results show a stark gap in current capabilities for this task. Humans prefer human-written reactions two times more often than machine-generated ones. This starkly contrasts recent findings where humans fare poorly to tell the difference between machine-generated and AI-generated news articles [27] or images [19]. This illustrates that while large generative models may be good at producing believable articles or images, it lacks the nuance required for emotional awareness. Furthermore, when we separate the generations into two groups - good generations as rated by humans (when humans couldn’t tell the difference between human-written and machine-generated) and poor generations (when humans picked human-written over machine-generated), we observe a negligible difference in the scoring by commonly used metrics like BART and CLIP score, which are also based on large language and vision models. This illustrates the difficulty of the problem since we need better emotionally aware models to make better metrics.

Hence, we hope these initial results spark further research and discussion into improving the emotional awareness of large language models. Adding to the recent discussion that large language models seem to lack rich social intelligence [20] and theory of mind [21], our benchmark shows that they lack nuanced emotional awareness as well.

2. SOCRATIS Dataset

We propose a benchmark for evaluating the emotional awareness of vision-language models. Specifically, our task is to predict the reaction a human may have of a certain emotion while looking at an image and caption (IC) pair. To this end, we collect a dataset showing human workers an IC pair and ask them to write the emotion words they feel and the reasons for feeling them. We interchangeably call these reasons as “reactions” the humans feel of a certain emotion. We have 18,378 annotated reactions for 980 emotions to 2075 IC pairs rated by an average of 8 independent workers per IC pair. The most common emotion words follow that of standard emotion datasets [14, 3] - happy, sad, excited, angry. We also have a variety of more subtle emotion

words such as inspire, hopeful, and nostalgic that feature among the top 30 emotion words. However, the novelty of our dataset isn’t the variety of emotion words, but the free-form reasons for feeling the different emotions for the same IC pair. Some examples are shown in Figure 2. We will make our dataset publicly available.

2.1. Data Annotation

Image-Caption pair collection We collect image-caption pairs from the Visual News [11] and the Conceptual Captions [22] dataset. We first randomly sample 1800 images from the Visual News Dataset [11], which consists of an image and a news headline caption from 4 news datasets - BBC, USA-Today, Guardian and The Washington Post. We specifically choose news datasets since news headlines and visuals usually elicit stronger reactions from humans than generic stock images and are of more practical use. Additionally, to also understand how people react to generic IC pairs, we sample images and captions from a widely used image-captioning dataset, Conceptual Captions [22]. IC pairs that explicitly convey a certain emotion (eg, stock photo of someone smiling) are less interesting due to a lack of ambiguity from diverse people. Hence, to sample IC pairs that are likely to elicit diverse emotional reactions, we choose 500 samples where the emotion of the image does not match the emotion of the caption. The emotion of the image is predicted by a CLIP [16] model, fine-tuned on the WebEmo [14] dataset, and a T5 model [17] predicts the text emotions.

Human Reaction Annotation We show users the IC pairs we collect as described above. We then ask them to write up to three emotions a person is likely to feel while viewing the IC pair, along with a reason for why they might feel each of the emotions they entered. We collect our dataset on Amazon Mechanical Turk. All data is cleared of personally identifiable information and only aggregate statistics of model training are reported in the paper. We incentivize the annotators by awarding them a small bonus if they can match the most popular emotion word for an IC pair. This encourages workers not to enter noisy reactions (or extremely uncommon or made-up words) since they are not likely to match other workers’ responses. To control the quality of responses, we choose workers with a greater than 98% approval rate (on at least 50 Human Intelligence Tasks or HIT’s). We also restrict the geographical location to the US or UK since the images and captions are sourced from news articles from these two countries. We collect 10 independent annotations for each image-caption pair to get a representative, diverse set of reactions.

2.2. Annotation Quality

We compute some heuristic automatic measures to judge the quality of the reactions.






Image	Caption	Emotion	Human Written	BLIP-2 Generated
	Alexei Navalny has complained about a crackdown on opposition members in Kostroma	Fear	They would fear for this politician's life as the pro-democracy faction carry on the fight in Russia	A person is being beaten
		Hope	They would feel hope because someone is standing up for the Russian people	A person feels hope because the opposition is fighting for their rights
	A margarita available soon in powdered form thanks to Palcohol	Refreshing	The drink looks tasty and cool	A margarita is a refreshing drink
		Curious	Wondering if a powder based margarita would have the same flavor as a "real" margarita	A powdered margarita is a new idea
	Customers using prepaid electricity meters pay on average 80 more per year on their energy bills than direct debit customer	Interested	I would be interested in knowing how the task shown works out.	The person is interested in the topic of electricity
		Upset	Someone would feel like they are being taken advantage of and could have saved more money.	The person feels that they are being ripped off
	Polar bear cubs were last raised in the UK in the early 1990s	Affection	They would feel affectionate towards this adorable animal	A person feels affection after seeing this image and reading the news caption
		Happy	The cub is a baby and baby's make people happy.	Polar bears are endangered
	Pope Francis celebrates Christmas Eve midnight Mass at St Peter's Basilica	Hope	They would feel hope for the the coming year and the Birth of Jesus on Christmas.	He is a religious leader
		Excited	Excited because the pope is celebrating Christmas.	A person is a religious leader

Figure 2. Qualitative examples from our Socratis dataset and a state-of-the-art multimodal model, BLIP-v2 generations.

Emotion-reaction match We check how often the entered emotion word matches the emotion of the reaction. For instance, the emotion of “sad” reaction should also be “sad”. To compute the sentiment, we use a T5 sentiment model [17], which predicts positive or negative sentiment. We observe an 87.15% accuracy of sentiment match between the emotion word and the reaction. We manually check some examples where the sentiments of reaction and emotion words mismatch and find that the T5 model [17] may be noisy. For instance, for the emotion word “hungry”, the reaction annotated is “the cocoa actually makes me want to eat something sweet.”. We believe this is a reasonable reaction for “hungry”. However, the sentiment predicted for “hungry” is “negative”, whereas the same for the reaction is “positive”, resulting in a mismatch. This further illustrates that coarse emotion buckets like positive and negative don’t capture the nuances of reactions of what people feel.

Agreement on reactions To judge the agreement of humans over the reactions for a given image, caption, and emotion, we further compute the BART [26] scores of the reactions for the same (I,C,emotion) tuple. We contrast this with the BART score [26] of 1000 randomly sampled reaction pairs from different IC-pairs and emotions. We note that the BART scores between reactions for the same emotion for an IC pair are higher 78% of the time than that of the random pairs.

3. Approach

To understand the emotional awareness of large multimodal models, we benchmark the capability of a state-of-the-art vision-language model on our proposed benchmark without further fine-tuning. Specifically, we

use the FLAN-T5 variant of BLIP-2 [10] from Huggingface [24]. Our task is to predict the reason of feeling a certain emotion given an image-caption pair. For a given image, caption text and emotion word, we prompt BLIP-2 with the image and a query using the following template: `Question: Why does a person feel {emotion} after seeing this image and reading the news caption '{caption}'?` `Answer:`

We input this prompt along with the image to the BLIP-2 model, and use a greedy approach to generate the response following the standard procedure outlined on the Huggingface [24] model page. Given an image i , caption c and emotion e , we can formulate the likelihood of the generated reaction text t as:

$$P(t|i, c, e) = \prod_{j=1}^n P(t_{<j}|i, c, e) \quad (1)$$

where n denotes the number of tokens in the generated text.

4. Experiments

Our goal is to evaluate whether large multimodal models are emotionally aware to generate plausible reasons for why humans might feel a certain reaction for a given image and text. A pragmatic formulation of this evaluation is whether a human finds it hard to distinguish a machine generated reason from a human-written reason. Hence, we first conduct a human evaluation on the multimodal BLIP-2 generated reasons to see how often humans prefer machine-generated reasons, human-written reasons, or both.

Human Evaluation To understand human preference, we conduct a preliminary human study on 500 randomly sampled data points from our dataset. We show an image, cap-

Table 1. Number of times majority ($\frac{2}{3}$) humans prefer human-written reactions, BLIP-2 (machine) reactions, or both out of 382 examples.

Machine-better	Human-better	Both-Good	Both-Bad
91	233	47	11

tion, emotion tuple to a user. We then show two choices of reactions - human-written and machine-generated. The user is unaware of which of the reactions is machine-generated or human-written. We ask 3 independent workers to choose the best reaction for the given tuple of image, caption and emotion. They also have the choice of either choosing that both reactions are reasonable or that neither is reasonable. Hence, there are four choices for each image-caption-emotion tuple: We split the human study image-caption-emotion tuples into three groups:

- **Machine-better:** machine was picked over human, indicating that the machine-generations are good.
- **Both-Good:** both human-written and machine-generated reactions are equally valid.
- **Both-Bad:** We discard these examples from further study since the data annotations are likely noisy.

Metrics Since human evaluations are slow and expensive, we aim to determine if commonly used captioning metrics can be used to judge good reaction generations from poor ones to speed up research. We define a good generation as one which is aligned to or better than human preferences. Hence, a machine generation is good if a human cannot tell the difference from a human-written one, or if a human prefers it over a human-written one. Hence, a good metric should score such cases (machine-preferred and both-good) higher than the cases where a human-written reason was preferred over a machine-generated one (human-preferred). We compute three commonly used metrics. First, we compute **BART Score** between the machine-generated reaction and the human-generated reference to measure human-likeness of generations as described in [26]. Next, we compute **CLIP-Score** and **RefCLIP-Score** [8] to see if image-relevance plays a major factor in distinguishing good generations from poor ones. We use a prompt like “Human feels {emotion}, when seeing this image with {caption} because {explanation}” and compute the cosine similarity to the image as described in [8].

5. Results

Humans prefer human-written reactions to machine generations. In Table 1, we observe that workers pick human reactions over two times more often than machine-generated ones (233 vs 91). This suggests that state-of-the-art large vision-and-language models are still limited at extrapolating contextual information beyond simply correlat-

Table 2. Evaluations with BART and CLIP-Score on subsets where humans prefer human-written reactions, BLIP-2 (machine) reactions, or both. We want machine-better or both-good generations scored higher than human-better.

	Machine-better	Human-better	Both-Good
BART	-5.48	-5.42	-5.54
CLIP-Score	0.74	0.76	0.75
RefCLIP	0.42	0.43	0.42

Table 3. Multimodal vs a text-only model on relevance of generation to the image

	CLIP \uparrow	Ref-CLIP \uparrow
BLIP-2	0.75	0.42
FLAN-T5	0.73	0.42

ing the visual concepts in the image with relevant words.

Current captioning metrics cannot distinguish between good and bad reactions In Table 2, we see that BART scores do not follow human preferences. The scores of the machine generations for when the machine was rated better or equally good (both-good) by humans are not higher than the scores when the machine generations are poor (human-better). CLIP and RefCLIP also do not seem to differ across the three sets. This suggests that visual similarity may also not be important in distinguishing good from poor reactions. Further investigation is required to check if we can train a custom BART metric based on a few examples in our dataset.

Multimodal models are slightly more image-relevant We also check the performance of a language-only model compared to the multimodal BLIP-2 by using only the language model in BLIP2, which is a FLAN-T5 model. In Table 3, we see that the relevance to the image is understandably higher (acc to CLIP scores) for BLIP-2. However, based on results in Table 2, this doesn’t necessarily mean that the generations are more preferred by humans.

Discussion and Conclusion We propose the Socratis benchmark to generate reactions for why images and captions elicit certain emotions in humans. Our initial experiments indicate that a state-of-the-art vision-language model is unable to extrapolate the required information to generate reasonable reactions. However, further research is required to investigate biases that exist in our dataset and how these are perpetuated in the models that are benchmarked on it. Further investigation is also required to see if quick fixes like changing the generation strategy, adapting a few layers, or in-context learning (showing examples in prompts) can make these models more emotionally aware. We hope our Socratis benchmark will encourage future research on training and evaluating emotionally aware AI algorithms.

Acknowledgements: We are thankful to Praneeth Chandra Bogineni for valuable discussions in the initial phase of the project. This material is based upon work supported, in part, by DARPA under agreement number HR00112020054. The findings in the paper do not reflect the opinions of the US Government or DARPA.

References

- [1] Panos Achlioptas, Maks Ovsjanikov, Kilichbek Haydarov, Mohamed Elhoseiny, and Leonidas J Guibas. Artemis: Affective language for visual art. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11569–11579, 2021.
- [2] Kezhen Chen, Qiuyuan Huang, Daniel McDuff, Xiang Gao, Hamid Palangi, Jianfeng Wang, Kenneth Forbus, and Jianfeng Gao. Nice: Neural image commenting with empathy. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4456–4472, 2021.
- [3] Vaidehi Dalmia, Hongyi Liu, and Shih-Fu Chang. Columbia mvso image sentiment dataset, 2016.
- [4] Ding Ding, Edward W Maibach, Xiaoquan Zhao, Connie Roser-Renouf, and Anthony Leiserowitz. Support for climate policy and societal action are linked to perceptions about scientific agreement. *Nature Climate Change*, 1(9):462–466, 2011.
- [5] Ajay Divakaran, Karan Sikka, Arijit Ray, Xiao Lin, and Yi Yao. User targeted content generation using multimodal embeddings, Sept. 23 2021. US Patent App. 17/191,698.
- [6] Saadia Gabriel, Skyler Hallinan, Maarten Sap, Pemi Nguyen, Franziska Roesner, Eunsol Choi, and Yejin Choi. Misinfo reaction frames: Reasoning about readers’ reactions to news headlines. *arXiv preprint arXiv:2104.08790*, 2021.
- [7] Daniel Goleman. Emotional intelligence. why it can matter more than iq. *Learning*, 24(6):49–50, 1996.
- [8] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021.
- [9] Liwei Jiang, Jena D Hwang, Chandra Bhagavatula, Ronan Le Bras, Jenny Liang, Jesse Dodge, Keisuke Sakaguchi, Maxwell Forbes, Jon Borchardt, Saadia Gabriel, et al. Can machines learn morality? the delphi experiment. *arXiv preprint arXiv:2110.07574*, 2021.
- [10] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023.
- [11] Fuxiao Liu, Yinghan Wang, Tianlu Wang, and Vicente Ordonez. Visual news: Benchmark and challenges in news image captioning. *arXiv preprint arXiv:2010.03743*, 2020.
- [12] Youssef Mohamed, Mohamed Abdelfattah, Shyma Alhuwaidar, Feifan Li, Xiangliang Zhang, Kenneth Ward Church, and Mohamed Elhoseiny. Artelingo: A million emotion annotations of wikiart with emphasis on diversity over language and culture. *arXiv preprint arXiv:2211.10780*, 2022.
- [13] OpenAI. Gpt-4 technical report, 2023.
- [14] Rameswar Panda, Jianming Zhang, Haoxiang Li, Joon-Young Lee, Xin Lu, and Amit K Roy-Chowdhury. Contemplating visual emotions: Understanding and overcoming dataset bias. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 579–595, 2018.
- [15] Xavier Puig, Tianmin Shu, Shuang Li, Zilin Wang, Yuan-Hong Liao, Joshua B Tenenbaum, Sanja Fidler, and Antonio Torralba. Watch-and-help: A challenge for social perception and human-ai collaboration. *arXiv preprint arXiv:2010.09890*, 2020.
- [16] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [17] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [18] Carley Reardon, Sejin Paik, Ge Gao, Meet Parekh, Yanling Zhao, Lei Guo, Margrit Betke, and Derry Tanti Wijaya. BU-Nemo: an affective dataset of gun violence news. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2507–2516, Marseille, France, June 2022. European Language Resources Association.
- [19] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [20] Maarten Sap, Ronan Le Bras, Daniel Fried, and Yejin Choi. Neural theory-of-mind? on the limits of social intelligence in large LMs. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3762–3780, Abu Dhabi, United Arab Emirates, Dec. 2022. Association for Computational Linguistics.
- [21] Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi, and Yulia Tsvetkov. Minding language models’ (lack of) theory of mind: A plug-and-play multi-character belief tracker. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13960–13980, Toronto, Canada, July 2023. Association for Computational Linguistics.
- [22] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018.
- [23] Bertie Vidgen, Austin Botelho, David Broniatowski, Ella Guest, Matthew Hall, Helen Margetts, Rebekah Tromble, Zeerak Waseem, and Scott Hale. Detecting east asian prejudice on social media. *arXiv preprint arXiv:2005.03909*, 2020.
- [24] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim

Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.

- [25] Jingyuan Yang, Qirui Huang, Tingting Ding, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Emoset: A large-scale visual emotion dataset with rich attributes. *arXiv preprint arXiv:2307.07961*, 2023.
- [26] Weizhe Yuan, Graham Neubig, and Pengfei Liu. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277, 2021.
- [27] Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news. *Advances in neural information processing systems*, 32, 2019.