

ArtELingo: A Million Emotion Annotations of WikiArt with Emphasis on Diversity over Language and Culture

Youssef Mohamed^{1*} Mohamed Abdelfattah¹ Shyma Alhuwaider¹ Feifan Li¹

Xiangliang Zhang² Kenneth Ward Church³ Mohamed Elhoseiny^{1*}

¹KAUST ² University of Notre Dame ³ Northeastern University

{youssef.mohamed, mohamed.abdelfattah, shyma.alhuwaider, feifan.li}@kaust.edu.sa
 xzhang33@nd.edu, k.church@northeastern.edu, mohamed.elhoseiny@kaust.edu.sa







a)		شلال طبيعي جميل. مشاعر النمو والحيوية والطاقة موجودة. Translation: Beautiful natural waterfall. Feelings of growth, vitality and energy.	The water that's rushing downward looks like a bride's wedding veil.	瀑布就像四蹄生风的白马如潮水涌来，非常的壮观 Translation: The waterfall is like a white horse and wind, it is spectacular.
	Excitement Arabic		Awe English	
			Contentment Chinese	
b)		Translation: Girls sitting with their mother outside the house, exchanging love and affection, pigeons flying over a tree.	The women relaxing while birds are flying about makes me feel relaxed and calm as well.	Translation: Three sisters lying on a bench and watching the birds fly comfortably.
	Contentment Arabic		Contentment English	
			Contentment Chinese	
c)		Translation: The use of black and white for painting the forests with all its details brings out a feeling of satisfaction.	The trees are dead and exposing their roots due to erosion and lack of water.	Translation: After the snow in winter, there is snow everywhere, and the dead trees look very depressed.
	Contentment Arabic		Sadness English	
			Sadness Chinese	

Figure 1: ArtELingo, a multilingual dataset and benchmark of WikiArt with captions & emotions

Abstract

This paper introduces ArtELingo, a new benchmark and dataset, designed to encourage work on diversity across languages and cultures. Following ArtEmis, a collection of 80k artworks from WikiArt with 0.45M emotion labels and English-only captions, ArtELingo adds another 0.79M annotations in Arabic and Chinese, plus 4.8K in Spanish to evaluate “cultural-transfer” performance. More than 51K artworks have 5 annotations or more in 3 languages. This diversity makes it possible to study similarities and differences across languages and cultures. Further, we investigate captioning tasks, and find diversity improves the performance of baseline models. ArtELingo is publicly available¹ with standard splits and baseline models. We hope our work will help ease future research on multilinguality and culturally-aware AI.

* Corresponding Authors

¹www.artelingo.org

1 Introduction

Figure 1 compares and contrasts annotations on WikiArt across language/culture. We believe these differences are interesting and important, and far from random. One might suggest using machine translation to translate English captions to many other languages, but we believe that doing so would miss much of the opportunity. Building human-compatible AI that is more aware of our emotional being is important for increasing the social acceptance of AI. ArtEmis (Achlioptas et al., 2021) is an important step in this direction, introducing a collection of 0.45M emotion labels and affective language explanations in *English* on more than 80,000 artworks from WikiArt. However, by design, ArtEmis is limited to English, lacking coverage of other cultures and languages.

Cultural differences are a major source of diversity (Meyer, 2014). The customs, social values,

lifestyles, and history of different countries and cultures greatly influence human behavior. Emotional experiences are no exception; people from different countries respond differently to similar scenarios. For example, a person born and raised in a Nordic country would be more comfortable in a lush forest than in a desert, but a Bedouin may be more comfortable in a desert than in a forest.

Consider Figure 1c, where an Arabic annotator assigned the image the label **contentment**, but the other two annotators used the label: **sadness**. Captions are useful for diving deeper into these differences. The **sadness** annotations mention *death*² and disasters,³ in contrast with the **contentment** annotation that ends with: *feeling of satisfaction*.

There can be interesting differences between languages/cultures even when annotators use the same label. Consider Figure 1b, where all three labels are **contentment**. Although the three captions agree on the label, two of the captions imply that some/all of the girls are sisters, but there is no such implication in the English caption.

We believe deep nets will be viewed as more culturally aware, if they can capture linguistic/cultural patterns such as these. Emotions are based on past experience, and play an integral role in determining human behavior. Not only they reflect our internal state but also directly effect how we perceive, interpret external stimuli (Izard, 2009), and how to act based on them (Lerner et al., 2015). Hence, studying emotions is essential to exploring a confounding aspect of human intelligence.

In summary, our **contributions** are:

1. 0.79M annotations (labels + captions) in Arabic and Chinese, plus 4.8k in Spanish,
2. a benchmark with standard splits, and
3. baseline models for two tasks: (1) label prediction and (2) affective caption generation.

The rest of the paper is organized as follows: related work is discussed in §2, followed by our main motivation in §3, and data collection in §4. §5 provides qualitative and quantitative analyses of ArtELingo. Baseline models for emotion label prediction and caption generation are presented in §6 and §7, respectively.

²冬天下雪后到处白雪皑皑，枯树显得很萧条。(snow everywhere and dying trees is depressing)

³no me gusta el ambiente, lo primero que me vino a la mente fué un desastre natural con destrucción a su paso (mentions a natural disaster)



(a) ArtEmis: *I love everything about this painting of a mother and her two children lovingly interacting with the family pet cat.*



(b) COCO: *A man and a woman holding a little kid while sitting at a table outside*

Figure 2: COCO captures the facts, and ArtEmis enhances those facts with emotion/commentary.

2 Related Work

2.1 Captions with Emotions

Work on captioning is moving beyond factual captions in early benchmarks such as COCO (Lin et al., 2014). Figure 2 shows two images of families, one from ArtEmis and the other from COCO. Both captions capture the facts, but ArtEmis enhances the facts with emotion/commentary.

Table 1 compares three benchmarks: COCO (Lin et al., 2014), ArtEmis and ArtELingo. ArtEmis encourages work on emotions by replacing COCO photos with WikiArt,⁴ and by introducing 9 emotion classes, 4 positive,⁵ 4 negative⁶ and Other. ArtELingo encourages researchers to work on visually grounded multilinguality by providing affective annotations in three languages (henceforth, ACE/ACES): Arabic, Chinese and English. In addition, we provide a small set of Spanish (S). Figure 3 shows that positive emotions are more frequent than negative emotions, especially in Arabic.

⁴<https://www.wikiart.org/>

⁵Positive: **Contentment, Awe, Amusement, Excitement**

⁶Negative: **Sadness, Fear, Disgust, Anger**

	COCO	ArtEmis	ArtELingo
Image Source	Photos	WikiArt	WikiArt
#Images	328k	80k	80k
#Annotations	2.5M	0.45M	1.2M
#Annot/Image	7.6	5.68	15.3
Emotions	0	9	9
Languages	E	E	ACES

Table 1: A Comparison of Three Datasets. ArtELingo has a million annotations from ACES: Arabic (A), Chinese (C), English (E) and Spanish (S).




Figure 3: In general, positive emotions are more frequent than negative emotions. The 9 plots are sorted by probability. The log scale on y-axis highlights relative probabilities. A (Arabic) is relatively high for some classes (**awe**), and low for others (**sadness**).

2.2 Related Work in Other Fields

There is a considerable literature on emotions, especially in Psychology (Russell and Barrett, 1999). One can find quite a few benchmarks on emotion in HuggingFace: (Saravia et al., 2018; Demszky et al., 2020; Xiao et al., 2018).⁷ There are a number of papers in computational linguistics on emotion and Chinese (Chen et al., 2020; Quan and Ren, 2009; Wang et al., 2016; Lee et al., 2010), and on emotion and Arabic (Abdullah and Shaikh, 2018). There is also considerable work on emotion in other fields such as vision (Mittal et al., 2021).

Many datasets have been collected to study emotional responses to modalities such as:

- Text (Strapparava and Mihalcea, 2007; Demszky et al., 2020; Mohammad et al., 2018; Liu et al., 2019),⁸
- Image (Mohammad and Kiritchenko, 2018; Kosti et al., 2017), and
- Audio (Cowen et al., 2019, 2020).

Bias is the flip side of inclusiveness. There has been considerable discussion recently about biases (Bender et al., 2021; Bolukbasi et al., 2016; Buolamwini and Gebru, 2018; Mehrabi et al., 2021;

⁷<https://huggingface.co/datasets?sort=downloads&search=emotion>

⁸<https://data.world/crowdflower/sentiment-analysis-in-text>

Liu et al., 2021). Some of this work is more relevant to our interest in Chinese (Jiao and Luo, 2021; Liang et al., 2020), and Arabic (Abid et al., 2021). Many machine learning methods will, at best, learn what is in the training data. There have been some attempts to remove biases in corpora, but it might also be constructive to create more inclusive benchmarks such as ArtELingo.

Awareness of different cultures is becoming increasingly important. Gone are the days when it was sufficient for datasets to focus on a single culture. Recently, the Vision & Language community has been producing more multicultural multilingual datasets (Bugliarello et al., 2022; Srinivasan et al., 2021; Armitage et al., 2020). ArtELingo contributes cultural diversity over emotional experiences. The effect of culture on psychology has been studied in separate studies (Henrich et al., 2010; Abu-Lughod, 1990; Norenzayan and Heine, 2005). ArtELingo provides empirical evidence that might motivate cultural psychology studies.

3 Opportunities for Improvement

Many of the resources mentioned above have advanced our understanding of the relationship between emotion and various stimuli, through there are always opportunities for improvement. We are particularly interested in three such opportunities: scale, multimodality and multilinguality/multiculturalism. As for scale, demand for larger training sets is expected to continue to increase, given the rise of large scale foundation models (Bommasani et al., 2021).

As for multimodality, although most benchmarks mentioned above focus on a single modality, there are a few multimodal exceptions such as IEMOCAP (Busso et al., 2008), COCO and ArtEmis. IEMOCAP collected speech and facial and hand movements of 10 actors. Unfortunately, this approach may be expensive to scale up.

The use of Amazon Mechanical Turk in ArtEmis is easier for scaling, however, ArtEmis is limited to English. ArtELingo addresses multilinguality/multi-culturalism by adding Arabic and Chinese annotations. We use languages as a proxy to reflect different cultures. English is a representative sample of the West, and Chinese is a representative sample of the East, and Arabic is a representative sample of the Middle East.

Region	#Artworks	%
West (Non English)	142.8k	57.1%
West (English)	54.0k	21.6%
Other	38.0k	15.2%
Middle East (Non Arabic)	12.2k	4.8%
Middle East (Arabic)	1.6k	0.6%
East (Chinese)	1.4k	0.5%
Total	250.0k	100%

Table 2: WikiArt is more representative of the West

3.1 Representation of Regions in WikiArt

ArtELingo assumes that WikiArt is a representative sample of the cultures of interest. While WikiArt is remarkably comprehensive, Table 2 suggests the WikiArt collection has better coverage of the West than other regions of the world. This table is based on WikiArt’s assignment of artworks to nationalities.⁹ We assigned each nationality to West (English¹⁰ and Non English¹¹), Middle East (Arabic¹² and Non Arabic¹³), East (Chinese) and Other.

4 ArtELingo

Following ArtEmis, we employ Amazon Mechanical Turk (AMT) platform to collect our data using interfaces (see Figures 8, 9, 10 in the appendix). We faced a lack of Arabic and Chinese speaking annotators on AMT which led us to devise different strategies to recruit annotators. Arabic speakers were recruited by advertising the task in middle eastern universities encouraging students and their families to join our data collection efforts. Whereas Chinese speakers were recruited through Baidu who we’d like to thank.

Annotators are asked to carefully examine each artwork before selecting the dominant emotion induced by it from a list of four positive, four negative

⁹<https://www.wikiart.org/en/artists-by-nation>

¹⁰West (English): *Americans, Australians, British and Canadians*

¹¹West (Non English): *Albanians, Armenians, Austrians, Azerbaijanis, Belarusians, Belgians, Bosnians, Bulgarians, Croatians, Czechs, Dutch, Estonians, Finnish, French, Georgians, Germans, Greeks, Hungarians, Icelandic, Irishes, Indigenous North Americans, Italians, Kazahstani, Latvians, Lithuanians, Luxembourgers, Maltese, Montenegrins, Polish, Portuguese, Romanians, Scottish, Slovaks, Serbians, Slovenians, Spanish, Swiss, Swedish, Ukrainians, Uruguayans, Uzbek and Venezuelans*

¹²Middle East (Arabic): *Algerians, Bahraini, Egyptians, Emiratis, Moroccans, Libyans, Lebanese, Iraqi, Palestinians, Qatari, Saudis, Syrians and Tunisians*

¹³Middle East (Non Arabic): *Kenyans, Jewish, Israeli, Iranians and Turkish*




Figure 4: Most (>60%) annotations are from long tail (workers who annotated less than 1K artworks).

	E	C	A	S
#Annotators	6377	745	656	31
#Annotations	429k	426k	369k	4.8K
#Work Hours	10k	13k	9.0k	178

Table 3: Size of the annotation effort by language.

emotions, and *Other* to indicate a different emotion. Annotators are then asked to write captions that reflects the content of the artwork and explains their choice of emotion. Similar to ArtEmis, we collect annotations from five annotators for each artwork.

For a better cultural representation in ArtELingo, we restrict the collection of different languages annotations to countries with large numbers of native speakers. Chinese data is collected from China. For Arabic, we collect our data mainly from Saudi Arabia and Egypt. Finally, Spanish is collected from Latin America and Spain. Figure 4 shows that most of the annotations are from a long tail of workers who annotated less than 1000 artworks ensuring a diverse representation of cultures.

Quality Control. Annotations were rejected if they are too short, or if they are too similar to captions for other artworks. In addition, a manual review was conducted by multiple reviewers, ensuring captions reflect the selected emotion label and the details of the artwork. Table 3 reports some statistics on annotations that passed this review process.

5 Dataset Analysis

5.1 Qualitative

There are some interesting similarities and differences between language and culture, as discussed in Figure 1. There is a considerable inter-annotator agreement (IAA) in the dataset, and there are also some interesting disagreements. There is agreement in Figure 2a that a mother’s love is universally




Figure 5: 8 artworks with genre. Green indicates high agreement in Table 4; red indicates high disagreement.

warm and pleasant. It is an instinct for mothers to be loving, caring and protective of their children.¹⁴ On the other hand, there is a difference in Figure 1a. All three annotators agree to observe a waterfall though some mention energy and growth, while others saw horses and wedding veils.

5.2 Quantitative

Table 4 reports multicultural agreement over the 9 emotions¹⁵ in each genre. WikiArt classifies artworks into 10 genres,¹⁶ as well as 27 styles¹⁷. Agreement is computed as a log likelihood agreement score, $A = \log_2(Pr(G|D)/Pr(G|U))$, where G is one of the 10 genres, and U and D are two sets of artworks. Let $Pr(G|U)$ be the fraction of artworks in U with genre G , and $Pr(G|D)$ be the fraction of artworks in D with genre G .

Let U be the universal set of artworks. That is, U contains all artworks in ArtELingo with 5

¹⁴ English caption for Figure 2a highlights the cat, whereas the Arabic and Chinese focus on the family and do not mention the cat:

أم تجلس مع طفلة صغيرة تنظر إلى طفل يمسك بيدها
وتتحدث وتبادل الحب والمودة.

女人看着自己的孩子，让人觉得很开心。

¹⁵The 9 emotion classes are: Amusement, Awe, Contentment, Excitement, Anger, Disgust, Fear, Sadness, and Other

¹⁶The 10 genres are: portrait, landscape, genre painting (misc), religious painting, abstract painting, cityscape, sketch and study, still life, nude painting and illustration.

¹⁷The art styles are: Abstract Expressionism, Action painting, Analytical Cubism, Art Nouveau Modern, Baroque, Color Field Painting, Contemporary Realism, Cubism, Early Renaissance, Expressionism, Fauvism, High Renaissance, Impressionism, Mannerism Late Renaissance, Minimalism, Naive Art Primitivism, New Realism, Northern Renaissance, Pointillism, Pop Art, Post Impressionism, Realism, Rococo, Romanticism, Symbolism, Synthetic Cubism and Ukiyo-e

Genre (G)	$Pr(G U)$	$Pr(G D)$	A
landscape	0.206	0.097	-1.08
cityscape	0.071	0.036	-0.98
still life	0.043	0.042	-0.03
illustration	0.029	0.029	-0.01
misc	0.167	0.177	0.08
portrait	0.217	0.233	0.10
nude	0.030	0.032	0.11
religious	0.101	0.133	0.40
abstract	0.076	0.112	0.55
sketch	0.061	0.109	0.85

Table 4: Genre sorted by agreement (A). Most agreement: landscapes; Most disagreement: sketches.

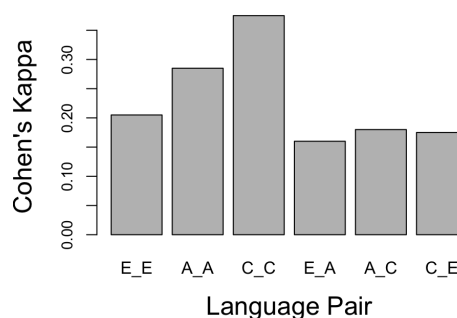


Figure 6: Cohen's Kappa for inter-annotator and cross-annotator agreement. Higher value means more agreement.

annotations in each of the 3 languages. D is a disagreement set of 2000 artworks. D was selected by computing Cohen Kappa scores (Cohen, 1960)¹⁸ for artworks in U . Let D be the 2000 artworks with the most disagreement (based on Kappa).

Table 4 shows that there is more agreement for some genres (landscapes), and more disagreement for other genres (sketches). When the agreement score is near 0, then the genre is about equally likely in U and D . This is to be expected for genres near the middle of the list such as misc. Figure 5 shows 8 artworks in genres with high agreement and high disagreement. Figure 6 reports the Cohen's Kappa score of annotations from language pairs. Annotators belonging to the same language have higher agreement.

We created D for zero-shot experiments to be reported in §6. The 4.8k Spanish annotations in Table 3 are on the set of D artworks with low IAA (inter-annotator agreement) in ACE (Arabic, Chinese and English).

¹⁸https://scikit-learn.org/stable/modules/generated/sklearn.metrics.cohen_kappa_score.html

6 Emotion Label Prediction

Baseline models for two tasks, emotion label prediction and caption generation, will be discussed in this section and the following section. These discussions assume familiarity with deep nets including fine-tuning BERT (Devlin et al., 2019) and cross language models XLM (Conneau et al., 2020), as well as HuggingFace (Wolf et al., 2019).

Emotion Classification. Given an input caption, c , we wish to predict an output emotion label, \hat{e} , where \hat{e} is one of the 9 emotions. The model starts with a pretrained language model, LM , and a tokenizer. The tokenizer converts c into a sequence of L tokens x . The language model converts x into more useful representation, $LM(x) \in R^{L \times d}$, where d is the number of hidden dimensions (a property of the LM). Finally, we feed $LM(x)$ into a linear layer to predict the emotion label, \hat{e} .

Majority Baseline. We use the majority emotion label for each artwork as the predicted emotion for all captions belonging to that artwork. Concretely, each artwork, I , has a set of caption-emotion pairs, S . The majority classifier outputs the most frequent emotion, \hat{e} , in the set S for all of the captions in the set, $c \in S$,

Language Models. We finetune 3 models based on BERT (BERT-E, BERT-A and BERT-C), where BERT-E is tuned for English, and BERT-A is tuned for Arabic and BERT-C is tuned for Chinese. Section 11.2 discusses more pretraining and finetuning details. We also finetune 4 models based on cross language models, XLM-roBERTa (Conneau et al., 2020), where XLM-E, XLM-A and XLM-C correspond to English, Arabic, and Chinese languages, as before. In addition, we create XLM-ACE by training on the combination of all 3 languages.

3-Headed Transformer. Finally, we create a model with XLM-R backbone but replace the single classifier head with 3 classifier heads, one for each of the 3 languages. While training, we feed the captions from each language to the shared backbone and then use the corresponding head to predict an emotion that would ultimately reflect the culture of that language. Geva et al. (2021) analyzed similar multi-headed transformers and showed how the non-target heads can be used to interpret the results of the target head. Similarly, our 3-headed transformer can be used to predict 3 different emotions each one reflecting the culture norms represented in each language. We can then use these predictions to better understand the similarities and differences

Backbone	Test Set				
	E	A	C	ACE	S (0-Shot)
Majority	0.474	0.491	0.604	0.525	-
BERT-E	0.644	-	-	-	-
BERT-A	-	0.558	-	-	-
BERT-C	-	-	0.922	-	-
XLM-E	0.662	0.345	0.781	0.606	0.513
XLM-A	0.446	0.556	0.695	0.569	0.437
XLM-C	0.482	0.349	0.926	0.599	0.415
XLM-ACE	0.663	0.558	0.927	0.724	0.519
3-Headed-E	0.660	0.478	0.914	0.694	0.529
3-Headed-A	0.597	0.542	0.854	0.672	0.501
3-Headed-C	0.630	0.474	0.924	0.687	0.495
3-Headed-M	0.653	0.498	0.917	0.700	0.525

Table 5: **Emotion Label Classification Baselines.** Majority baseline output the most frequent emotion for each artwork. Models are fine-tuned on BERT and XLM backbones. Accuracy is best for XLM-ACE. ‘‘ACE’’ combines Arabic (A), Chinese (C), and English (E). ‘‘M’’ stands for mode where the majority vote between the 3 heads is used. For Spanish we evaluate the models without any finetuning (Zero-Shot prediction).

between cultures.

Experimental Setup. We use the base versions of both the BERT and XLM-R models with their default tokenizers from HuggingFace. We use the standard finetuning procedure where we use the ADAM optimizer to finetune the model for 5 epochs on batches of size 32 with learning rate of 2×10^{-5} . We use cross entropy as the loss function for updating the full model parameters, including the transformer backbone. We follow the standard ArtEmis (Achlioptas et al., 2021) splits introduced in (Mohamed et al., 2022) and adopt them for both Arabic and Chinese datasets. The same training and testing images are used in all cases. For BERT models, we only evaluate on the same language as the training set because BERT tokenizers are language specific.

Baseline Results. Table 5 reports accuracy for several BERT/XLM models. There are 4 test sets, one for each language, plus ACE (a combination of 3 languages). XLM models perform better than BERT, because there is no data like more data, as well as the cross language setup used during pre-training. Interestingly, scores on the Chinese test set are higher than for English and Arabic, suggesting that Chinese captions are easier to classify. Finally, notice that XLM-ACE (XLM trained on 3 languages) outperforms other conditions, showcasing benefits of multiple languages. Note that XLM-ACE even outperforms matching conditions,

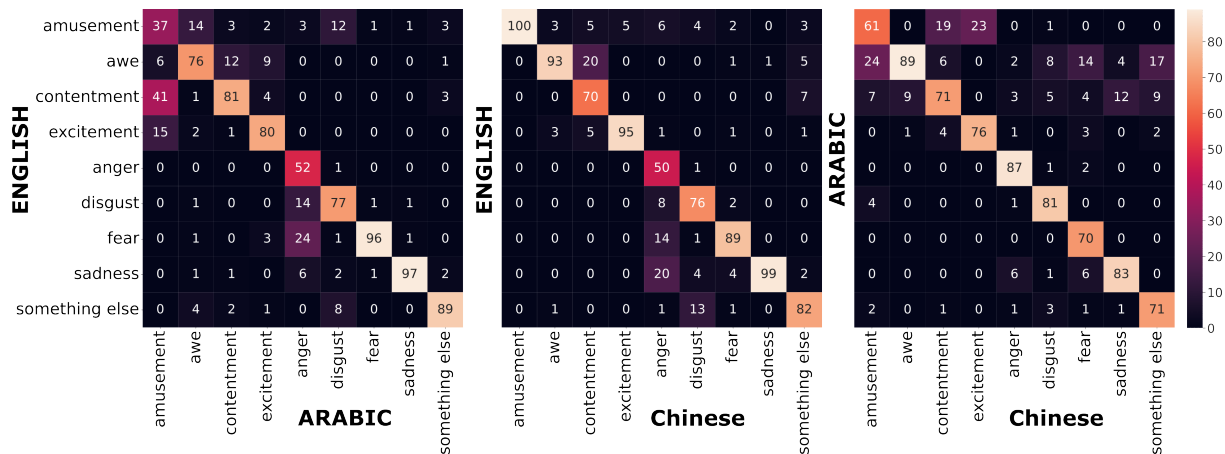


Figure 7: **Confusion Matrices** The heatmaps show confusion matrices comparing predictions from the 3-Headed Transformer with ground truth.

where training language = test language.

3-Headed Transformer Analysis. Although the 3-Headed transformer did not improve accuracy, the 3 classification heads are useful for error analysis. We feed the entire ArtELingo dataset to the model and predict 3 \hat{e} values, one for each head/language. Confusion matrices are reported in Figure 7. There is more agreement on negative emotions, and less agreement on positive emotions.

We are interested in large off-diagonal values in Figure 7, especially between positive and negative emotions. For example, Arabic **disgust** is often confused with English **amusement**.

Upon further investigation, we found nude paintings contributed $\sim 15\%$ of these confusions. Explicit content and alcohol are frowned upon in some Arabic speaking communities, as illustrated by the second and third rows of Table 6, where the label is positive in English and Chinese, but not in Arabic.

Religious symbols are also associated with large off-diagonal values in confusion matrices. The first row in Table 6 mentions Jesus and how a beautiful girl holds his cross and stomps on the devil. The annotation is positive (**awe**) in English and Arabic, but negative (**fear**) in Chinese. In China, the cross holds less meaning, and stomping on the devil is more scary than reassuring. Many symbols are associated with religion, holidays and legends that mean more in some places than others.¹⁹

While there are a few off-diagonal cells with large values, most of the large values in the confusion matrices are on the main diagonal. That is, the similarities across languages tend to dominate the differences. Consider the last row in Table 6,

¹⁹Dragons are positive in East, but negative in West.

Input Caption (Gloss)	Transformer Head		
	E	A	C
[A] A beautiful girl holding a Jesus cross stomping on the devil	Awe	Awe	Fear
[E] The woman on the ground isn't wearing any clothes	Amu.	Dis.	Amu.
[E] The man looks like he's drunk since his expression is so wired out	Amu.	Sad	Exc.
[C] Countless babies have descended into the world, giving life to the world and making people feel happy.	Cont.	Cont.	Cont.

Table 6: **Predictions from 3-Headed Transformer:** The input is a caption in Arabic (A), Chinese (C) or English (E). The first column shows the language and a gloss. The last three columns show predictions for each head (with interesting differences across heads).

which receives a positive label (**contentment**) in all 3 languages. Babies make people feel happy (nearly) everywhere. In this case, all 3 heads of our 3-headed transformer predict positive labels for this caption. For training models across multiple languages, similarities across languages may be more useful than differences.

Zero-Shot Evaluation. We use Spanish annotations in ArtELingo to evaluate models mentioned above in a zero-shot setting. The last column in Table 5 reveals two interesting relations:

1. 3-Headed-E > XLM-ACE
2. 3-Headed-E > 3-Headed-A > 3-Headed-C

The first relation suggests that 3-Heads may not perform as well as XLM when there is plenty of data, but 3-Heads may have advantages in low-

resource and zero-shot settings. 3-Heads are better for capturing interactions between languages.

The second relation suggests that language transfer may be more effective across some language pairs than others. Historically, Spanish and English are both relatively close Indo-European languages,²⁰ compared to Semitic languages such as Arabic. There has been much less contact (Thomason, 2001) between those languages and Chinese.

7 Affective Caption Generation

The previous section described baseline models for the first task: label prediction. This section will describe baseline models for the second task: affective caption generation.

To this end, we follow Achlioptas et al. (2021) and train two affective captioning models: Show, Attend, and Tell (SAT) (Xu et al., 2015) and Meshed Memory Transformer (M^2) (Cornia et al., 2020). We use *Affective Captioning Models* to refer to captioning models that generate affective captions. These captions connect the dots between input paintings and emotions.

SAT is a LSTM (Hochreiter and Schmidhuber, 1997) based captioning model with an attention module, it consists of a visual encoder and a text decoder. The visual encoder extracts visual features from an input image. The decoder then uses a stack of an attention module and LSTM recurrent unit to generate a caption autoregressively. M^2 is a transformer based model (Vaswani et al., 2017) which utilizes a pretrained Faster-RCNN (Ren et al., 2015) object detector to extract visual region features. These features are used as an input sequence to a multi-layer attention based encoder. M^2 differs from basic transformers by feeding the encoded features from all encoder layers to the cross attention module in each decoder’s layer. In order to include Emotion and Language grounding, we use a simple embedding layer to convert the emotion and language labels into feature vectors and then concatenate them to the visual features.

Experimental Setup. For both models, we use the default parameters proposed in (Achlioptas et al., 2021). We train four different versions of each model, three versions are trained on English, Arabic, and Chinese only datasets, while the fourth version is trained on the three languages combined. We then test all the models on all the languages. In

²⁰<http://www.sssscomic.com/comicpages/196.jpg>

Test Set	SAT				M^2				
	E	A	C	ACE	E	A	C	ACE	
E	B_4	6.2	0	0	6.9	8.7	0	0	8.1
	M	13.9	0	0	14.2	12.9	0	0	12.4
	R	26.5	0	0	26.5	28.0	0	0	27.4
	C	6.4	0	0	6.3	9.2	0	0	9.4
A	B_4	0	3.1	0	3.2	0	3.5	0	3.7
	M	0	30.2	0	30	0	30.9	0	30.7
	R	0	15.4	0	15.4	0	15.1	0	15.5
	C	0	7.7	0	7.7	0	7.8	0	8.0
C	B_4	0	0	11.9	10.9	0	0	8.3	8.7
	M	0	0	16.1	15.8	0	0	15.1	14.6
	R	0	0	34.3	33.6	0	0	31.1	31.1
	C	0	0	9.5	8.5	0	0	8.9	7.8
ACE	B_4	6.0	0	11.3	9.6	8.9	3.9	8.3	27.4
	M	13.5	0.42	15.2	10.5	11.8	30.8	14.8	21.2
	R	28.9	94.8	33.3	51.8	27.6	45.1	30.5	32.1
	C	2.4	0.06	3.0	2.0	3.1	14.6	3.1	5.6

Table 7: **Affective Captioning Baseline.** SAT and M^2 are trained on English (E), Arabic (A), Chinese (C), and all languages (ACE). The trained models are evaluated on a test set from each language as well as a combined test set. For metrics, we use BLEU-4 (B_4), METEOR (M), ROUGE (R), and CIDEr (C). Each row corresponds to a test set in a particular language. Meanwhile, columns correspond to model trained on a given language.

order to allow the models to work on an arbitrary languages during testing, we create our custom tokenizer which is based on xlm-roberta-base tokenizer from HuggingFace. The available tokenizer has a vocabulary of size 200K tokens which makes the training inefficient. To mitigate this, we use the same *xlm-roberta-base*²¹ tokenizer training strategy to create a tokenizer with 60K vocabulary size on ArtELingo.

Results. We report the results of our baseline models in Table 7. Models trained using all the languages perform very similarly to their language specific counterparts on every metric except for the Chinese language. This provides additional evidence that English and Arabic speaking cultures are more closely related to one another than either is to Chinese ones. In other words, English captioning models do not lose much performance when Arabic data is added to the training set and vice versa. On the other hand, Chinese models suffer when such data is added. Moreover, we also observe that for models trained on single languages, the scores on the combined test set is proportional to the language specific test sets.

²¹<https://huggingface.co/xlm-roberta-base>

8 Conclusion

This paper introduced ArtELingo, a multilingual dataset and benchmark on WikiArt images with more than 1.2M captions and emotion labels. The benchmark has diverse emotional experiences constructed over different cultures, and communicated in four languages (English, Chinese, Arabic, and Spanish). We found more agreement for some genres such as landscapes and more disagreement for other genres such as sketches. These differences are interesting and important, and far from random. Annotations for trees in Figure 1c are labeled as sadness in English and Chinese but contentment in Arabic. People are likely to feel more comfortable with what they know. People raised in countries with lush forests are likely to prefer that, whereas people brought up in less humid environments are likely to prefer that.

Towards building more socially and multi-culturally aware AI, we created baseline models for two tasks on ArtELingo: (1) emotion label prediction and (2) affective caption generation. For emotion label prediction, our best baseline model trained XLM on a combination of training data from all three languages (XLM-ACE). We also created 3-headed transformers, training three heads for three languages (Arabic, Chinese, and English) at the same time. The performance of this model is close to XLM-ACE, but generalizes better in a zero-shot experiment on Spanish. For the caption generation task, we trained two models on ArtELingo, SAT and M^2 . For English and Arabic, models on all three languages have a similar performance to language specific models, but for Chinese, it is best to train without the other languages since the performance drop is significant.

We hope our benchmark and baselines will help ease future research in visually-grounded language models that can communicate affectively with us. In addition, ArtELingo can provide empirical examples of cross-cultural similarities and differences. Sociologists and Cultural Psychologists may formulate hypotheses and conduct field studies based on ArtELingo. Data, code, and models are publicly available at www.artelingo.org/.

9 Limitations

ArtELingo’s artworks are extracted from WikiArt. Although ArtELingo is diverse in language and culture, it inherits WikiArt’s bias toward western artworks as discussed in Table 2 in §3.1. There

is room to improve the representation of certain regions of the world. Due to globalisation, people tend to follow similar trends around the world, causing others to follow their lead (for better and for worse).

Many cultures, such as Arabic, do not have a rich heritage of oil paintings. Instead, they have other forms of Art like poetry and calligraphy. Such art forms are interesting to study on their own, but mixing them with paintings is not obvious. Based on the original ArtEmis dataset, we chose WikiArt with the intent to be a continuation of their work. Also, artworks are more accessible and can be interpreted easier by different cultures compared to poetry and other art forms.

The addition of affective captions for Arabic, Chinese, as well as a small set of Spanish is a step toward cultural diversity. However, more than four regions and languages are indeed needed to cover the world. Scalability can be a challenge. However, we hope that progress can be accelerating by developing affective vision and language models that can learn with limited data for each additional language by distilling knowledge from language-only models as in (Chen et al., 2022; Alayrac et al., 2022).

ArtELingo was also collected through AMT’s online platform²². This suggests that the workers are familiar with technology and social media, imposing an influence on the data. Social media influences many concepts such as: trending news, and standards, which may lead to the presence of similarities between cultures. There have been, of course, other concerns about the use of AMT and the so-called “gig” economy and workers’ rights.

10 Acknowledgements

The authors would like to thank Baidu for soliciting Chinese annotators, all the annotators for their effort in the data collection, Eric Macedo Esparza for reviewing the Spanish dataset, and the anonymous reviewers for their valuable comments. We also would like to thank all the middle eastern universities, mainly in Egypt, who contributed to collecting the Arabic version.

This work was supported by King Abdullah University of Science and Technology (KAUST), under Award No. BAS/1/1685-01-01.

²²Chinese data collection was done on AMT as a tool with help from Baidu who helped recruit the human participants from China

References

- Malak Abdullah and Samira Shaikh. 2018. [TeamUNCC at SemEval-2018 task 1: Emotion detection in English and Arabic tweets using deep learning](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 350–357, New Orleans, Louisiana. Association for Computational Linguistics.
- Abubakar Abid, Maheen Farooqi, and James Zou. 2021. [Persistent anti-muslim bias in large language models](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 298–306.
- Lila Abu-Lughod. 1990. The romance of resistance: Tracing transformations of power through bedouin women. *American ethnologist*, 17(1):41–55.
- Panos Achlioptas, Maks Ovsjanikov, Kilichbek Haydarov, Mohamed Elhoseiny, and Leonidas Guibas. 2021. [Artemis: Affective language for visual art](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*.
- Jason Armitage, Endri Kacupaj, Golsa Tahmasebzadeh, Maria Maleshkova, Ralph Ewerth, and Jens Lehmann. 2020. Mlm: a benchmark dataset for multitask learning with multiple languages and modalities. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2967–2974.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. [Man is to computer programmer as woman is to homemaker? debiasing word embeddings](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Emanuele Bugliarello, Fangyu Liu, Jonas Pfeiffer, Siva Reddy, Desmond Elliott, Edoardo Maria Ponti, and Ivan Vulić. 2022. Iglue: A benchmark for transfer learning across modalities, tasks, and languages. *arXiv preprint arXiv:2201.11732*.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. [IEMOCAP: Interactive emotional dyadic motion capture database](#). *Language resources and evaluation*, 42(4):335–359.
- Jun Chen, Han Guo, Kai Yi, Boyang Li, and Mohamed Elhoseiny. 2022. Visualgpt: Data-efficient adaptation of pretrained language models for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18030–18040.
- Ying Chen, Wenjun Hou, Shoushan Li, Caicong Wu, and Xiaoqiang Zhang. 2020. [End-to-end emotion-cause pair extraction with graph convolutional network](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 198–207, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. 2020. [Meshed-memory transformer for image captioning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10578–10587.
- Alan S Cowen, Hillary Anger Elfenbein, Petri Laukka, and Dacher Keltner. 2019. [Mapping 24 emotions conveyed by brief human vocalization](#). *American Psychologist*, 74(6):698.
- Alan S Cowen, Xia Fang, Disa Sauter, and Dacher Keltner. 2020. [What music makes us feel: At least 13 dimensions organize subjective experiences associated with music across different cultures](#). *Proceedings of the National Academy of Sciences*, 117(4):1924–1934.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. [GoEmotions: A dataset of fine-grained emotions](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages

- 4040–4054, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mor Geva, Uri Katz, Aviv Ben-Arie, and Jonathan Berant. 2021. [What’s in your head? Emergent behaviour in multi-task transformer models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8201–8215, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Joseph Henrich, Steven J Heine, and Ara Norenzayan. 2010. The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3):61–83.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural computation*, 9(8):1735–1780.
- Carroll E Izard. 2009. [Emotion theory and research: Highlights, unanswered questions, and emerging issues](#). *Annual review of psychology*, 60:1.
- Meichun Jiao and Ziyang Luo. 2021. [Gender bias hidden behind Chinese word embeddings: The case of Chinese adjectives](#). In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 8–15, Online. Association for Computational Linguistics.
- Ronak Kosti, Jose M. Alvarez, Adria Recasens, and Agata Lapedriza. 2017. [EMOTIC: Emotions in context dataset](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Sophia Yat Mei Lee, Ying Chen, Shoushan Li, and Chu-Ren Huang. 2010. [Emotion cause events: Corpus construction and analysis](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Jennifer S Lerner, Ye Li, Piercarlo Valdesolo, and Karim S Kassam. 2015. [Emotion and decision making](#). *Annual review of psychology*, 66(1).
- Sheng Liang, Philipp Dufter, and Hinrich Schütze. 2020. [Monolingual and multilingual reduction of gender bias in contextualized representations](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5082–5093, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. [Microsoft COCO: Common objects in context](#). In *European conference on computer vision*, pages 740–755. Springer.
- Chen Liu, Muhammad Osama, and Anderson De Andrade. 2019. [DENS: A dataset for multi-class emotion analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6293–6298, Hong Kong, China. Association for Computational Linguistics.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021. [Visually grounded reasoning across languages and cultures](#). *arXiv preprint arXiv:2109.13238*.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. [A survey on bias and fairness in machine learning](#). *ACM Comput. Surv.*, 54(6).
- Erin Meyer. 2014. *The culture map: Breaking through the invisible boundaries of Global Business*. Public Affairs.
- Trisha Mittal, Puneet Mathur, Aniket Bera, and Dinesh Manocha. 2021. [Affect2MM: Affective analysis of multimedia content using emotion causality](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5661–5671.
- Youssef Mohamed, Faizan Farooq Khan, Kilichbek Haydarov, and Mohamed Elhoseiny. 2022. [It is okay to not be okay: Overcoming emotional bias in affective image captioning by contrastive data collection](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, volume abs/2204.07660.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. [SemEval-2018 task 1: Affect in tweets](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 1–17, New Orleans, Louisiana. Association for Computational Linguistics.
- Saif Mohammad and Svetlana Kiritchenko. 2018. [WikiArt emotions: An annotated dataset of emotions evoked by art](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Ara Norenzayan and Steven J Heine. 2005. Psychological universals: What are they and how can we know? *Psychological bulletin*, 131(5):763.
- Changqin Quan and Fuji Ren. 2009. [Construction of a blog emotion corpus for Chinese emotional expression analysis](#). In *Proceedings of the 2009 Conference*

- on Empirical Methods in Natural Language Processing*, pages 1446–1454, Singapore. Association for Computational Linguistics.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. [Faster R-CNN: Towards real-time object detection with region proposal networks](#). *Advances in neural information processing systems*, 28.
- James A Russell and Lisa Feldman Barrett. 1999. [Core affect, prototypical emotional episodes, and other things called emotion: dissecting the elephant](#). *Journal of personality and social psychology*, 76(5):805.
- Elvis Saravia, Hsien-Chi Toby Liu, Yen-Hao Huang, Junlin Wu, and Yi-Shin Chen. 2018. [CAREER: Contextualized affect representations for emotion recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3687–3697, Brussels, Belgium. Association for Computational Linguistics.
- Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021. [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*, pages 2443–2449.
- Carlo Strapparava and Rada Mihalcea. 2007. [SemEval-2007 task 14: Affective text](#). In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 70–74, Prague, Czech Republic. Association for Computational Linguistics.
- Salah G Thomason. 2001. *Language contact*. Edinburgh University Press.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, ukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *Advances in neural information processing systems*, 30.
- Zhongqing Wang, Yue Zhang, Sophia Lee, Shoushan Li, and Guodong Zhou. 2016. [A bilingual attention network for code-switched emotion prediction](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1624–1634, Osaka, Japan. The COLING 2016 Organizing Committee.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *arXiv preprint arXiv:1910.03771*.
- Zhongzhe Xiao, Ying Chen, Weibei Dou, Zhi Tao, and Liming Chen. 2018. [MES-P: an emotional tonal speech dataset in mandarin chinese with distal and proximal labels](#). *arXiv preprint arXiv:1808.10095*.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In *International conference on machine learning*, pages 2048–2057. PMLR.

11 Appendix

11.1 GitHub Repo

You can find the dataset and more visuals in artelingo.org or our github repo github.com/Vision-CAIR/artelingo

11.2 Pretrained BERT models

In the emotion prediction experiment, we finetune pretrained BERT models. For each language, we use a BERT model pretrained only on that language. In particular, we use “bert-base-uncased”²³ for English; “CAMEL-Lab/bert-base-arabic-camelbert-mix”²⁴ for Arabic; and “bert-base-chinese”²⁵ for Chinese.

Language specific models are finetuned on subsets of ArtELingo having captions written in the same language. On the other hand, multilingual models are pretrained “XLMroBERTa”²⁶ and they are finetuned on the whole of ArtELingo.

For each model, we finetune the pretrained model for 5 epochs. We use an ADAMW optimizer²⁷ with a learning rate of 2×10^{-5} with a linear schedule²⁸. We use cross-entropy as the loss function²⁹. Please check our GitHub repo for all of the implementation details³⁰.

11.3 Ethical Concerns

We received approval for the data collection from KAUST Institutional Review Board (IRB). The IRB requires informed consent; in addition, there are terms of service in AMT. We respected fair treatment concerns from EMNLP (compensation) and IRB (privacy). We compensated the workers well above the minimum wage (<\$1 USD/hour in Egypt and \$2.48 USD/hour in China). We paid our workers \$0.07 USD per completed task. Each

task takes on average 50 seconds to complete. In addition, we paid bonuses (mostly 30%) to workers who submitted high-quality work.

The workers were given full-text instructions on how to complete tasks, including examples of approved and rejected annotations (please refer to §11.4). Participants’ approvals were obtained ahead of participation. Due to privacy concerns from IRB, comprehensive demographic information could not be obtained.

²³<https://huggingface.co/bert-base-uncased>

²⁴<https://huggingface.co/CAMEL-Lab/bert-base-arabic-camelbert-mix-ner>

²⁵<https://huggingface.co/bert-base-chinese>

²⁶<https://huggingface.co/xlm-roberta-base>

²⁷https://huggingface.co/docs/transformers/main_classes/optimizer_schedules#transformers.AdamW

²⁸https://huggingface.co/docs/transformers/main_classes/optimizer_schedules#transformers.get_linear_schedule_with_warmup

²⁹<https://pytorch.org/docs/stable/generated/torch.nn.CrossEntropyLoss.html>

³⁰github.com/Vision-CAIR/artelingo

11.4 User Interfaces

ماهي العاطفة التي تشعر بها عندما تتمعن في اللوحة؟ (انقر هنا للاخفاء التعليمات)

تعليمات اختيار العاطفة وتكوين وصف قصير للوحة الفنية:

1. ماهي العاطفة الرئيسية التي تشعر بها من هذه اللوحة؟
2. أكتب وصف تفصيلي (ثمانية كلمات كحد أدنى) عن سبب اختيارك لهذه العاطفة مستخدماً تفاصيل دقيقة في الصورة

أمثلة وصف جيدة:

- " السماء قاتمة والظلال الطويلة مخيفة والقارب مهجور ومهترئ "
- " العلامات الحمراء على الطاولة تبدو كأنها قطرات دماء " (تصبح بكتلة تشبهت جميلة مثل هذه)
- " لون البحيرة الزرقاء متباين بشكل جميل مع القبعات البرتقالية التي يرتديها الرجال "

(أ) لا تكتب اوصاف هامشية لاتضيف اي معلومة مفيدة , مثل: " السماء حلوة " , " اللون جميلة " , يجب ان تشرح بتفصيل وتوضح اكثر

(ب) لا تبدأ الوصف بـ

- " أنا أشعر..."

(ج) لا تكتب اسم اللوحة او اي معلومات اخرى خارجية في الوصف. يجب كتابة الوصف من محتوى الصورة فقط

(د) نسخ ولصق وصف الصور او التلاعب على قوانين كتابة الوصف ستؤدي إلى الرفض

(هـ) من المهم جداً ذكر تفاصيل محتوى الصورة في الوصف

(و) الكتابة باللغة العربية فقط لا غير

شكراً لكم على وقتكم!


اختر اكثر عاطفة تشعر انها تتناسب مع اللوحة مع وصف السبب مستخدماً التفاصيل الموجودة في اللوحة

الغضب

الاشمزاز

الخوف

الحزن



الحماس

الشمعة

الرضا

التفوق

أخرى

أكتب في ثمانية كلمات او أكثر عن سبب اختيارك لعاطفة أخرى

اكتب الوصف هنا

You must ACCEPT the HIT before you can submit the results.

Figure 8: Arabic Interface

您对此作品的感受是什么？请描述原因。

指导语——怎样选择您的感受并给出解释:

1. 您对该作品有何主要感受？如果有多种感受，则选择最强烈的那一个。(只能选择一种)
2. 请根据作品的特定细节，具体解释您为什么会有此感受。(仅限中文,不少于8个字)


范例——好的描述:

- “天空看起来很阴暗，阴影部分吓人”
- “桌子上的红色标记看起来像血滴”(欢迎使用类比比喻!)
- “蓝色的湖水与男人橙色的帽子相映成趣”

- (a) 请不要使用笼统的描述，例如“有趣”、“颜色很棒”等。这种解释不够明确。
- (b) 请不要以“我觉得/感觉到”等方式开头。
- (c) 请不要在您的解释中写出作品名称或引用其他外部资料。您的解释只需提到作品的细节。
- (d) 复制粘贴将导致您被拉黑!
- (e) 请务必在您的解释中提到作品的细节!

感谢您抽出时间完成任务!

请选出最符合您对当前作品感受的一种情感，并解释您选择这一情感的原因。



其它情感

请解释为什么该作品使您产生这种感受 (仅限中文,不少于8个字) **嫌恶**

请在此处输入您的描述

Submit

Figure 9: Chinese Interface

¿Cómo te hace sentir esta pintura? ¡Descríbelo! (haz clic aquí para colapsar esta pestaña)

Normas:

1. ¿Cómo te hace sentir **principalmente** esta pintura? (elige uno botón)
2. Da una descripción detallada (**al menos 8 palabras**) sobre **POR QUÉ** te sientes así, basado en detalles **ESPECÍFICOS** de la pintura.

Ejemplos de BUENAS descripciones:

- "el cielo se ve sombrío y las sombras dan miedo"
- "las marcas rojas en la mesa parecen gotas de sangre"(nos gustan las analogías!)"
- "el azul del lago contrasta bien con los sombreros naranjas de los hombres"

(a) No utilices descripciones poco informativas , como: "es divertido", "bonitos colores"; por ejemplo: NO explicar POR QUÉ de una manera específica.

(b) No empieces tu frase con








- "Me siento..."
- "En la pintura ..."

(c) Si no sientes "nada" o te sientes "aburrido", aún tienes que explicar POR QUÉ te sientes así.

¡Si no dominas el español, no aceptes este HIT!

¡Muchas gracias por tu arduo trabajo!

¿Cómo te hace sentir esta pintura? ¡Descríbelo!



Algo más

Da una descripción detallada (al menos 8 palabras) sobre **POR QUÉ** te sientes **Algo más**

Escribes aquí

You must ACCEPT the HIT before you can submit the results.

Figure 10: Spanish Interface