

# Leveraging Domain Knowledge for Inclusive and Bias-aware Humanitarian Response Entry Classification

Nicolò Tamagnone<sup>1,2</sup>, Selim Fekih<sup>1</sup>, Ximena Contla<sup>1</sup>, Nayid Orozco<sup>1</sup>, Navid Rekabsaz<sup>3</sup>

<sup>1</sup>Data Friendly Space

<sup>2</sup>ISI Foundation, Turin, Italy

<sup>3</sup>Johannes Kepler University Linz, LIT AI Lab, Austria

<sup>1</sup>{nico,selim,ximena,nayid}@datafriendlyspace.org, <sup>3</sup>navid.rekabsaz@jku.at

## Abstract

Accurate and rapid situation analysis during humanitarian crises is critical to delivering humanitarian aid efficiently and is fundamental to humanitarian imperatives and the Leave No One Behind (LNOB) principle. This data analysis can highly benefit from language processing systems, e.g., by classifying the text data according to a humanitarian ontology. However, approaching this by simply fine-tuning a generic large language model (LLM) involves considerable practical and ethical issues, particularly the lack of effectiveness on data-sparse and complex subdomains, and the encoding of societal biases and unwanted associations. In this work, we aim to provide an effective and ethically-aware system for humanitarian data analysis. We approach this by (1) introducing a novel architecture adjusted to the humanitarian analysis framework, (2) creating and releasing a novel humanitarian-specific LLM called HumBERT, and (3) proposing a systematic way to measure and mitigate biases. Our results show the better performance of our approach on zero-shot and full-training settings in comparison with strong baseline models, while also revealing the existence of biases in the resulting LLMs. Utilizing a targeted counterfactual data augmentation approach, we significantly reduce these biases without compromising performance.

## 1 Introduction

The efficiency of humanitarian aid during crises, be it natural disasters, conflicts, or pandemics such as COVID-19, relies heavily on quick and accurate analysis of relevant data within the first 72 hours of a disaster. During these crucial hours, humanitarian response analysts from international organizations and non-governmental organizations (NGOs) sift through a large volume of data, related to the crisis, to obtain a comprehensive understanding of the situation. A significant amount of this information is from secondary sources such as reports, news articles, and text-based data. During this phase, the process of analysis involves extracting critical information from those sources and categorizing it based on established humanitarian frameworks and guidelines. The analysis of this infor-

mation plays a crucial role in determining the appropriate relief efforts to implement, and can highly benefit from technological solutions based on natural language processing (NLP) and Deep Learning. In line with the Leave No One Behind (LNOB) principle – the central and transformative promise of the 2030 Agenda for Sustainable Development Goals – such systems can assist humanitarian organizations with efficient and accurate information analysis while supporting the generation of evidence and data disaggregation that goes beyond gender, geography, and age.

To enable training and benchmarking NLP models in the humanitarian domain, the recently released HUMSET dataset [Fekih *et al.*, 2022] provides a rich resource of humanitarian document analysis and entry (text excerpt) classification. The dataset originated from the Data Entry and Exploration Platform (the DEEP)<sup>1</sup> developed and maintained by Data Friendly Space, a U.S.-based international non-governmental organization (INGO),<sup>2</sup> in collaboration with several international organizations.<sup>3</sup> HUMSET is a multilingual dataset, covering several disasters around the globe, where the informative entries of each document are extracted by humanitarian analysts, and classified/tagged according to a designated humanitarian information analysis framework. Fekih *et al.* [2022] further use this data to fine-tune large pre-trained language models (LLMs) such as XLM-RoBERTa [Conneau *et al.*, 2020], observing the overall effectiveness of these baseline approaches. However, we spot several challenges and concerns regarding this approach, especially when one aims to utilize these models in the real workflow of humanitarian response. These challenges are due to the inherent complexity of the humanitarian framework, the lack of training data on new crises, and the presence of societal biases and unwanted associations in the models. Our aim in this work is to address these challenges and provide an effective and ethically-aware solution for humanitarian response entry classification, as explained in the following.

The first challenge is rooted in the fact that the various cat-

<sup>1</sup><https://thedeep.io/>

<sup>2</sup><https://datafriendlyspace.org/>

<sup>3</sup>Such as the International Federation of Red Cross (IFRC), the United Nations High Commissioner for Refugees (UNHCR), and the United Nations Office for the Coordination of Humanitarian Affairs (UNOCHA), DEEP is funded by USAID’s Bureau for Humanitarian Assistance (BHA).

egories of the humanitarian framework share topical conceptual similarities, while each category also requires its own learning capacities. Learning all categories together in one model – as done in the baseline approach – leads to overwriting of information and catastrophic forgetting, as similarly reported in the context of multi-task learning [Parisi *et al.*, 2019; Chen *et al.*, 2020]. We approach this challenge by taking into account the composition of the humanitarian framework and introduce a combinatorial model consisting of several shared layers across categories/tasks, as well as task-specific Transformer [Vaswani *et al.*, 2017] layers. Our experiments, using various backbone LLMs, show the overall better performance of our proposed architecture in comparison with the baseline approaches.

The second challenge regards the (expected) weak performance of the models in the cases with no or little training data available (zero-/few-shot setting). This scenario is particularly common in the humanitarian domain, as each crisis – while commonly happening abruptly and unexpectedly – often differs from the previous incidents. We approach this challenge by creating and releasing a novel LLM called HumBERT, tailored to the characteristics of the humanitarian domain. Our experiments show that utilizing HumBERT significantly improves the performance on both zero-shot entry classification using prompting, and the standard fully-supervised setting when compared with the various baseline models using generic LLMs.

Finally, the third challenge concerns the issue of encoding societal biases and stereotypes in LLMs [Sheng *et al.*, 2019; Blodgett *et al.*, 2020; Rekabsaz *et al.*, 2021b], and carrying and exaggerating these biases in downstream tasks [Rekabsaz and Schedl, 2020; Rekabsaz *et al.*, 2021a; Kumar *et al.*, 2023; Hauenberger *et al.*, 2023]. This issue is particularly critical, as a central principle in the humanitarian and development domains is to support minorities and underrepresented populations, in accordance with the LNOB principle, which might be neglected by a biased model. To address this challenge, we first create the HUMSETBIAS dataset, a subset of HUMSET consisting of the data points that can potentially cause or reflect the biases with respect to gender and specific countries. Using HUMSETBIAS, we then measure the gender/country bias in a model as the discrepancies of the model’s prediction probabilities for a data point and its counterfactual form. Our results reveal the existence of a considerable amount of biases in the models and the effects of these biases on the decisions of the model. We further approach these biases using a counterfactual data augmentation (CDA) [Zhao *et al.*, 2018; Lu *et al.*, 2018] approach and achieve significant mitigation in the biases without any decrease in overall performance.

The remainder of the work is organized as follows: in Section 2, we provide the background on the humanitarian domain and review the related work. Section 3 explains the proposed architecture, followed by Section 4 to introduce the novel LLM and dataset resources. We explain the setup of our experiments and report their results in Section 5. Our code and models are available at <https://github.com/the-deep-nlp/bias-aware-humanitarian-entry-classification>.

Single-leveled Tasks:	
<b>Sectors (11)</b>	Agriculture, Cross-sector, Education, Food Security, Health, Livelihoods, Logistics, Nutrition, Protection, Shelter, WASH (Water, Sanitation & Hygiene)
<b>Pillars 1D (7)</b>	Context, COVID-19, Displacement, Humanitarian Access, Information & Communication, Casualties, Shock/Event
<b>Pillars 2D (6)</b>	Capacities & Response, Humanitarian Conditions, Impact, At Risk, Priority Needs, Priority Interventions
Two-leveled Tasks:	
<b>Subpillars 1D (36)</b>	<i>Casualties:</i> Dead, Injured, Missing <i>Context:</i> Demography, Economy, Environment, Legal & Policy, Politics, Security & stability, Socio-cultural <i>COVID-19:</i> Cases, Contact tracing, Deaths, Hospitalization & care, Prevention campaign, Research and outlook, Restriction measures, Testing, Vaccination <i>Displacement:</i> Intentions, Local integration, Pull factors, Push factors, Type/numbers/movements <i>Humanitarian Access:</i> People facing humanitarian access constraints/humanitarian access gaps, Physical constraints, Population to relief, Relief to population <i>Information and Communication:</i> Communication means and preferences, Information challenges and barriers, Knowledge and info gaps (hum), Knowledge and info gaps (pop) <i>Shock / Event:</i> Hazard & threats, Mitigating factors, Type and characteristics, Underlying/aggravating factors
<b>Subpillars 2D (19)</b>	<i>At Risk:</i> Number of people at risk, Risk and vulnerabilities <i>Capacities &amp; Response:</i> International response, Local response, National response, People reached/response gaps, Red cross/red crescent <i>Humanitarian Conditions:</i> Coping mechanisms, Living standards, Number of people in need, Physical and mental well-being <i>Impact:</i> Driver/aggravating factors, Impact on people, Impact on systems, services, and networks, Number of people affected <i>Priority Needs:</i> Expressed by humanitarian staff, Expressed by population <i>Priority Interventions:</i> Expressed by humanitarian staff, Expressed by population

Table 1: Overview of the humanitarian analysis framework. The number of tags present for each category is reported in parentheses.

## 2 Background and Related Work

Previous work introduces a number of resources for humanitarian data analysis. The HUMSET dataset [Fekih *et al.*, 2022] – used in the work at hand – is created by humanitarian analysts on official documents and news from the most recognized humanitarian agencies. These data are organized according to the proposed analytical framework in the humanitarian domain, which defines a set of guidelines and structures to collect and organize data. This humanitarian analytical framework – shown in Table 1 – consists of five categories of Sectors, Pillars/Subpillars 1D, and Pillars/Subpillars 2D. The Subpillars 1D and Subpillars 2D categories are organized into a two-leveled tree hierarchy (Pillars to Subpillars).

As other resources in this domain, Alam *et al.* [2021] provide CrisisBench, a benchmark of social media data labeled with the relevance of the text to humanitarian action, beneficial for filtering social media content as an additional resource of situational awareness. Similarly, Imran *et al.* [2016] release human-annotated Twitter corpora accumulated during various crises in the span of 2013 to 2015, and [Adel and Wang, 2020] and later [Alharbi and Lee, 2021] publish Arabic Twitter classification datasets for crisis events.

Using deep learning methods, Yela-Bello *et al.* [2021]

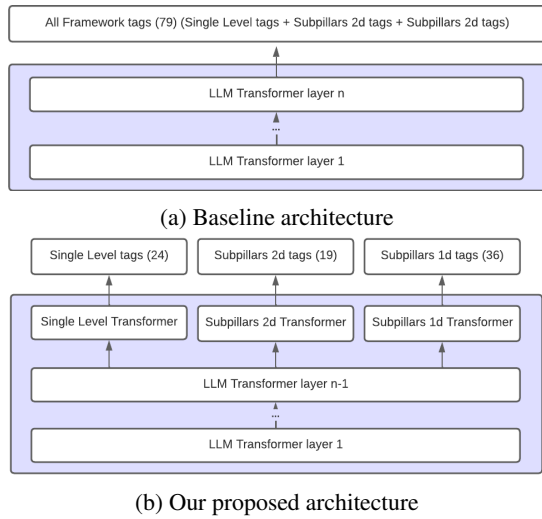


Figure 1: Model architectures used in this study

study neural extractive summarization techniques to spot informative excerpts in humanitarian data. More recently, Lai et al. [2022] propose a hybrid named-entity recognition model that utilizes a set of features to extract the flooding information and risk reduction projects from newspapers. Our work contributes to this line of research by introducing a novel humanitarian-specific LLM leveraged in zero-shot and low-resource settings and approaching underlying biases.

### 3 Model Architecture

In this section, we explain the design of our model architecture. Our aim is to provide an effective (multi-label) entry classification model, which takes into account the complexity of the analytical framework consisting of categories with many labels (see Table 1), as well as the possible inter-category conceptual relationships. In the following, we first explain the baseline entry classification model and then extend it to our proposed approach. Both architectures leverage the power of pre-trained LLMs as the backbone. The schematic view of the two architectures is shown in Figure 1.

The baseline approach follows the standard classification approach of LLMs [Devlin *et al.*, 2019], namely by first encoding the input sequence into a vector embedding, used for predicting the labels. Since the classification task is multi-label, a linear classification head is defined to predict each class (79 heads in total). The baseline architecture in fact does not take into account the hierarchy or the relationships of the framework’s labels and treats all the categories with the same encoded embedding.

Our proposed architecture aims to account for these aspects. Similar to the baseline, our approach shares  $N - 1$  Transformer layers of the LLM across all the categories, while the  $N$ th layer is replicated three times, forming three sub-layers. The first sub-layer is assigned to the single-leveled tasks namely Sectors, Pillars 1D, and Pillars 2D, and each of the other two sub-layers is dedicated to each of the two-leveled tasks, Subpillars 1D and Subpillars 2D. A linear classification head is then defined on top of each of the

Transformer sub-layers, predicting only the labels belonging to the corresponding task. The resulting architecture is in fact a combination of the shared parameters and specific components and follows the relations and hierarchy of the analytical framework’s label space.

## 4 Novel Humanitarian Resources

In this section, we present a suite of innovative resources designed to improve humanitarian response classification. Our resources include HumBERT, a new LLM model that incorporates insights of the latest research on language modeling specifically trained on a vast corpus of humanitarian text data, and HUMSETBIAS, a dataset for analyzing and mitigating biases in humanitarian response classification.

### 4.1 HumBERT

The advancement of LLMs in recent years have provided immense benefits to the field of NLP, particularly by enabling effective transfer learning approaches [Malte and Ratadiya, 2019]. LLMs are trained on vast amounts of general text data, and have shown remarkable performance across a wide range of language downstream tasks. However, such general LLMs may not perform as well as on specific domains, containing particular topics, specialized vocabulary, and complex concepts not well-represented in general language texts. To address this challenge, recent works focus on developing pre-trained language models for different domains, as for instance SciBERT [Beltagy *et al.*, 2019], BioBERT [Lee *et al.*, 2020], and FinBERT [Araci, 2019] are trained on scientific, biomedical and financial textual data, respectively. In the work in hand, we introduce HumBERT, a multilingual large language model trained on a large corpus of humanitarian text data, including reports, news articles, and other materials related to humanitarian crises and aid efforts.

**Humanitarian Corpus Data Collection** The text data for creating the humanitarian corpus is collected from three main sources: ReliefWeb,<sup>4</sup> UNHCR Refworld,<sup>5</sup> and Europe Media Monitor News Brief (EMM).<sup>6</sup> ReliefWeb is a humanitarian information portal that provides news, reports, maps, and data related to crises. UNHCR Refworld is a database mainly focused on refugee-related documentation, including country of origin information, policy documents, and legal analysis. EMM is a real-time news monitoring tool that aggregates articles and media posts, with a specific focus on issues related to the European Union. While the main focus of Reliefweb and UNHCR Refworld is related to humanitarian crises and refugees, EMM broadcasts a wider range of news articles. We hence apply a filtering criterion on themes such as *Crisis Response*, *Humanitarian Aid*, *Food Security & Food Assistance* and *Asylum*, to retrieve the humanitarian-relevant documents.

We collect the mentioned text articles, resulting in a collection of  $\sim 2$ M documents, of which 90% are published articles from 2003 to 2022. The documents are in three languages of English (82%), French (10%), and Spanish (8%). We extract

<sup>4</sup><https://reliefweb.int/>

<sup>5</sup><https://www.refworld.org/>

<sup>6</sup><https://emm.newsbrief.eu/>

Female	Male	Neutral
she	he	they
woman	man	person, individual
women	men	persons, individuals
mother	father	person, individual
mothers	fathers	persons, individuals
girl	boy	child
girls	boys	children
her	his	their
her	him	them
female	male	person, individual
females	males	persons, individuals
wife	husband	person, individual
wives	husbands	persons, individuals

Table 2: Keywords indicating gender information, and the corresponding mappings for creating counterfactual samples.

the text from the documents (web pages, related PDF attachments, etc.). The documents have on average 26 sentences. We keep the structure of the text consistent with one of the paragraphs in the original document. We execute text cleaning and simple pre-processing of the text (details available in the code). After concatenating all text together, the resulting corpus consists of  $\sim 50M$  sentences with approximately 2B tokens. We further process the corpus to remove personal information such as e-mails and phone numbers.

**Training** HumBERT is a domain-specific LLM that has been fine-tuned on our humanitarian corpus, starting from the checkpoint of XLM-RoBERTa<sub>Base</sub> [Conneau *et al.*, 2020]. Starting from this model provides us the advantage of its pre-trained multilingual capabilities, which is then further fine-tuned to capture the linguistic nuances specific to the humanitarian domain. We train HumBERT for 3 epochs with 10% of the corpus used as the validation set. We further fine-tune the model based on its masked language modeling objective, using the training pipeline provided by the HuggingFace library [Wolf *et al.*, 2020].<sup>7</sup> The training is performed on a single v2-8 TPU.<sup>8</sup>

## 4.2 HUMSETBIAS

As described in Fekih *et al.* [2022] and also touched upon in the previous sections, HUMSET dataset is created by the teams of humanitarian experts in various international organizations, through annotating the relevant entries of a text document according to the analysis framework. However, the distribution of the annotations across the labels/categories is influenced by the types of resources and the information involved. Given the diverse range of humanitarian crises and topics covered by HUMSET, the annotations might be skewed toward the context in which the contents are produced. As a result, the annotated data – and hence the models trained on the dataset – may exhibit biases towards attributes such as

<sup>7</sup>[https://github.com/huggingface/transformers/blob/main/examples/pytorch/language-modeling/run\\_mlm.py](https://github.com/huggingface/transformers/blob/main/examples/pytorch/language-modeling/run_mlm.py)

<sup>8</sup>Research supported with Cloud TPUs from Google’s TPU Research Cloud (TRC)

Attribute	Bias label	Train	Val.	Test	All
Gender	Female	1,173	92	80	1,345
	Male	604	57	65	726
	Neutral	15,271	1,443	1,524	18,238
	Sum	17,048	1,592	1,669	20,309
Country	Venezuela	261	20	13	294
	Syria	1,686	182	187	2,055
	Sum	1,947	202	200	2,349

Table 3: Statistics of the HUMSETBIAS dataset for each bias type, as subsets of the train/validation/test sets of HUMSET.

gender and country (geographical), which could impact the accuracy and fairness of downstream applications.

To enable measuring and (potentially) addressing these biases, we introduce HUMSETBIAS, a subset of the English part of the HUMSET dataset, created by searching for specific sensitive English keywords related to genders and countries within the annotated text. In addition, we extended this subset by incorporating *targeted* counterfactual [Roese and Olson, 1995] samples, generated by modifying the original entries in order to create the altered versions of each text with gender/country information. The purpose of HUMSETBIAS is to provide a more targeted resource for analyzing and addressing potential biases in humanitarian data and to enable the development of accurate and bias-aware NLP applications in the humanitarian sector. In the following, we explain the process of creating HUMSETBIAS, how it is extended with targeted counterfactual samples.

**Dataset Creation** To study country bias, we select Syria and Venezuela, given their frequent occurrence in HUMSET and also the fact that these countries encompass distinct types of humanitarian crises, namely the Syrian civil war<sup>9</sup> and the Venezuelan socioeconomic and political crisis<sup>10</sup>. Through keyword search, we selected the data points that contain the name of these countries. We further aim to ensure that the selected entries are focused on countries, and do not contain other information, for instance about the nationalities, population, or geographical locations inside the countries. To this end, we exclude entries that mention any nationality (including Venezuelan and Syrian), other countries, and all cities of the chosen countries. To identify the entries with potential biases towards genders, we create a list of keywords indicating female, male, and neutral terms, reported in Table 2. We use these keywords to select the gender-related subset of entries. After reviewing this subset, we exclude the entries that contain concepts specific to one gender (such as the ones with words like *pregnant* or *lactating*), as well as the entries referring to more than one gender. Table 3 reports the statistics of the resulting potentially biased excerpts according to the train, validation, and test sets, selected from the corresponding sets in HUMSET.

**Targeted Counter-factual Samples** Various bias measurement and mitigation approaches can benefit from the exist-

<sup>9</sup>[https://en.wikipedia.org/wiki/Syrian\\_civil\\_war](https://en.wikipedia.org/wiki/Syrian_civil_war)

<sup>10</sup>[https://en.wikipedia.org/wiki/Crisis\\_in\\_Venezuela](https://en.wikipedia.org/wiki/Crisis_in_Venezuela)

LLM	Sectors		Pillars 1D		Subpillars 1D		Pillars 2D		Subpillars 2D		Avg.	
	Prec.	F1	Prec.	F1	Prec.	F1	Prec.	F1	Prec.	F1	Prec.	F1
m-BERT	0.16	0.10	0.18	0.16	0.07	0.05	0.25	<b>0.21</b>	0.08	0.05	0.15	0.11
XLm-R	0.27	0.23	0.24	0.21	0.11	0.09	<b>0.26</b>	0.18	0.09	0.06	0.19	0.15
HumBERT	<b>0.28</b>	<b>0.24</b>	<b>0.27</b>	<b>0.24</b>	<b>0.13</b>	<b>0.10</b>	0.22	0.17	<b>0.10</b>	<b>0.08</b>	<b>0.20</b>	<b>0.17</b>

Table 4: Evaluation results of zero-shot classification on the five categories/tasks of the HUMSET dataset according to Precision and F1 metrics. ‘‘Avg.’’ reports the (macro) average of the results over the tasks. The best results in each task-metric are shown in bold.

LLM	Model	Sectors		Pillars 1D		Subpillars 1D		Pillars 2D		Subpillars 2D		Avg.	
		Prec.	F1	Prec.	F1	Prec.	F1	Prec.	F1	Prec.	F1	Prec.	F1
m-BERT	BASE	0.75	0.77	0.68	<b>0.70</b>	0.51	0.51	0.7	0.69	0.54	0.48	0.71	0.72
	OURS	0.76	0.76	0.65	0.68	0.49	0.51	0.69	0.69	0.48	0.47	0.7	0.71
XLm-R	BASE	0.74	0.76	<b>0.70</b>	0.68	0.48	0.49	0.69	0.68	0.51	0.47	0.71	0.71
	OURS	0.75	0.77	0.68	0.69	0.5	0.51	0.69	0.7	0.51	0.48	0.71	0.72
HumBERT	BASE	0.76	<b>0.78</b>	0.69	<b>0.70</b>	0.52	0.53	0.7	0.7	0.52	0.5	<b>0.72</b>	<b>0.73</b>
	OURS	<b>0.77</b>	<b>0.78</b>	0.69	<b>0.70</b>	<b>0.57</b>	<b>0.54</b>	<b>0.71</b>	<b>0.71</b>	<b>0.55</b>	<b>0.51</b>	<b>0.72</b>	<b>0.73</b>

Table 5: Evaluation results on fully training the models on the HUMSET dataset. The best results across LLMs are shown in bold.

tence of counterfactual samples, namely when the protected attribute of a given data point is represented in other possible forms. To enable this research, we additionally provide targeted counterfactual samples for each data point in HUMSETBIAS dataset. In particular for gender attribute, we replace each female/male/neutral word with the two other cases (for instance female  $\rightarrow$  {male, neutral}), according to the mapping in Table 2. We follow a similar procedure for the country-related data points by selecting Canada as a contrastive (neutral) case, as this country is not associated with any crisis in the HUMSET dataset. The counterfactual samples are hence created by replacing each source country in each data entry with the other source country and Canada, namely Syria  $\rightarrow$  {Venezuela, Canada}, and Venezuela  $\rightarrow$  {Syria, Canada}. The data of targeted counterfactual samples is provided along with the HUMSETBIAS dataset in our repository.

## 5 Results and Analysis

To conduct our experiments, we use the HUMSET dataset for training and evaluating the models on the classification task for the five humanitarian categories (see Section 2). We utilize three LLMs, namely m-BERT [Devlin *et al.*, 2019], XLm-RoBERTa [Conneau *et al.*, 2020], and HumBERT. Each LLM is used as the backbone for both the baseline architecture and our proposed one (see Section 3), referred to as BASE and OURS, respectively. We fine-tune the hyperparameters on the validation set and report the evaluation results on the test set. Details of the models’ hyperparameters are reported in Appendix A in supplementary materials<sup>11</sup>.

In the following, we first explain our zero-shot classification method and the achieved results. We then report the performance of the models on the full-training scenario. Finally,

<sup>11</sup>Supplementary materials are available at <https://github.com/the-deep-nlp/bias-aware-humanitarian-entry-classification>

leveraging the HUMSETBIAS dataset we describe our method to measure the existence of gender/country bias in the models and report the results before and after applying the CDA bias mitigation method.

### 5.1 Zero-shot Learning Results

We utilize prompting to evaluate the abilities of the LLMs (in comparison with HumBERT) on entry classification zero-shot scenario. Previous studies have utilized prompts to perform downstream tasks by directly exploiting LLM objectives, showing promising results on few-/zero-shot settings [Schick and Schütze, 2021; Liu *et al.*, 2021]. As a common prompting method, the input text data point followed by a template with a masked token is given to an LLM to retrieve the probabilities of possible tokens to fill the mask. In our experiment, we utilize the simple template of ‘‘It is about [MASK]’’, appended at the end of each input entry. To perform classification concerning the labels of the humanitarian tasks (see Table 1), we create a *verbalizer* to map the predicted [MASK] token to the labels. Previous work proposes various ways to formulate a verbalizer, namely manually [Schick and Schütze, 2021], manually with pre-existing knowledge bases extensions [Hu *et al.*, 2022], or through automatic searching methods [Gao *et al.*, 2021] applicable in presence of training and validation examples. In our experiments, we follow the first approach and build our verbalizer manually, by associating each label — whose class name is already semantically informative — with a set of words with similar meanings or semantic relation. For example, the Sector label *Agriculture* is associated with the terms *farming*, *cultivation*, *land*, and *farm*, in addition to the word *agriculture*. The process results in an average of 4.6 words per label.

Table 4 reports the results of the zero-shot classification on the five categories/tasks of the humanitarian analytical framework. The zero-shot prompting inference is performed on the test set of HUMSET for the three LLMs (without any training

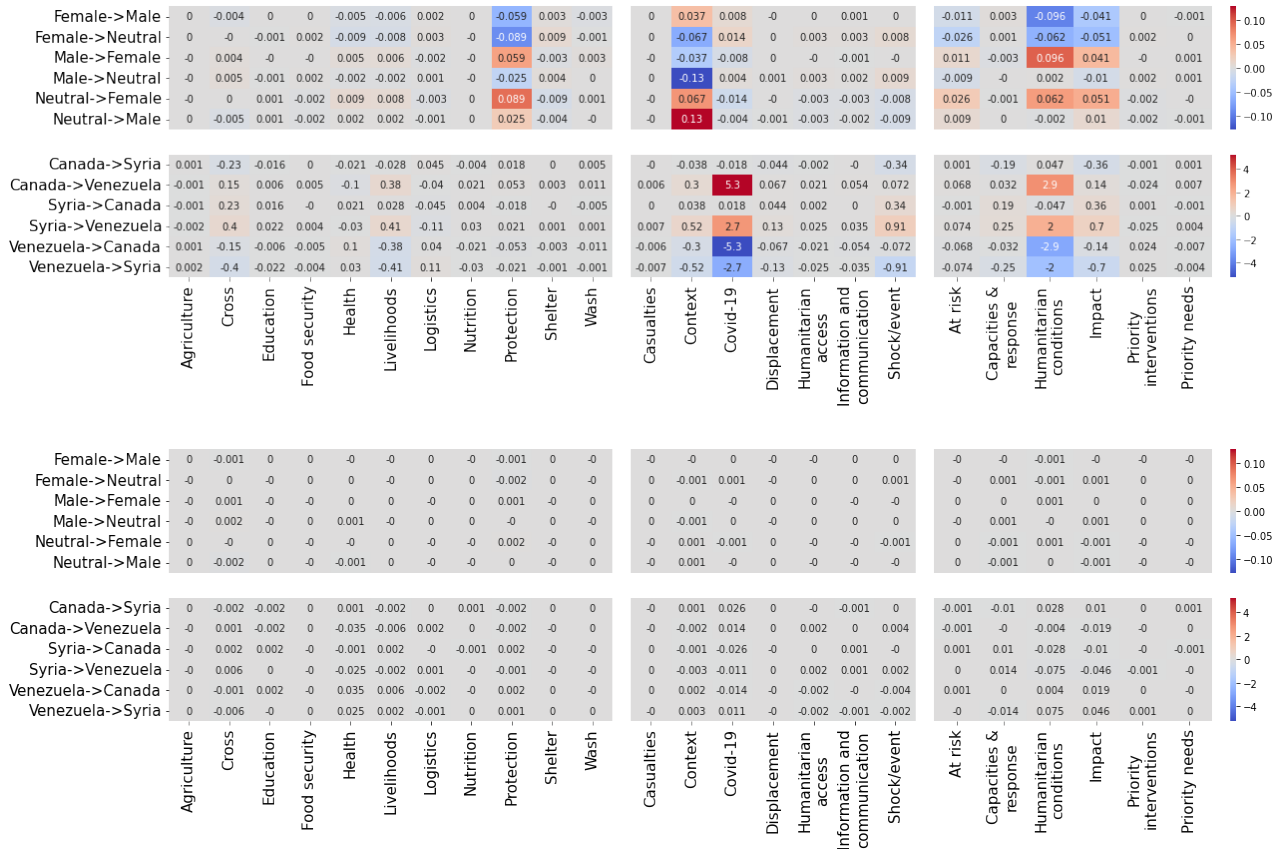


Figure 2: The results of Tag-Shift bias metric for OURS architecture using the HumBERT as the backbone. (Top) Original model without debiasing; (Bottom) Counterfactual debiasing.

and validation). Based on the results, we observe that overall our domain-specific HumBERT outperforms the other two LLMs except for Pillars 2D. Generally, the XLM-R model shows better results than m-BERT, while this model is further improved by being fine-tuned on the humanitarian text data in HumBERT. One special case is the overall weak performance of all models on the Subpillars 1D/2D categories. We assume that this is due to the complexity associated with the semantics of the labels of these tasks, and the limitation of the manual verbalizer to comprehensively address the diverse facets of the meaning associated with these labels. We consider further investigation in this area as a line of research in our future work.

## 5.2 Full-training Performance

We continue with evaluating the LLMs and the underlying model architectures (discussed in Section 3) on the fully supervised classification setting. The results are reported in Table 5, where the best results per LLM are shown in bold, and the overall best results are indicated with underlining.

As noted, our proposed architecture overall shows higher performance, particularly in the F1-scores. It is important to note that the results particularly increase for the Subpillars 1D task (especially for XLM-R and HumBERT back-

bones), which is the category with the largest number of classes (36 tags). This observation supports the benefits of using the combination of shared and task-specific Transformer layers in our proposed architecture, to reduce the sparsity of classification heads without losing the shared knowledge across the categories. Comparing across the LLMs, HumBERT achieves the best performance in all the categories.

## 5.3 Bias Measurement and Mitigation

To measure the bias of the models, we use HUMSETBIAS to calculate the sensitivity of the trained models to the changes in gender and countries keywords concerning the changes in the predicted probabilities of the labels. A non-biased model should not be sensitive to such changes so that the predicted probabilities should remain the same when the respective words of gender/country are swapped.

To this end, our bias measurement is conducted on an extended version of the test sets, where each test set also includes the counterfactual samples of each data point (described in Section 4.2). Using this data, we first define the P-Shift metric as the shift in the predicted probability of the entry  $x$  concerning the tag  $t$  with the bias label  $m$  when replacing the label with  $n$ , formulated below:

$$P\text{-Shift}(x, t, m \rightarrow n) = (P(x_n, t) - P(x_m, t)) \times 100 \quad (1)$$

LLM	Model	Tasks Avg.↑		Bias↓	
		Prec.	F1	Gender	Country
m-BERT	BASE	0.71	0.72	0.57	9.79
	OURS	<b>0.72</b>	0.72	0.48	6.37
	BASE-CDA	0.7	0.71	0.02	0.17
	OURS-CDA	<b>0.72</b>	0.72	0.02	0.12
XLM-R	BASE	0.71	0.71	0.67	8.69
	OURS	0.71	0.72	0.4	3.67
	BASE-CDA	0.71	0.72	0.05	0.14
	OURS-CDA	<b>0.72</b>	0.72	<b>0.01</b>	<b>0.11</b>
HumBERT	BASE	<b>0.72</b>	<b>0.73</b>	0.71	8.37
	OURS	0.71	<b>0.73</b>	0.28	6.52
	BASE-CDA	<b>0.72</b>	<b>0.73</b>	0.02	<b>0.11</b>
	OURS-CDA	<b>0.72</b>	<b>0.73</b>	<b>0.01</b>	0.13

Table 6: Results of Overall-Shift bias metric, and average tasks performance (taken from Table 5). The best results per LLM are shown in bold.

Here,  $P(x_m, t)$  is the prediction probability of  $x$  on the tag  $t$  in the form of bias label  $m$ , and  $P(x_n, t)$  is the probability on the same tag  $t$  of a corresponding counterfactual sample of  $x$  swapped to bias label  $n$ . The values are multiplied by 100 for improving readability. The total predicted probability shift from the attribute  $m$  to the attribute  $n$  for a tag  $t$  is therefore calculated as the median of the probability shifts:

$$\text{Tag-Shift}(t, m \rightarrow n) = \underset{x \in X_m}{\text{Median}} \left\{ \text{P-Shift}(x, t, m \rightarrow n) \right\} \quad (2)$$

where  $X_m$  refers to the set of data points with the bias label  $m$ . In our formulations, we choose the median (instead of the mean) to mitigate the effect of the outliers in P-Shift in the final statistic. We should note that since the test sets are augmented with counterfactual forms concerning each bias label,  $X_m$  has the same size across various bias labels  $m$  (size of the full test dataset). Figure 2 top reports the results of Tag-Shift for both bias attributes, calculated on our proposed architecture with HumBERT as the backbone LLM, for all  $m \rightarrow n$  cases of the single-leveled tags. The results without debiasing (top) reveal the existence of both gender and country biases in the model, particularly to specific tags. For example, in the Covid-19 Pillars 1D, Venezuela  $\rightarrow$  Canada and Venezuela  $\rightarrow$  Syria are negative (-5.3 and -2.7), indicating that the model assigns higher probabilities on this tag for the data points related to Venezuela compared to Syria and Canada.

**Bias Mitigation with Counterfactual Data Augmentation (CDA)** To conduct bias mitigation with CDA, we extend the training set of HUMSETBIAS data with the respective counterfactual data points concerning both gender and country. This results in new training set with 166,990 data points, which extends the original 129,268 training instances of HUMSET. Figure 2 bottom shows the results of applying CDA. The results of other combinations are provided in Appendix B in supplementary materials. Looking at the results, we observe a significant and consistent decrease in the degree of both gender and country bias. These results show the effectiveness of balancing the training data concerning the

protected attributes.

**Overall Probability Shift** The Tag-Shift metric provides useful information regarding the bias of a specific tag according to a specific counterfactual transition. However, we are also interested in an overall metric, enabling straightforward comparison across different models. We hence define the aggregation metric by summing over all tags the average of Tag-Shift absolute values across all  $m \rightarrow n$  transitions. This metric is formulated as follows:

$$\text{Overall-Shift} = \frac{1}{|C|} \sum_{m \rightarrow n \in C} \sum_{t \in T} |\text{Tag-Shift}(t, m \rightarrow n)| \quad (3)$$

where  $C$  indicates the set of all possible bias label transitions. For both gender and country attributes,  $|C|$  is equal to 6. The results of the Overall-Shift metrics for all backbone LLMs, architectures, and debiasing settings are reported in Table 6. For a better comparison, the table also reports the average task performance, taken from Table 5. These results similarly indicate that despite the existence of biases in all models, the degrees of biases significantly decrease after applying CDA using the provided data of HUMSETBIAS. This bias mitigation does not cause any drop in the models’ performance while observing slight improvements in some cases. Among various models, the HumBERT LLM together with our architecture demonstrates effective results in terms of both classification performance (the highest classification results on F1-measure) and bias measurement (lowest gender bias). These observations indicate the benefits of using our proposed model with CDA, as an effective and bias-aware classification solution for the humanitarian sector.

## 6 Conclusion

During humanitarian crises, processing and classification of the data greatly assist analysts and decision-makers in gaining a comprehensive understanding of the situation in line with humanitarian imperatives and the Leave No One Behind (LNOB) principle. In this work, we propose a set of resources and methods for building a practice-oriented, effective, and bias-aware solution for entry classification in the humanitarian domain. We first introduce a novel parameter-sharing model architecture, designed according to the structure of the humanitarian analysis framework, followed by creating and releasing HumBERT– a novel LLM adapted to the topics and nuances of this domain. Our experiments on the HUMSET dataset in both zero-shot and full-training settings, show the overall performance improvements of the models when using HumBERT in comparison with generic LLMs with comparable size, and also the overall benefits of our proposed architecture in comparison with baseline fine-tuning, particularly on the tasks with many target classes. In addition, we introduce the HUMSETBIAS dataset, consisting of data points with specific gender and country attributes. Utilizing the counterfactual variations of the data points in HUMSETBIAS, we observe the significant effects of these attributes on the predictions of the models, which we successfully mitigate utilizing a targeted data augmentation method.

## Ethical Statement

This work aims to promote more equitable and inclusive humanitarian response practices, by providing a sector-specific model and a comprehensive dataset for analyzing and mitigating biases. However, we emphasize that the proposed data, model, and classification architecture have inherent limitations explained below, which should be considered thoroughly, particularly when they are utilized in practice.

HUMSETBIAS is created using rule-based mappings. Hence, we further review the entries to exclude the ones with negative keywords for both gender and country attributes and conduct several sanity checks iterations. However, the dataset should not be assumed as completely error-free, as there might be some minor shifts in the results we obtained. We indeed aim to continue improving the dataset, particularly through the common effort of the humanitarian community.

Next, while the CDA method significantly reduces bias, the models still contain bias. We should consider that such biases may still lead to some issues depending on the usage scenario.

Finally, the HUMSETBIAS dataset does not cover all possible societal biases in the humanitarian domain. We hence encourage more effort to enlarge this dataset to other critical attributes such as specific needs groups, affected groups, and age categories. We believe that these features are central to any humanitarian technological solution in support of the analysis, particularly given the humanitarian imperatives of neutrality and impartiality and the LNOB principle part of the 2030 agenda.

## Acknowledgments

We are immensely thankful to the DEEP users, taggers, and project owners for their dedicated work in making this dataset possible. Furthermore, we acknowledge the DEEP board members, which guidance has been crucial in shaping the DEEP. We express our sincere appreciation to the entire DFS team, the ToggleCorp team, ISI Foundation, and the Johannes Kepler University Linz, for their continuous support in enabling the use of NLP in the humanitarian community. Their collective contributions have been instrumental in the success of this research. Furthermore, we also express our sincere appreciation to the USAID’s Bureau for Humanitarian Assistance (BHA) for their trust in DFS and their generous grant support that allows us to work on these topics.

NT gratefully acknowledges the support from the Lagrange Project of the ISI Foundation funded by CRT Foundation. NT thanks Ciro Cattuto, Kyriaki Kalimeri, Yelena Mejova, and Rossano Schifanella for their supervision.

## References

[Adel and Wang, 2020] Ghadah Adel and Yuping Wang. Detecting and classifying humanitarian crisis in arabic tweets. In *2020 3rd International Conference on Artificial Intelligence and Big Data (ICAIBD)*, 2020.

[Alam et al., 2021] Firoj Alam, Hassan Sajjad, Muhammad Imran, and Ferda Ofli. Crisisbench: Benchmarking crisis-related social media datasets for humanitarian information processing. In *Proceedings of the International AAAI Conference on Web and Social Media*, 2021.

[Alharbi and Lee, 2021] Alaa Alharbi and Mark Lee. Kawarith: an Arabic Twitter corpus for crisis events. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 42–52, Kyiv, Ukraine (Virtual), April 2021. Association for Computational Linguistics.

[Araci, 2019] Dogu Araci. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*, 2019.

[Beltagy et al., 2019] Iz Beltagy, Kyle Lo, and Arman Cohan. SciBERT: A pretrained language model for scientific text. 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)*, 2019.

[Blodgett et al., 2020] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. Language (technology) is power: A critical survey of “bias” in nlp. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, 2020.

[Chen et al., 2020] Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. Recall and learn: Fine-tuning deep pretrained language models with less forgetting. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.

[Conneau et al., 2020] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.

[Devlin et al., 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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*, 2019.

[Fekih et al., 2022] Selim Fekih, Nicolo’ Tamagnone, Benjamin Minixhofer, Ranjan Shrestha, Ximena Contla, Ewan Oglethorpe, and Navid Rekabsaz. HumSet: Dataset of multilingual information extraction and classification for humanitarian crises response. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4379–4389, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics.

[Gao et al., 2021] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. In *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)*, Online, 2021.

[Hauzenberger et al., 2023] Lukas Hauzenberger, Shahed Masoudian, Deepak Kumar, Markus Schedl, and Navid Rekabsaz. Modular and on-demand bias mitigation with



- attribute-removal subnetworks. In *Findings of the Association for Computational Linguistics: ACL (Findings of ACL)*, 2023.
- [Hu *et al.*, 2022] Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022.
- [Imran *et al.*, 2016] Muhammad Imran, Prasenjit Mitra, and Carlos Castillo. Twitter as a lifeline: Human-annotated twitter corpora for nlp of crisis-related messages. *arXiv preprint arXiv:1605.05894*, 2016.
- [Kumar *et al.*, 2023] Deepak Kumar, Oleg Lesota, George Zerveas, Daniel Cohen, Carsten Eickhoff, Markus Schedl, and Navid Rekasaz. Parameter-efficient modularised bias mitigation via AdapterFusion. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2738–2751, 2023.
- [Lai *et al.*, 2022] Kelvin Lai, Jeremy R Porter, Mike Amodeo, David Miller, Michael Marston, and Saman Armal. A natural language processing approach to understanding context in the extraction and geocoding of historical floods, storms, and adaptation measures. *Information Processing & Management*, 59(1):102735, 2022.
- [Lee *et al.*, 2020] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.
- [Liu *et al.*, 2021] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. GPT understands, too. *CoRR*, abs/2103.10385, 2021.
- [Lu *et al.*, 2018] Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. Gender bias in neural natural language processing. *arXiv preprint arXiv:1807.11714*, 2018.
- [Malte and Ratadiya, 2019] Aditya Malte and Pratik Ratadiya. Evolution of transfer learning in natural language processing. *arXiv preprint arXiv:1910.07370*, 2019.
- [Parisi *et al.*, 2019] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- [Rekasaz and Schedl, 2020] Navid Rekasaz and Markus Schedl. Do neural ranking models intensify gender bias? In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2065–2068, 2020.
- [Rekasaz *et al.*, 2021a] Navid Rekasaz, Simone Kopeinik, and Markus Schedl. Societal biases in retrieved contents: Measurement framework and adversarial mitigation of BERT rankers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 306–316, 2021.
- [Rekasaz *et al.*, 2021b] Navid Rekasaz, Robert West, James Henderson, and Allan Hanbury. Measuring societal biases in text corpora via first-order co-occurrence. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*, 2021.
- [Roese and Olson, 1995] Neal J Roese and James M Olson. Counterfactual thinking: A critical overview. *What might have been: The social psychology of counterfactual thinking*, pages 1–55, 1995.
- [Schick and Schütze, 2021] Timo Schick and Hinrich Schütze. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online, April 2021. Association for Computational Linguistics.
- [Sheng *et al.*, 2019] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. The woman worked as a babysitter: On biases in language generation. 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 3398–3403, 2019.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [Wolf *et al.*, 2020] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October 2020. Association for Computational Linguistics.
- [Yela-Bello *et al.*, 2021] Jenny Paola Yela-Bello, Ewan Oglesby, and Navid Rekasaz. MultiHumES: Multilingual humanitarian dataset for extractive summarization. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1713–1717, Online, April 2021. Association for Computational Linguistics.
- [Zhao *et al.*, 2018] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of North American Chapter of the Association for Computational Linguistics*, 2018.