

OATH-Frames: Characterizing Online Attitudes Towards Homelessness with LLM Assistants

Jaspreet Ranjit[♥] Brihi Joshi[♥] Rebecca Dorn[♥] Laura Petry[♦]
Olga Koumoundouros[♦] Jayne Bottarini[♥] Peichen Liu[♦]
Eric Rice[♦] Swabha Swayamdipta[♥]

[♥]Thomas Lord Dept. of Computer Science, University of Southern California

[♦]Suzanne-Dwork School of Social Work, University of Southern California

jranjit@usc.edu

Abstract

Warning: Contents of this paper may be up-setting. Public attitudes towards key societal issues, expressed on online media, are of immense value in policy and reform efforts, yet challenging to understand *at scale*. We study one such social issue: homelessness in the U.S., by leveraging the remarkable capabilities of large language models to assist social work experts in analyzing millions of posts from Twitter. We introduce a framing typology: Online Attitudes Towards Homelessness (OATH) Frames: nine hierarchical frames capturing critiques, responses and perceptions. We release annotations with varying degrees of assistance from language models, with immense benefits in scaling: $6.5\times$ speedup in annotation time while only incurring a 3 point F_1 reduction in performance with respect to the domain experts. Our experiments demonstrate the value of modeling OATH-Frames over existing sentiment and toxicity classifiers. Our large-scale analysis with predicted OATH-Frames on 2.4M posts on homelessness reveal key trends in attitudes across states, time periods and vulnerable populations, enabling new insights on the issue. Our work provides a general framework to understand nuanced public attitudes at scale, on issues beyond homelessness.

1 Introduction

Public opinion towards social issues is used as a signal by advocacy organizations to garner support for their relief and reform efforts (Burstein, 2020). Online social media platforms, such as Twitter¹, serve as a great repository of public opinion from a diverse (e.g., socio-economic and political) user base (Hu et al., 2021). However, understanding broader public opinion necessitates processing online discourse *at a large scale*, a daunting challenge in social science research. Large language models (LLMs) have recently been used in social science

(Ziems et al., 2023; Cho et al., 2023; Li et al., 2023; Bohra et al., 2023), and may offer a promising solution to the scale challenge. However, can LLMs be reliably used to process discourse with nuanced latent attitudes towards complex societal issues?

One such issue is homelessness in the U.S.², which inspires complex and nuanced public attitudes and beliefs (Robertson and Greenblatt, 1992). For instance, it is common for people to simultaneously feel both sympathy towards the plight of people experiencing homelessness (PEH), as well as anger and resentment towards them (Cronley, 2010; Moore et al., 1988; Narayan and Petesch, 2002). A majority of work from social science has studied discourse about homelessness via ethnographic studies, qualitative surveys, and smaller scale content analyses (Kim et al., 2023; Tsai et al., 2019; Kingree and Daves, 1997). However, understanding the societal attitudes about homelessness *at a large scale* via online discourse has rarely been attempted in prior efforts. Our work investigates the role of LLMs in reliably assisting social work experts in characterizing such large-scale online discourse on homelessness.

In this work, we introduce a novel framework to understand, synthesize and characterize large-scale public attitudes towards societal issues through a collaboration between social work experts and LLMs, illustrated in Figure 1. Our framework consists of three components designed to study written discourse on social issues: (i) the discovery of a framing typology (§2), (ii) annotation of data at scale with this typology (§3), and (iii) the analysis of these collected annotations (§4). While we focus our study on homelessness, our framework is generalizable to other domains of social discourse.

At the heart of our framework is our novel framing typology to characterize Online Attitudes To-

²Over 653,104 U.S. individuals experienced homelessness in 2023 (Sousa et al., 2023).

¹Now, X: <https://x.com>.

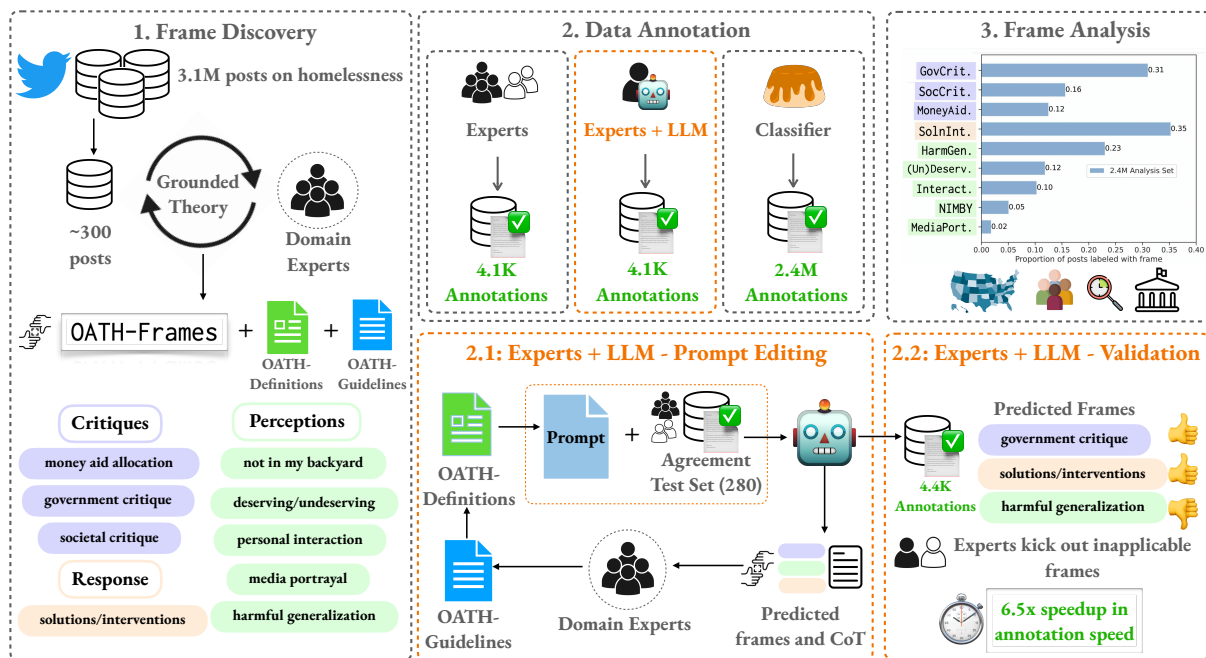


Figure 1: **Overview of OATH-Frames Collection.** (1.) Domain experts applied grounded theory to surface nine Issue-specific frames, their corresponding definitions (OATH-Definitions), and annotation guidelines (OATH-Guidelines). (2.) We annotated posts with OATH-Frames via Experts, via Experts + LLM (GPT-4), and via Multilabel Classifier (Flan-T5-Large). Our Experts + LLM annotations pipeline consists of (2.1) prompt editing based on insights from domain experts and chains-of-thought, and (2.2) validation of predicted frames. (3.) We analyzed annotated posts for attitudes across social and political dimensions.

wards Homelessness (OATH-Frames) in the U.S., as expressed in social media posts from Twitter (§2). Our frames are developed using framing theory (Entman, 1993) and grounded theory (Glaser and Strauss, 1967) from social science, and in collaboration with social work researchers studying homelessness. OATH-Frames can be broadly categorized into three themes: Critiques, Perceptions and Responses to PEH, and further into nine finer-grained Issue-specific frames; Table 1 illustrates our typology with examples. We provide three kinds of OATH-Frames annotations for posts from Twitter: 4.1K expert-only (§3.1), 4.1K LLM-assisted expert (§3.2) and 2.4M predicted annotations from a multi-label classification model (§3.3). Our LLM-assisted expert annotation for socially sensitive tasks follows a new scalable framework which incorporates expert insights (OATH-Guidelines) and chain-of-thought (CoT) explanations for designing better LLM prompts. This results in a $6.5\times$ annotation speedup, while incurring only a 3 point F_1 reduction in accuracy compared to experts—an attractive trade-off for our multi-label annotation task.

We observe that the most salient theme in expert-

only annotations is Perceptions of PEH while the rest of our annotations highlight Critiques of the government, normative structures and (lack of) aid distribution (Figure 2).

Our 2.4M OATH-Frames-annotated posts enable a large-scale analysis across states and time periods, revealing changing trends in attitudes with key sociopolitical events (§4). OATH-Frames surface harmful language towards PEH, which we show are often mislabeled by popular sentiment and toxicity classifiers, highlighting the value of our typology. We empirically establish trends where differences in attitudes are observed, when vulnerable populations are compared with PEH; for instance, comparisons with immigrants elicit harmful rhetoric whereas those with Ukrainians elicit critiques towards the government for funding allocation. OATH-Frames can assist in understanding nuanced issues concerning marginalized PEH populations, often ignored in language technologies (Kiritchenko et al., 2023). Overall, our framework provides a blueprint for fast and scalable analysis of public opinion on key social issues of our time; we believe this will aid advocacy groups and political constituents in work towards meaningful social

Definition	Example and Labels from Different Annotators
GovCrit. Government Critique: Criticism about government body, policies and laws including discussion of homelessness through the lens of political parties and values.	these people are already housed possibly over crowded but at least housed .there are up to 130000 homeless living rough on our streets every night women and children in the back of vans and cars and the gov gives billions away to help third world countries.? time he looked here All: GovCrit. MoneyAid. (Un)Deserv.
MoneyAid. Money Aid Resource Allocation: Discussion of money, aid or resource disbursement on addressing the homelessness issue. Examples include the allocation of necessities such as essential items for emergency relief supplies, or government budgeting with respect to competing priorities.	you propose atlanta/fulton county fund a homeless city. i'm asking what return are we getting for our tax dollars? police officers and firefighters provide a service which can justify us building a training for them. what justification do you have for your city for the homeless? All: MoneyAid. GovCrit. / Expert, GPT-4, Expert+GPT-4: SolnInt. / Expert: (Un)Deserv.
SocCrit. Societal Critique: Criticism of social norms, systems and society at large in how homelessness is being addressed and perceived. Also includes pointing out hypocrisy and performative activism.	i really hate people who have mango trees but don't eat mangos so they just let them go to waste. i be walking by peoples yards and see hella mangos just laying there for days. give them to the homeless, don't just let the mangos sit there and rot. All: SocCrit., SolnInt.
SolnInt. Solutions and Interventions: Discussion of solutions, interventions, charitable acts and remedies to address the homelessness crisis.	Ever think that instead of jail, A MENTALLY ILL HOMELESS PERSON needs actual help? Yes, I have compassion for his victims. But I also have compassion for this guy. He needs help. He doesn't need to be thrown in jail where he'll rot and probably be left to die. You pointed out All: SolnInt. / Expert: HarmGen. / GPT-4: GovCrit., Interact.
NIMBY Not in my Backyard: Opposition by residents to proposed developments in their local area, as well as support for strict land use regulations against wanting to see homelessness in their local area and neighborhood.	just found out the head of the opposition to a local affordable housing for the homeless project is a senior planner in the neighboring county Expert, GPT-4, Expert+GPT-4: NIMBY, GovCrit. / Expert, Flan-T5-Large: SolnInt.
Interact. Personal Interaction: Anecdote describing a direct personal exchange with PEH.	i'm legit sad af. i saw a homeless woman and her daughter. i really wanted to help but i had no cash. i hope they are still there when i circle back around. All: Interact. / Expert, GPT-4, Expert+GPT-4: SolnInt.
MediaPort. Media Portrayal: Portrayal of (fictional or real) PEH as described in the media (e.g. discussing PEH in a TV show or in the news).	is the national news reporting on the homeless on the streets and encampments as well as random crime exaggerated? beautiful city and location. All: MediaPort. / Expert: NIMBY
(Un)Deserv. Deserving vs. Undeserving of Resources: Perpetuating a hierarchy of PEH with other marginalized communities or the use of harmful generalizations such as substance use and mental illness to justify that PEH that are more or less deserving of aid. Includes nationalistic rhetoric.	Look at the illegals 500\$ a night hotel rooms , destroying them, wasting food instead of giving to the homeless All: (Un)Deserv. MoneyAid. / Expert: SolnInt. / Flan-T5-Large: GovCrit.
HarmGen. Harmful Generalization: Blanket statements that ascribe an undesirable characteristic to PEH that include but are not limited to generalizing all PEH as having an unkempt appearance, or being violent, racist, thieves, or sexual predators.	@mention you're more confused than a homeless man in a house arrest. you can't even put your words together. please get out. All: HarmGen., SocCrit.

Table 1: **OATH-Frames:** Definitions and examples for our 9 Issue-specific OATH-Frames, accompanied by our Expert labels and GPT-4, Expert+GPT-4, and Flan-T5-Large predictions. Each post is associated with multiple frames, at times even across our three overarching themes: (i) Critiques, (ii) Perceptions, and (iii) Responses.

reforms (Burstein, 2020; Applebaum, 2001). Our data and code will be publicly available³.

2 OATH-Frames Discovery

We discovered a novel typology of pragmatic frames for capturing and characterizing societal attitudes towards homelessness and people experiencing homelessness (PEH). To this end, we applied concepts from framing theory (Entman, 1993), which has been used to identify how certain aspects in an issue's portrayal are emphasized to influence broader public opinion for topics like immigration (Mendelsohn et al., 2021), politics in media (Boydston and Gross, 2013), eating disorders (Arseniev-Koehler et al., 2016), and so on. In social media discourse, prior work has focused on the framing of stigma towards PEH (Kim et al., 2023) and surveys that study attitudes and perceptions towards PEH (Toro and McDonnell, 1992; Agans and Liu, 2015). In our work, we drew in-

spiration from Kim et al. (2023) and their application of grounded theory (Glaser and Strauss, 1967) to characterize societal attitudes on homelessness at scale, going beyond stigma. Following Boydston and Gross (2013), our *Frame Discovery* process uncovered salient themes (frame types) in a text corpus on a topic; in the context of our study, this resulted in OATH-Frames (Frames for Online Attitudes Towards Homelessness).

Using a keyword search for “homeless” on the Twitter API,⁴ we sample 5M posts. We preprocessed this sample (i.e. deduplicated, replaced emojis with unicode, anonymized user tags, and removed non-English posts) to retain 3.1M posts from 2.1M unique users spanning from 2021 to 2023.

As a first step in *Frame Discovery*, we only considered posts discussing public attitudes and opinions on homelessness, following Boydston and

⁴<https://developer.X.com/en/docs/X-api> accessed in February 2023.

³<https://dill-lab.github.io/oath-frames/>

Gross’s (2013) Public Opinion frame. We discarded posts which we deemed ambiguous, under-specified or unrelated (e.g. political homelessness). We also discarded posts containing first person lived experiences of PEH, since our focus was on studying public attitudes on homelessness.

Next, we followed an extensive iterative ‘coding’ process based on grounded theory from social science (Glaser and Strauss, 1967) to surface our OATH-Frames where ‘codes’ are 1-2 word descriptors of the data. The coding process involved our four domain experts independently analyzing a sample of 100 posts to conceptualize a unique categorization of attitudes. This was followed by a discussion where all posts were reviewed and similar frames were grouped together till we reached consensus on definitions, edge cases and representative examples. We repeated this process thrice on different sets of posts with refined iterations of our frames until we reached thematic saturation which surfaced nine fine-grained Issue-specific frames under homelessness; see our OATH-Definitions in Table 1. Appendix A.1 provides our experts’ OATH-Guidelines developed during *Frame Discovery* that contain detailed definitions and steps on how to address edge cases.

Our Issue-specific frames can be categorized into three overarching themes, as illustrated in Figure 1 (*Frame Discovery*). First, we find Critiques of systems and institutions enabling homelessness and also the use of homelessness as a vehicle to critique others (e.g. actors, programs or spending). Second, we find posts reflecting different Perceptions of PEH which attribute undesirable characteristics to PEH to make harmful generalizations. Third, we find many Responses to the issue, posed as solutions, interventions, calls to action and charitable acts. Our three themes are consistent with those in prior work on homelessness (Baiocchi and Arguello, 2019; Gowan, 2010), but our work is the first to surface nine Issue-specific (homelessness) frames providing more fine-grained insights.

3 Annotation of OATH-Frames

3.1 Expert Annotation

Our four domain experts leading the frame discovery effort (§2) trained two other annotators (trained experts). The trainees first familiarized themselves with OATH-Definitions and OATH-Guidelines as detailed in Appendix A.1. Next, the trainees

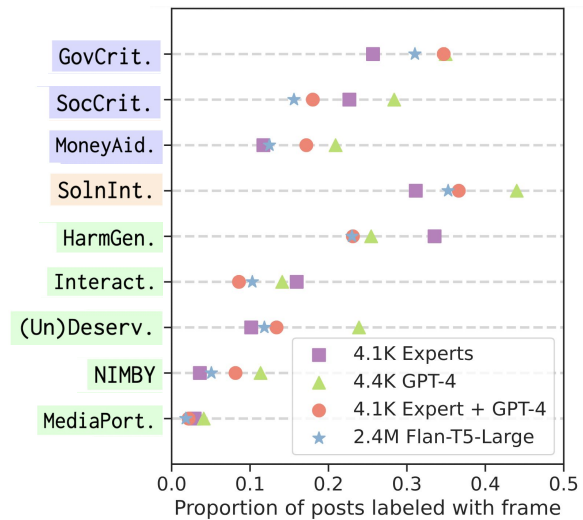


Figure 2: **Proportion of Issue-specific frames** across different annotation methods. Distribution of OATH-Frames across annotation strategies is similar. See more details in Table 4 in Appendix C.

labeled 20 posts with OATH-Frames where the experts clarified questions about OATH-Definitions. Finally, the trainees performed an independent annotation exercise on a set of 280 posts annotated by the domain experts. The training concluded once the trainee achieved moderate agreement (Fleiss’ $\kappa > 0.4$; Fleiss, 1971) with the experts.

For all relevant posts, each expert (domain and trained) annotator performs a 9-way-multilabel classification of Issue-specific frames, where 54% of the posts were labeled with more than one frame. Co-occurrence of frames represented multiply expressed attitudes highlighting the complexity of our task. All annotators achieved an average inter-annotator agreement of 0.46 (Fleiss’ κ) and a macro F_1 of 0.55 (see Table 2) across nine Issue-specific frames.⁵

Overall, we considered 5.3K posts, of which we filtered out 1.2K posts, leaving 4.1K expert-annotated posts. See Appendix B.1 for details on our inter-annotator agreement, co-occurrence of frames and annotation interface.

3.2 LLM-Assisted Expert Annotation

Given the nature of the rhetoric on homelessness, annotating 5.3K posts was laborious, and emotion-

⁵Our agreement across 9 frames is ‘moderate’; prior work has reported similar agreement of 0.3-0.4 across 8 frames and 4 annotators in a multi-label classification task (Stede et al., 2023). We calculate the agreement across six annotators, a much higher number than prior work that considers only two annotators (Card et al., 2015).

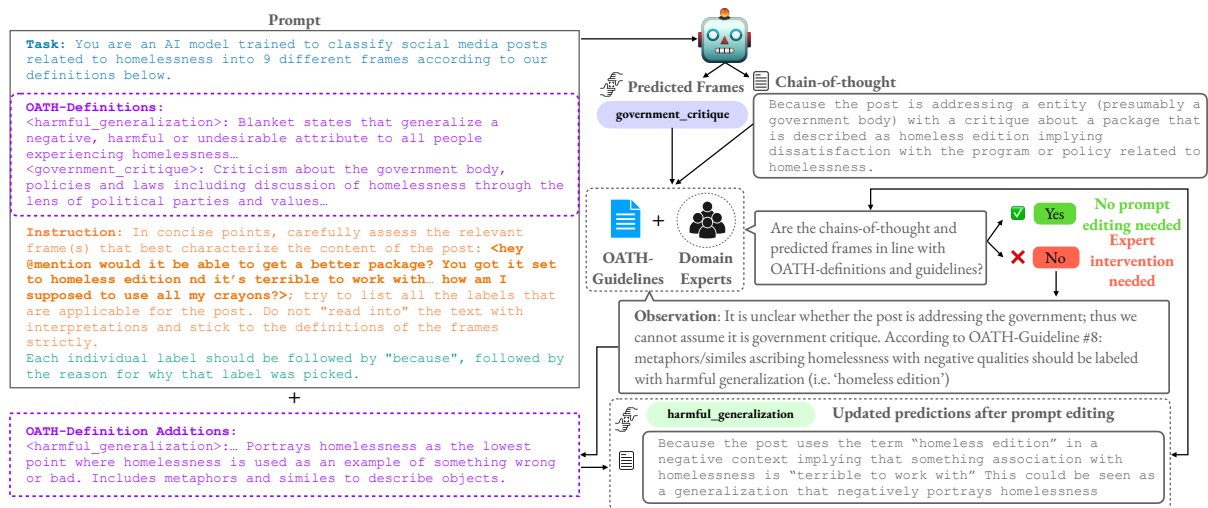


Figure 3: **Experts+LLM-Prompt Editing:** We prompt GPT-4 with our task, OATH-Definitions, and our instruction as well as our own OATH-Guidelines and GPT-4’s Chain-of-Thought reasoning (CoT; Wei et al., 2022), in an iterative setup. We observe that CoT serves as a signal to understanding misinterpretations of OATH-Definitions and benefits from addition of expert annotator guidelines (OATH-Guidelines).

ally challenging for the human experts. LLMs are now being deployed as annotation assistants in social science (Li et al., 2023; Bohra et al., 2023), but questions on how best to incorporate insights from social science into LLMs still remain (Nelson et al., 2024; Gondimalla et al., 2024); also see Appendix E.3. We explore a collaborative setting for annotating OATH-Frames with LLM assistants to our expert annotators to ameliorate their manual and emotional labor. We first investigate the ability of GPT-4 (Achiam et al., 2023) to do our annotation task *stand-alone* (§3.2.1), before describing our collaborative framework (Expert+GPT-4; §3.2.2).

3.2.1 Predicting OATH-Frames with only GPT-4

We design our prompts for GPT-4 by incorporating instructions and examples as well as our own OATH-Guidelines and GPT-4’s own Chain-of-Thought reasoning (CoT; Wei et al., 2022), in an iterative setup as shown in Figure 3. Building on prior work (Pangakis et al., 2023; Rytting et al., 2023), we observe that CoT serves as a signal to understanding misinterpretations of OATH-Definitions and benefits from addition of expert annotator guidelines (OATH-Guidelines). Our prompt refinement pipeline could be generalized to other social science tasks, where practitioners could use LLM CoT and their expert annotator guidelines to leverage LLMs as annotation assistants. Appendix B.2 details the prompts we used.

We report GPT-4’s performance on a subset of

280 multiply-annotated posts⁶ ($|D_{\text{agreement}}| = 280$) in Table 2. GPT-4 achieves a precision of 0.39 ± 0.05 and recall of 0.64 ± 0.05 across all nine Issue-specific frames showing that it may have reasonable coverage of the predicted labels but also produces false positives due to the inherent social nuances in the posts that may be easier for a expert to resolve. For example, GPT-4 often confuses the SocCrit. and GovCrit. frames, seen in the following post: ‘@mention it appears the problem in el paso with illegal aliens were solved overnight, so that you could have a photo op without anyone seeing hundreds of homeless illegals on the streets’. Here, although the government is not explicitly mentioned, politicians often enact displacement sweeps and photo-ops to disguise the severity of the homelessness crisis. We observe that GPT-4 is not able to resolve domain-specific references (i.e. displacement sweeps) and thus predicts both SocCrit. and GovCrit. (Table 7). In these cases, it’s important to have experts validate the predictions.

3.2.2 Predicting OATH-Frames with Expert+GPT-4

Although LLMs serve as a useful tool to efficiently generate annotations, the inherent subjectivity and sensitivity of our task poses a challenge to using GPT-4 as an annotator on its own (Rytting et al., 2023; Heseltine and Clemm von Hohenberg, 2024).

⁶This allows us to compute performance within a margin of error across all 6 (expert) annotators.

Annotators →	Experts*	GPT-4	Expert+GPT-4	Flan-T5-Large	Flan-T5-Large	Llama-2-7b
Test Set →	$ D_{\text{agreement}} = 280$			$ D_{\text{test}} = 1280$		
GovCrit.	0.65 ± 0.06	0.70 ± 0.08	0.72 ± 0.07	0.65 ± 0.07	0.63	0.35
MoneyAid.	0.60 ± 0.16	0.58 ± 0.14	0.61 ± 0.15	0.59 ± 0.18	0.61	0.22
SocCrit.	0.43 ± 0.09	0.46 ± 0.06	0.50 ± 0.05	0.42 ± 0.09	0.49	0.31
SolnInt.	0.69 ± 0.06	0.62 ± 0.06	0.65 ± 0.06	0.68 ± 0.03	0.57	0.48
Interact.	0.45 ± 0.13	0.45 ± 0.07	0.52 ± 0.07	0.54 ± 0.11	0.70	0.18
MediaPort.	0.70 ± 0.10	0.50 ± 0.08	0.63 ± 0.11	0.69 ± 0.10	0.54	0.13
NIMBY	0.17 ± 0.21	0.30 ± 0.16	0.39 ± 0.21	0.34 ± 0.26	0.26	0.10
(Un)Deserv.	0.58 ± 0.19	0.46 ± 0.06	0.63 ± 0.08	0.67 ± 0.01	0.61	0.22
HarmGen.	0.53 ± 0.07	0.49 ± 0.06	0.51 ± 0.06	0.55 ± 0.05	0.58	0.41
Public Opinion	0.75 ± 0.05	0.04 ± 0.05	0.04 ± 0.05	0.0 ± 0.0	0.0	0.02
Aggregated _{Micro}	0.62 ± 0.01	0.54 ± 0.05	0.59 ± 0.05	0.57 ± 0.04	0.55	0.30
Aggregated _{Macro}	0.55 ± 0.04	0.46 ± 0.04	0.52 ± 0.04	0.51 ± 0.03	0.50	0.24

Table 2: Mean (\pm s.d. across 6 expert annotator labels) F_1 for Issue-specific frames and the Public Opinion filter, and overall macro and micro F_1 (Aggregated). Expert * reports the average F_1 computed using each annotator as a reference. Expert+GPT-4 and Flan-T5-Large consistently outperform GPT-4, while lagging behind Expert *. Performance of Flan-T5-Large and Llama-2-7b-chat-hf (zero-shot) are reported on 1280 singly annotated test instances. See detailed precision and recall for each row in Appendix Table 12.

We propose an Expert+GPT-4 approach that leverages both GPT-4’s annotation efficiency and human expertise to verify and correct GPT-4’s generations. We generate GPT-4 predictions on a set of 4.4K posts. Through expert validation, we filter out 300 posts, resulting in 4.1K Expert+GPT-4 validated posts. Our expert validation involves discarding false positive labels by GPT-4; this results in a 12 point boost in precision and a slight drop in recall compared to GPT-4 generations. We note that by themselves, experts take 187.49 (\pm 30.71) seconds, whereas with Expert+GPT-4, they take 28.80 (\pm 8.55) seconds, which is a 6.5 \times boost in annotation speed. Overall, Expert+GPT-4 results in a 6 point increase in macro F_1 over GPT-4; see Table 2.

3.3 Scaling OATH-Frames Annotations

While Expert+GPT-4 reduces manual labor, annotations by Expert and Expert+GPT-4 are relatively cost-intensive at a larger scale. To this end, we scale OATH-Frames annotations using a model trained with Expert and Expert+GPT-4 annotations. We obtain 10.4k training examples: 5.3K Expert and 4.1K Expert+GPT-4 annotations and 1K filtered posts. We split these into a train-val-test split containing 8217, 913 and 1280 posts respectively and train a Flan-T5-Large (Chung et al., 2022) model for 10-way multilabel classification (nine OATH-Frames and filter).

Results in Table 2 show that Flan-T5-Large performs better than GPT-4 with a Macro F_1 of 0.51

\pm 0.05, and is on par with the Expert+GPT-4 setting on the Agreement Test Set ($|D_{\text{agreement}}| = 280$). We also find that Flan-T5-Large obtains a Macro F_1 score of 0.50 on a larger test set of 1280 posts ($|D_{\text{test}}| = 1280$). Among four fine-tuned size-variants of Flan-T5, and zero-shot Llama-2-7b (Touvron et al., 2023), Flan-T5-Large performs the best; detailed results and data settings are provided in Appendix C.

Using our trained Flan-T5-Large model, we generate OATH-Frames predictions on a set of 2.4M unlabeled posts. Figure 2 illustrates that the distribution of frames under different annotation strategies is similar. Next, we use the 2.4M predicted annotations in a large-scale socio-political analysis of public attitudes.

4 Analyzing Posts with OATH-Frames

Synthesizing public opinion using online discourse is useful for advocacy groups and policymakers to inform homelessness services and equitable housing policies to local communities (Burstein, 2020). We demonstrate the utility of OATH-Frames to such end users by analyzing attitudes with respect to state mentions, marginalized populations, and time period on the set of 2.4M posts. We also highlight the utility of our OATH-Frames over sentiment and toxicity labels.

4.1 OATH-Frames Versus Toxicity and Sentiment Prediction

PEH experience harms online in unique ways through subtle stereotyping (i.e. HarmGen. : ‘you look homeless’; Kiritchenko et al., 2023; Hutchinson et al., 2020). We investigate how existing toxicity and sentiment classifiers capture harmful rhetoric against PEH as represented by the HarmGen. frame. We plot the distribution of

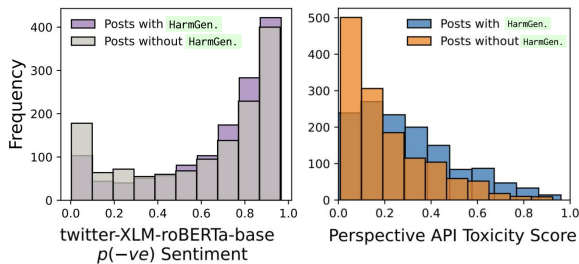


Figure 4: Distributions of sentiment and toxicity scores for posts labeled with and without HarmGen. in our subset of 4.1K Expert-annotated posts.

the Perspective API⁷ toxicity score and sentiment scores (Rosenthal et al., 2017) in Figure 4 for the subset of Expert-annotated posts labeled with HarmGen. . We find that a majority of HarmGen. posts are labeled with a toxicity of less than 0.5, indicating a high false negative rate. For example, the following post contains harmful stereotypes against PEH but was tagged with a toxicity score of less than 0.5: ‘@mention some people wear clothes at home i wear for an outing and here i dress up like a homeless person inside my home whos a drug addict and has been sleeping since forever’. Furthermore, the negative sentiment distributions for posts labeled with and without HarmGen. are significantly similar (*independent samples t-test*: $p < .001$) indicating that sentiment classifiers are not able to represent HarmGen. against PEH. This highlights that OATH-Frames provides a framework characterizing how PEH may experience harms and subtle stereotyping, going beyond sentiment and toxicity. Our framework can serve to improve existing classifiers in detecting such harms, especially for underrepresented communities (Dorn et al., 2023).

4.1.1 Case Study: Solutions and Interventions in Posts Mentioning California

Given the high rates of PEH in California (CA accounted for 49% of U.S. unsheltered PEH in

⁷<https://perspectiveapi.com>

2023), we investigate the proposed SolnInt. for posts that mention CA. We measure the statistical prevalence of a term in CA posts in two subsets: posts labeled with SolnInt. and posts labeled with a positive sentiment probability ($p(+ve) > 0.5$) (Rosenthal et al., 2017) in our set of 2.4M posts.

(+) SolnInt.	Score	$p(+ve)$	Score
rent	3.72	great	4.68
voucher	2.59	love	4.22
abandon	2.48	new	3.26
bass	4.68	thank	3.16
healthcare	2.03	help	2.32

Table 3: Log-odds ratios for terms that are statistically more likely to appear in posts about CA labeled with (+) SolnInt. and posts about CA with positive sentiment ($p(+ve) > 0.5$). Scores outside the range $[1.96, -1.96]$ are considered statistically significant. Full set of results are included in Table 13.

We calculate Monroe et al.’s (2017) weighted log-odds ratio with a Dirichlet informed prior for unigrams/bigrams. By incorporating a prior estimate over the expected distribution of the vocabulary, this approach accounts for words that are very commonly or rarely prevalent in the two subsets of posts. Table 3 shows 5 significant terms in each subset and their z-score.

We find that in posts labeled with (+) SolnInt. , OATH-Frames captures more comprehensive insights specific to homelessness in CA. Whereas posts labeled with a positive sentiment ($p(+ve)$) contain more generic insights (i.e. ‘great, love, help’). For example, ‘abandon, voucher’, and ‘bass’ refer to LA Mayor Karen Bass’s homelessness initiative: ‘Inside Safe’ that converted abandoned hotels to temporary shelters for PEH (Urevich, 2023) whereas ‘rent’ and ‘healthcare’ are considered as drivers of homelessness in CA and potential entry points for interventions. Although sentiment classifiers provide coarse-grained insights of conflating factors in relation to homelessness (e.g. substance use, crime), we find that OATH-Frames provide more targeted insights specific to CA.

4.2 Attitudes across State Mentions

Discourse on homelessness is closely tied to a state’s social and economic climate and the prevalence of PEH in each region (Tsai et al., 2019). We investigate the saliency of our Issue-specific

frames in posts that mention states in the U.S.⁸ Given less than 1% of tweets are geotagged (Ajao et al., 2015), we use NER⁹ to segment our data by posts mentioning different U.S. states. We find that HarmGen. is the most prominent Issue-specific frame in posts about California, Washington and Oregon and (Un)Deserv. and MoneyAid. are the most prominent in posts about New York (see Figure 12, Table 10 in Appendix D). To better understand the drivers of attitudes in a state, we also estimated a linear regression for two state level factors (i.e. cost of living index, percent of unsheltered PEH) in relation to the proportion of posts labeled with frames¹⁰ as shown in Figure 5. We find that a high cost of living index¹¹ correlates to a larger proportion of GovCrit. (R^2 of 0.61) and fewer SolnInt. whereas a larger unsheltered PEH population correlates to more HarmGen. (R^2 of 0.45) and less (Un)Deserv. We hypothesize that greater visibility of PEH in the west coast, which is due to a larger unsheltered population (68% in California compared to 4.9% in New York), results in HarmGen. and in contrast, posts about New York are more likely to contain (Un)Deserv. attitudes due to an influx of asylum seekers and immigrants in 2023 (Sousa et al., 2023). In addition, higher costs of living elicit more GovCrit. for economic reform and reducing rent prices.

4.3 Differences in Attitudes Across Vulnerable Populations and PEH

Attitudes of (Un)Deserv. (Solas, 2018) reflect comparisons between vulnerable populations and PEH as being more or less deserving of aid (Kiritchenko et al., 2023; Hoffmann, 2019) putting them at risk for being further ostracized from garnering public support for relief and reform measures (Lalor et al., 2022). We investigate the finer-grained differences in attitudes as characterized by

⁸Our analysis includes the following states which were most frequently represented in our set of 12K posts containing state mentions out of 2.4M: CA, NY, TX, WA, OR, FL, IL, CO, AZ and DC across 12K posts.

⁹https://huggingface.co/cjber/reddit-ner-place_names

¹⁰Understanding the socio-economic factors shaping attitudes extends beyond cost of living and unsheltered populations (Lee et al., 1990); we leave further exploration of these factors up to future work.

¹¹Cost of living index indicates the amount needed to cover basic expenses, (i.e. food, shelter) <https://worldpopulationreview.com/state-rankings/cost-of-living-index-by-state>

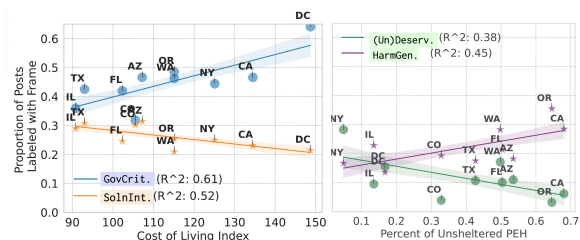


Figure 5: Linear regression results for state level factors: cost of living and percent of unsheltered PEH, predicting proportion of frames for each state. The shaded region represents the standard error of the fit to the true values.

OATH-Frames when PEH are compared to other vulnerable populations in our subset of posts labeled with (Un)Deserv. (Rosenheck et al., 1999; Chishti and Gelatt, 2022).

We identify the following subpopulations of PEH: *veterans, people, immigrants, refugees, americans* and the following vulnerable populations: *asylum seekers, immigrants, refugees, illegal aliens, Ukrainians* amongst the most frequently occurring bigrams in our data. We extract a subset of 153K posts (out of 2.4M posts) mentioning these terms. ‘Homeless Veteran’, ‘Homeless People’ and ‘Homeless American’ co-occur most frequently with ‘Ukrainian’ and ‘Immigrant’ (see Figure 13 in Appendix D). For each co-occurring (PEH, Vulnerable Population) pair, we examine the most statistically significant attitudes (see Table 11 in Appendix D). We find marked differences in attitudes when PEH are compared to Ukrainians versus when they are compared to immigrants. Public attitudes in posts that compare PEH with Ukrainians appear either as GovCrit. or appeals for MoneyAid. In contrast, posts comparing PEH to immigrants elicit more nationalistic sentiment and are oriented around relocation (NIMBY) and HarmGen.

4.4 Tracking Attitudes Across Time

Geo-political events serve as catalysts for reactive and inflammatory attitudes on social media. We find that the prevalence of certain attitudes across several months correlates with news cycles that dominated the discourse on homelessness. Figure 6 shows the frequency of attitudes by month across 2021 to 2023. We observe that GovCrit., SolnInt., (Un)Deserv., and MoneyAid. peaked in 12-2022, a time when tensions were rising in the U.S. due to Congress considering a spending bill with \$44.9 billion in aid for Ukrainians during the Russian-Ukraine

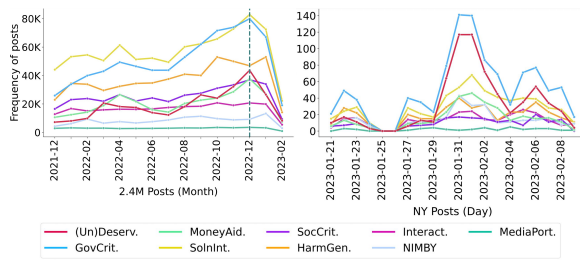


Figure 6: Frequency of posts labeled with frames across months in 2021-2023 in 2.4M annotated posts (Left) and across days in posts about NY (Right).

War (Weber, 2022). This sparked comments regarding the prioritization of government funding (GovCrit., MoneyAid.) amongst Ukraine and PEH ((Un)Deserv.). For example, the following post contains mentions of the spending bill followed by a comparison to the existing homelessness crisis in the country: 'They just approved \$45 Billion for Ukraine that'll be unaccounted for and to be spent however they please. We have homeless wandering the streets looking for a warm place to get out of the cold'.

Additionally, we investigate the prevalence of attitudes in posts about New York and find that GovCrit. and (Un)Deserv. attitudes spike significantly between 01-29-2022 and 02-02-2022. The spike correlates to a \$40M contract to convert the hotels into temporary housing for an influx of migrants (Zraick, 2023; Flanagan and Formoso, 2023) leading to protests, increased hostility towards migrants and public opposition against the government. Such comparisons between immigrants and the state's homeless population resulted in a spike in (Un)Deserv. and GovCrit. attitudes.

5 Conclusion and Future Work

We introduce OATH-Frames which characterize online attitudes towards homelessness under three overarching themes and nine Issue-specific frames. We annotate posts using Expert, GPT-4 and Expert+GPT-4 and use them to train a model to predict OATH-Frames on a large collection of 2.4M posts. We show that modeling OATH-Frames offers distinct advantages over existing models for toxicity and sentiment classification, and reveals nuanced attitudes in relation to states, marginalized populations, and over time.

We hope to expand OATH-Frames to include more context such as target groups, or correlated

factors (e.g., substance use and mental illness). Our work serves as a blueprint for building frameworks that incorporate insights from social science experts for different large-scale socially sensitive annotation tasks. We expect OATH-Frames to serve as a tool for social science practitioners in analyzing attitudes towards complex social issues *at scale*.

Acknowledgements

We would like to thank Julia Mendelsohn, Xuhui Zhou, Rob Voigt, Justin Cho, Marco Castro, Sunyoung Lie, Aida Davani, and Vinodkumar Prabhakaran for their valuable insights on our framing typology and annotation pipeline. We would also like to thank the USC NLP faculty, USC Center for AI in Society (CAIS) and students: Sayan Ghosh, Johnny Wei and Matthew Finlayson for their continuous feedback in the formulation of our project and the anonymous reviewers and area chairs in the ARR OpenReview cycle for their helpful suggestions. This research was supported by an award from the Allen Institute for AI and a USC Zumberge DEI in Research Fellowship. Brihi was supported by the USC + Amazon Center on Secure & Trusted Machine Learning.

Limitations and Ethical Considerations

LLMs as Annotation Assistants One limitation of using LLMs as annotation assistants for socially sensitive tasks is the potential for harmful biases to be reflected in the predictions. LLMs have been shown to reflect biases (Feng et al., 2023; Gallegos et al., 2023) that are considered harmful to underserved communities (Kiritchenko et al., 2023). We are still understanding the capabilities of LLMs in contextualizing and reasoning about sensitive topics such as homelessness that are contingent on complex social and political norms. As a result, we ensure our annotations are validated by experts in social science to mitigate the propagation of harms reflected by LLMs. An interesting avenue for future work would be to more effectively scale up large scale annotations for socially sensitive settings by better allocating expert intervention where necessary.

Diversity of Attitudes Given that we do not know the demographic background or geographical origin of the post, our framing typology may not necessarily be representative of attitudes towards homelessness outside the U.S. Homelessness is viewed differently depending on regional-

ity, ideology, political and economic status among other socio-political factors that might not be represented in our dataset. Our framing typology may have introduced our own prior biases in reasoning about these posts that may not be in direct alignment with the general public's attitudes towards homelessness.¹² The use of LLMs could further exacerbate societal biases in reasoning about an already marginalized population. This is an important consideration while using OATH-Frames in making generalizations about attitudes towards homelessness.

Data Collection via Keyword Search Our collection of Twitter posts was restricted to those that contained the keyword 'homeless' but there are many other terms to reference PEH, which could be explored in future work.¹³

Privacy Due to the diversity of our dataset and lexical diversity of social media posts, we anonymized user handles, and did not obtain or use any priors about the user from their profile or follower network; however, prior work has shown that demographic background could be useful (Tsai et al., 2019). We only rely on the content of the post and the timestamp (time at which the post was created) for our study.

Confounders We did not consider additional social, political and economic confounders that can point to why some attitudes are more salient than others. For example, substance use and mental illness may influence the public's opinion about PEH (Baiocchi and Arguello, 2019). OATH-Frames can serve as the foundation for identifying such factors which can further corroborate the statistical prevalence of certain attitudes over others. Furthermore, taking into consideration the target group of the post may provide more finer grained signals as to who the attitude is directed at and help better classify harmful rhetoric towards PEH (Kiritchenko et al., 2023). Initially, we considered the characterization of these additional factors which are shown

¹²Our annotation team consisted of domain and trained experts from social work and computer science from a mix of backgrounds including 4 from the U.S., 1 from India, and 1 from China.

¹³Initially, we started with a much larger set of seed words including: "homeless, homelessness, unhoused, unsheltered, encampment, tent, houseless, living on the streets, hobo, doss, bag lady, mendicant, couchsurf, housing benefit, street people, soup kitchen, roofless, vagabond, down-and-out, destitute", but we found that this returned content not necessarily related to homelessness.

in Figure 8, but it's extremely difficult and laborious to label posts at scale with more than 50+ frames and labels. We leave it up to future work to characterize additional social and political factors that can provide important and necessary context in elaborating the statistical prevalence of the frames in our dataset.

References

- Jihan Abdalla. 2022. [Why republican governors are relocating migrants across us](#).
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- P. Robert Agans and Guangya Liu. 2015. [Public attitudes toward the homeless](#). *Global Journal of Science Frontier Research*, 15(E3):1–8.
- Oluwaseun Ajao, Jun Hong, and Weiru Liu. 2015. [A survey of location inference techniques on twitter](#). *Journal of Information Science*, 41(6):855–864.
- Lauren D Applebaum. 2001. The influence of perceived deservingness on policy decisions regarding aid to the poor. *Political psychology*, 22(3):419–442.
- Alina Arseniev-Koehler, Hedwig Lee, Tyler H. McCormick, and Megan Andreas Moreno. 2016. [#proana: Pro-eating disorder socialization on twitter](#). *The Journal of adolescent health : official publication of the Society for Adolescent Medicine*, 58 6:659–64.
- Arturo Baiocchi and Tyler Arguello. 2019. *Street Talk: Homeless Discourses and the Politics of Service Provision*, pages 123–147. Springer US.
- M.A. Barnett, S.W. Quackenbush, and L.K. Pierce. 1997. [Perceptions of and reactions to the homeless: A survey of fourth-grade, high school, and college students in a small midwestern community](#). *Journal of Soc Distress Homeless*, 6:283–302.
- John R Belcher and Bruce R DeForge. 2012. [Social stigma and homelessness: The limits of social change](#). *Journal of Human Behavior in the Social Environment*, 22(8):929–946.
- Arth Bohra, Govert Verkes, Artem Harutyunyan, Pascal Weinberger, and Giovanni Campagna. 2023. [Byoc: Personalized few-shot classification with co-authored class descriptions](#). *arXiv preprint arXiv:2310.06111*.
- Amber E. Boydston and Justin H. Gross. 2013. [Identifying media frames and frame dynamics within and across policy issues](#). In *New Directions in Analyzing Text as Data Workshop*, London.
- Paul Burstein. 2020. The determinants of public policy: What matters and how much. *Policy Studies Journal*, 48(1):87–110.

- Dallas Card, Amber Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444.
- Muzaffar Chishti and Julia Gelatt. 2022. [Busing and flights of migrants by gop governors mark a new twist in state intervention on immigration.](#)
- Hyundong Cho, Shuai Liu, Taiwei Shi, Darpan Jain, Basem Rizk, Yuyang Huang, Zixun Lu, Nuan Wen, Jonathan Gratch, Emilio Ferrera, et al. 2023. Can language model moderators improve the health of online discourse? *arXiv preprint arXiv:2311.10781*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models.](#)
- Courtney Cronley. 2010. [Unraveling the social construction of homelessness.](#) *Journal of Human Behavior in the Social Environment*, 20(2):319–333.
- Georgina Curto, Svetlana Kiritchenko, Isar Nejadgholi, and Kathleen C Fraser. 2023. The crime of being poor. *arXiv preprint arXiv:2303.14128*.
- Kerry Dittmeier, Sharon H Thompson, Emma Kroger, and Nancy C Phillips. 2018. [Perceptions of homelessness: Do generational age groups and gender matter?](#) *College student journal*, 52:441.
- Rebecca Dorn, Negar Mokhberian, Julie Jiang, Jeremy Abramson, Fred Morstatter, and Kristina Lerman. 2023. Non-binary gender expression in online interactions. *arXiv preprint arXiv:2303.04837*.
- Robert M. Entman. 1993. [Framing: Toward Clarification of a Fractured Paradigm.](#) *Journal of Communication*, 43(4):51–58.
- Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. 2023. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair nlp models. *arXiv preprint arXiv:2305.08283*.
- Jacob Flanagan and Jessica Formoso. 2023. [Migrants fight plan to move them from hotel to 'inhumane' brooklyn shelter.](#)
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2023. Bias and fairness in large language models: A survey. *arXiv preprint arXiv:2309.00770*.
- Barney G. Glaser and Anselm L. Strauss. 1967. *Discovery of Grounded Theory Strategies for Qualitative Research*. AldineTransaction, London.
- Apoorva Gondimalla, Varshinee Sreekanth, Govind Joshi, Whitney Nelson, Eunsol Choi, Stephen C. Slota, Sherri R. Greenberg, Kenneth R. Fleischmann, and Min Kyung Lee. 2024. [Aligning data with the goals of an organization and its workers: Designing data labeling for social service case notes.](#) In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.
- T. Gowan. 2010. *Hobos, Hustlers, and Backsliders: Homeless in San Francisco*. University of Minnesota Press.
- Michael Heseltine and Bernhard Clemm von Hohenberg. 2024. Large language models as a substitute for human experts in annotating political text. *Research & Politics*, 11(1):20531680241236239.
- John E Hocking and Samuel G. Lawrence. 2000. [Changing attitudes toward the homeless: The effects of prosocial communication with the homeless.](#) *Journal of Social Distress and the Homeless*, 9:91–110.
- Anna Lauren Hoffmann. 2019. [Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse.](#) *Information, Communication & Society*, 22:900 – 915.
- Minda Hu, Ashwin Rao, Mayank Kejriwal, and Kristina Lerman. 2021. Socioeconomic correlates of anti-science attitudes in the us. *Future Internet*, 13(6):160.
- Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denyul. 2020. Social biases in nlp models as barriers for persons with disabilities. *arXiv preprint arXiv:2005.00813*.
- Nathan J Kim, Jessica Lin, Craig Hiller, Chantal Hildebrand, and Colette Auerswald. 2023. Analyzing us tweets for stigma against people experiencing homelessness. *Stigma and health*, 8(2):187.
- JB Kingree and Walter F Daves. 1997. Preliminary validation of the attitudes toward homelessness inventory. *Journal of community psychology*, 25(3):265–288.
- Svetlana Kiritchenko, Georgina Curto Rex, Isar Nejadgholi, and Kathleen C. Fraser. 2023. [Aporophobia: An overlooked type of toxic language targeting the poor.](#) In *The 7th Workshop on Online Abuse and Harms (WOAH)*, pages 113–125, Toronto, Canada. Association for Computational Linguistics.

- John Lalor, Yi Yang, Kendall Smith, Nicole Forsgren, and Ahmed Abbasi. 2022. [Benchmarking intersectional biases in NLP](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3598–3609, Seattle, United States. Association for Computational Linguistics.
- Barrett A Lee, Sue Hinze Jones, and David W Lewis. 1990. [Public beliefs about the causes of homelessness](#). *Social Forces*, 9:253–65.
- Lori Leibowitz and Joachim Krueger. 2005. [Attitudes and stereotypes about the homeless: A study on self-persuasion and stereotype accuracy](#). *Journal of Social Distress and the Homeless*, 14:125–150.
- Minzhi Li, Taiwei Shi, Caleb Ziems, Min-Yen Kan, Nancy F Chen, Zhengyuan Liu, and Diyi Yang. 2023. [Coannotating: Uncertainty-guided work allocation between human and large language models for data annotation](#). *arXiv preprint arXiv:2310.15638*.
- Bruce G Link, Barrett A Lee, and Paul A Toro. 1991. [Images of the homeless: Public views and media messages](#). *Housing Policy Debate*, 2(3):649–682.
- Bruce G Link, Sharon Schwartz, Robert Moore, Jo Phelan, Elmer Struening, Ann Stueve, and Mary Ellen Colten. 1995. [Public knowledge, attitudes, and beliefs about homeless people: evidence for compassion fatigue](#). *American journal of community psychology*, pages 533–55.
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. [Modeling framing in immigration discourse on social media](#). *arXiv preprint arXiv:2104.06443*.
- Burt L. Monroe, Michael P. Colaresi, and Kevin M. Quinn. 2017. [Fightin’ words: Lexical feature selection and evaluation for identifying the content of political conflict](#). *Political Analysis*, 16(4):372–403.
- Charles H. Moore, David W. Sink, and Patricia Hoban-Moore. 1988. [The politics of homelessness](#). *PS: Political Science*, 21(1):57–63.
- Lauren E Mullenbach, Nick A. D Pitas, and Benjamin Hickerson. 2023. [The role of social contact in public perceptions of homelessness in parks and public spaces](#). *International Journal of Public Administration*, 0(0):1–14.
- Deepa Narayan and Patti Petesch. 2002. [Voices of the poor from many lands](#).
- Whitney Nelson, Min Kyung Lee, Eunsol Choi, and Victor Wang. 2024. [Designing LLM-based support for homelessness caseworkers](#). In *AAAI-2024 Workshop on Public Sector LLMs: Algorithmic and Sociotechnical Design*.
- Nicholas Pangakis, Samuel Wolken, and Neil Fasching. 2023. [Automated annotation with generative ai requires validation](#). *arXiv preprint arXiv:2306.00176*.
- Marjorie J. Robertson and Milton Greenblatt. 1992. [Homelessness](#), pages 339–349. Springer US, Boston, MA.
- Robert Rosenheck, Ellen Bassuk, and Amy Salomon. 1999. [Special populations of homeless americans](#). In *Practical lessons from the 1998 national symposium on homeless research*, pages 1–27.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. [Semeval-2017 task 4: Sentiment analysis in twitter](#). In *Proceedings of the 11th international workshop on semantic evaluation (SemEval-2017)*, pages 502–518.
- Christopher Michael Rytting, Taylor Sorensen, Lisa Argyle, Ethan Busby, Nancy Fulda, Joshua Gubler, and David Wingate. 2023. [Towards coding social science datasets with language models](#). *arXiv preprint arXiv:2306.02177*.
- Megan Smith, Thomas Byrne, and Ann Elizabeth Montgomery. 2023. [Examining the relationships between exposure to homelessness and housing insecurity and attitudes toward homelessness among low- and moderate-income people](#). *Journal of Social Distress and Homelessness*, 0(0):1–12.
- John Solas. 2018. [Deserving to deserve: Challenging discrimination between the deserving and undeserving in social work](#). *Journal of Social Work Values and Ethics*, 15(2):62–70.
- Tanya Sousa, Alyssa Andrichik, Ed Prester, Katherine Rush, Colette Tano, and Micaiah Wheele. 2023. [The 2023 annual homelessness assessment report \(ahar\) to congress](#).
- Manfred Stede, Yannic Bracke, Luka Borec, Neele Charlotte Kinkel, and Maria Skeppstedt. 2023. [Framing climate change in nature and science editorials: Applications of supervised and unsupervised text categorization](#). *Journal of Computational Social Science*, 6(2):485–513.
- Carolyn J Tompsett, Paul A Toro, Melissa Guzicki, Manuel Manrique, and Zatakia Jigna. 2006. [Homelessness in the united states: assessing changes in prevalence and public opinion, 1993-2001](#). *American Journal of Community Psychology*, 37:47–61.
- Paul Toro and D McDonnell. 1992. [Beliefs, attitudes, and knowledge about homelessness: A survey of the general public](#). *American journal of community psychology*, 20:53–80.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Jack Tsai, Crystal Lee, Thomas Byrne, Robert Pietrzak, and Steven Southwick. 2017. [Changes in public attitudes and perceptions about homelessness between](#)

1990 and 2016. *American Journal of Community Psychology*, 60.

Jack Tsai, Crystal YS Lee, Jianxun Shen, Steven M Southwick, and Robert H Pietrzak. 2019. Public exposure and attitudes about homelessness. *Journal of Community Psychology*, 47(1):76–92.

Lindsey L. Upton. 2016. *Care, control, or criminalization?: Discourses on homelessness and social responses*. Ph.D. thesis, Old Dominion University.

Robin Urevich. 2023. [Residential hotels got contracts under the los angeles mayor’s homelessness program despite violations](#).

Angelina Wang, Jamie Morgenstern, and John P Dickerson. 2024. [Large language models cannot replace human participants because they cannot portray identity groups](#). *ArXiv*, abs/2402.01908.

Peter Weber. 2022. [Biden will announce patriot missiles for ukraine during oval office meeting with zelensky, officials say](#).

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

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2023. [Can large language models transform computational social science?](#)

Karen Zraick. 2023. [How manhattan hotels became refuges for thousands of migrants](#).

A OATH-Frames Discovery

We provide an overview of an earlier iteration of OATH-Frames in [Figure 7](#) where we had close to 50 Issue-specific frames across six different themes. Following the ‘coding’ process from grounded theory ([Glaser and Strauss, 1967](#)), we refined our frames after consensus discussions. In our primary iterations, we also considered annotating factors that are often conflated with homelessness. Examples of such factors include substance use and mental illness. Such factors provide important context for how public opinion can be influenced by the social and economic climate of a particular region. We provide details of these factors in [Figure 8](#) and leave this up to future work as an extension of OATH-Frames. In the following section, we outline our expert annotation OATH-Guidelines developed during *Frame Discovery* that contain clarifications of our definitions, common questions that came up during the annotation process and steps on how to address edge cases. We find that

incorporating such insights from domain experts in our prompt refinement guided by CoT explanations improves annotation performance on our task.

A.1 OATH-Guidelines

- References to homelessness as lowest point, or linguistic structure of “listing”: ‘the country is in shambles, immigration sucks, inflation is up, and homeless crisis is up’ → are **GovCrit**.
- Posts that fall under **HarmGen** are references to homelessness as lowest point, usually these posts read ‘i’d rather be homeless than ...’
- SolnInt** can be expressed in many different ways: volunteering, advocacy etc. If someone is defending PEH from harmful stereotypes, tag with both **SolnInt** and **HarmGen**.
- If someone is advocating against a specific harmful stereotype – label as **HarmGen** and **SolnInt**, otherwise, neutral sentiments and positive comments about PEH can be labeled as **Interact**.
- HarmGen** coupled with **SolnInt** means defending against a specific harmful stereotype
- HarmGen** on its own or coupled with **Interact** means person writing the post is making a harmful stereotype
- If no harmful stereotype is mentioned and person is just making a positive comment about PEH: **Interact**, possibly **SolnInt**.
- Metaphors/similes ascribing homelessness with negative qualities should be tagged with **HarmGen**.
- If comparing homelessness to other issues and discussing prioritization of issues: **(Un)Deserv**, potentially **SolnInt**.
- ‘looking homeless’ implying PEH have a dirty, unkempt appearance is **HarmGen**.
- There is a difference between one experience of PEH doing something harmful vs. generalizing this attribute to PEH as a whole. **HarmGen** referring to generalizations made about PEH population as a whole as opposed to one incident.

12. If specifically discussing how money, aid or resources should be allocated across PEH vs. other groups such as immigrants or refugees, tag with **MoneyAid**.
13. **GovCrit** also include posts which point out political parties by name explicitly → you're looking for a juxtaposition between parties, for e.g using democratic ideals specifically to critique the party's response to homelessness in comparison to republican ideals
14. **SolnInt** includes calls to action which are more directed and pointed and often more declarative in nature, For ex. 'Hurry up govsl!' or includes critiquing or discussing specific solutions
15. Statements like: "the city is so great, why do you have homelessness" should be tagged with **GovCrit**.
16. When the post is critiquing a solution or is questioning a solution of intervention, it should be augmented with additional labels from the **Critiques** theme along with **SolnInt**.
17. If there is critique of solutions funded by the government or government sponsored, then label as **GovCrit**.
18. Google names of politicians if they show up (**GovCrit** frame often names politicians, if you suspect that's what they're doing and don't recognize the name, then google to confirm)
19. Don't open links to external sources
20. If the post mentions policing of homelessness, consider tagging with **GovCrit** since this is government mandated
21. Critiquing of how the ultra-rich are donating money is **MoneyAid**.
22. discussion/debate around contributing factors to homelessness: **SolnInt** potentially
23. If someone is defending PEH from harmful stereotypes, tag with both **SolnInt** and **HarmGen**.
24. **HarmGen** on it own or coupled with **Interact** → person is making a harmful stereotype
25. If author of the post is expressing prejudice against PEH → tag with **HarmGen**. If the author of the post is calling out prejudice, tag with **HarmGen** and **SolnInt**.

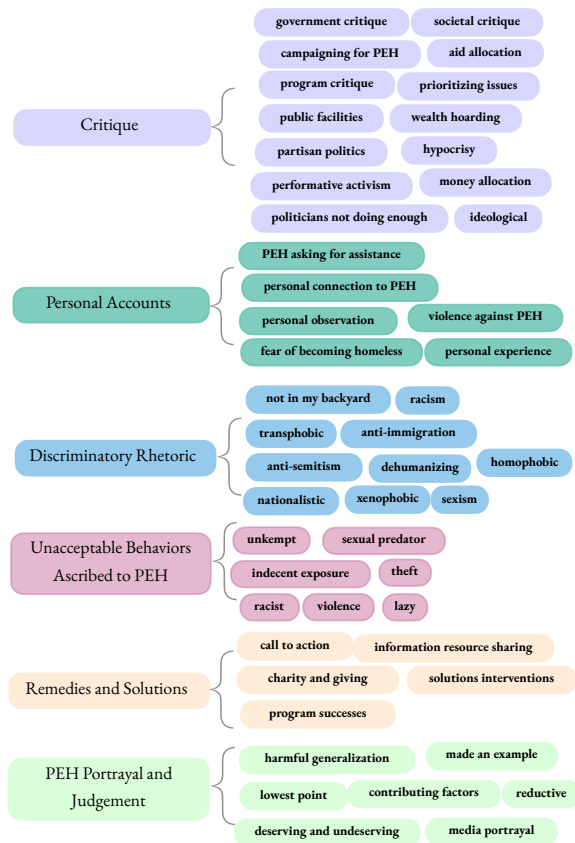


Figure 7: **OATH-Frames v1**: Our first iteration of OATH-Frames included more themes and Issue-specific frames. We went applied principles from grounding theory to iterate on our typology to more closely capture public opinions and attitudes towards homelessness.

B Annotation of OATH-Frames

We show our dataset splits for each annotation strategy in [Table 4](#).

B.1 Expert Annotations

54% of our posts are labeled with more than one Issue-specific frames. We detail the co-occurrence of our frames in [Figure 9](#) in our set of 4.1k expert annotated posts. We find that

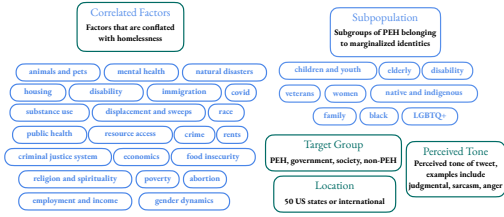


Figure 8: **Additive Labels:** We also considered additive labels that further contextualized the post with correlated factors often discussed with homelessness. We plan to incorporate additive labels in future work to further understand the social context of attitudes towards homelessness.

Annotation Method	Posts	Filtered Out	Remaining
Expert	5.3k	1.2k	4.1k
GPT-4	10k	3k	7k
Expert+GPT-4 (expert validated)	4.4k \in 7k GPT-4	300	4.1k
Flan-T5-Large (model predictions)	3.1M	700k	2.4M

Table 4: **Annotation Splits:** We include the number of posts originally annotated, the number of posts filtered out and the number of posts remaining for analysis.

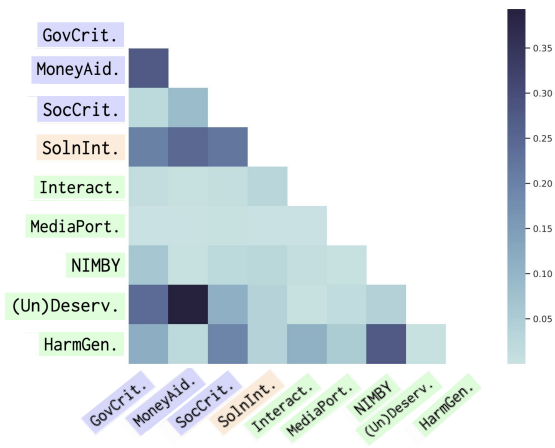


Figure 9: Proportion of posts containing co-occurrences of Issue-specific frames in our set of 2.4M annotated posts. (Each cell is normalized by the column).

(Un)Deserv. and MoneyAid. frequently co-occur. The time period of our posts is from 2021-2023, and in this time, a main agenda for the republican party was sending asylum seekers across the country, predominantly to blue states; this resulted in a lot of immigrants becoming homeless and also resulted in a news cycle that pitted immigrants against PEH in competition for jobs, housing and funding (Abdalla, 2022). Furthermore, GovCrit.

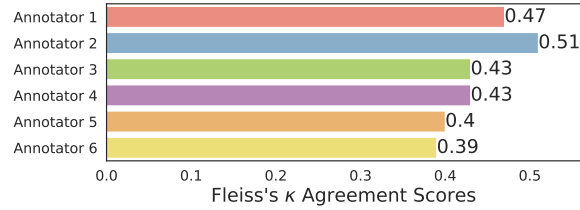


Figure 10: **Inter-annotator Agreement Scores:** We plot the average Fleiss' κ agreement scores for each expert (Annotators 1...6) where each Annotator's Fleiss' κ is computed by using every other annotator as a reference.

and MoneyAid. frequently co-occur as critique against government institutions and politicians is often in the context of funding and resource allocation. HarmGen. co-occurs frequently with NIMBY and Interact.. NIMBY rhetoric can be expressed in harmful and exclusionary ways towards PEH, and furthermore, personal interactions with PEH can elicit more harmful stereotyping.

We provide details of our inter-annotator agreement scores in Figure 10 and an example of our Amazon Mechanical Turk Sandbox annotation interface in Figure 11. The expert first determines whether the post should be filtered out if it is not relevant to public attitudes towards homelessness. For relevant posts, the expert must select at least one or more Issue-specific frames to label the post. We provide examples of our expert annotated posts in Table 5.

B.2 LLM-Assisted Expert Annotation

We include the final prompt that we achieve after iterating through the OATH-Frames instructions, for obtaining GPT-4 annotations (Table 6).

Our prompt consists of a two-step process: (1) Using the Public Opinion prompt to filter relevant posts, followed by (2) Using OATH-Frames to get multi-label Issue-specific frames for each post. We provide examples of using CoT and expert observations to incorporate OATH-Guidelines in Table 8. We also included examples of frames with the highest false discovery rates and highest precision as predicted by GPT-4 in Table 7.

C LLM-Predicted OATH-Frames: Training Details

Combining expert annotated posts along with posts that are annotated by GPT-4 and expert validated using Expert+GPT-4 in §3.2, we obtain a total of

Post	Expert Annotations
we used to have a fella come to our late night hockey games, sit in the stands and knit, one night we asked what he was knitting any why he chose a hockey arena, his response was he enjoyed watching the ladies play, and knitting scarfs and hats for the homeless was his goal.	SolnInt.
yes, there are rich people who own companies like tesla, who employ thousands of people who are a plus for our economy. they work, they get paid, they buy houses, cars, groceries, become citizens of law abiding communities, etc. they feel good about themselves. no jobs, homeless.	SolnInt. , MoneyAid.
borders were legally open (only need single lane), business closed cuz they were told to (one place tried to stay open & got harassed by police, homeless were fed & the truckers were the ones harassed, egged & called names.	(Un)Deserv.
considerable damage has already been done this isnt just for pain medication too bc so many black, brown, trans, queer poor, disabled (lol), homeless, immigrant, addicts, mentally ill, and beyond (most people in general) dont get believed by doctors and get pawnd off	HarmGen. , SocCrit.
what about our rights to a secure border!? those laws were passed by a democratic process. yet you violate your oath of office and give aid and comfort to those that break our laws. use funds to aid them,that could give our citizens m4all! look at the homeless in your state!!!	GovCrit. , (Un)Deserv. , SolnInt. , MoneyAid.
now, shall we show the millions of refugees russia created? tens of millions of homeless? yeah. russia staged a photo op where everything is fine while mariupol of half a mil has been destroyed completely by russia. good job on the ethical balance. murderers can be friendly.	GovCrit.
your not going to do anything, what makes you think most the homeless want to work, pay rent, pay taxes, most can't maintain long enough to hold or get a job. can't do drugs on the job. they would rather sit in a tent getting high. she loves her meth to much.	HarmGen.
i have a couple of ideas for the elderly homeless, and the homeless vets. too radical for most people, but, you might be interested to hear them.	SolnInt.
im not saying he is guilty or anything, i hanged around a lot ot of homeless people when i was still in high school, they are an odd bunch of people, they had public next to their sleeping, however no one used it everyone was shiting everywhere but inside the toilet	HarmGen. , Interact.
covid; mislead you, talk down the seriousness of the pandemic or letting it spread, gop, top 1%, and their lobbyist theory was; less people more money can go around, believing it would kill most of the homeless or unhealthy people that are a burden to the healthcare sector.	HarmGen. , MoneyAid. , GovCrit. , (Un)Deserv. , SolnInt.
my impression is that homelessness is highly comorbid with psychiatric illness, so i would *ideally* like to see the homeless being fed as if they were being treated for metabolic syndrome. (i'm actually involved in such an initiative, but we're temporarily on hold.)	HarmGen. , Interact. , SolnInt.
that guy his name is joey. he was part of an organization called center for political innovation, headed by a clown named caleb maupin but apparently lost all of his money, and became homeless, but still defense caleb till the end.	Interact. , GovCrit.
homeless to the shelters, edp's to the mental hospitals.	SolnInt.
why in the world are you giving homeless folks \$12,000 a piece? do denver voters know you are doing this? you are asking for drug addicts and drunks on the streets making the city look like a third-world country. another disgraceful democrat mayor.	HarmGen. , MoneyAid. , GovCrit. , NIMBY
man, that's some epic rambling bullsh** right there. what is it that you idolize about that guy? i mean, seriously, you could get the same level of coherence out of a random meth-head from the homeless shelter.	HarmGen. , SocCrit.
here we are - back to people proving who they are. should homeless people get a chance to vote? people who ran away from home and don't have a birth certificate? why can you buy a gun w/o an id, but you need one to vote? it's because the gop wants guns, but don't want poc voting.	HarmGen. , GovCrit. , SolnInt.
seriously?? he's saying the city is spending \$170 million on bike lanes, while the homeless count soars. \$170m is better spent pretty much on anything else.	MoneyAid.
and the people who still have to work, even if they get heatstroke? the homeless? those in poor housing? not everyone has the freedom to defend themselves against extreme heat - or a sociopath government and public discourse.	SocCrit. , GovCrit. , SolnInt.
i literally can't stand when ppl record themselves doing something for someone in need, like giving food to homeless ppl. it isn't ur situation to profit off of my luv	SocCrit.

Table 5: Examples of expert annotated posts with OATH-Frames.

10400 posts. We then split these posts into a train-validation-test split using the following criteria:

- We ensure roughly an 80:10:10 proportion for the above mentioned split.
- The test set contains *only* domain-expert anno-

tated posts. We pick 280 posts that were used in our agreement set in §2, and add another 1000 instances to the pool. This overall set of 1280 posts is used to pick the best performing Flan-T5 variation for final inference on a larger set of posts.

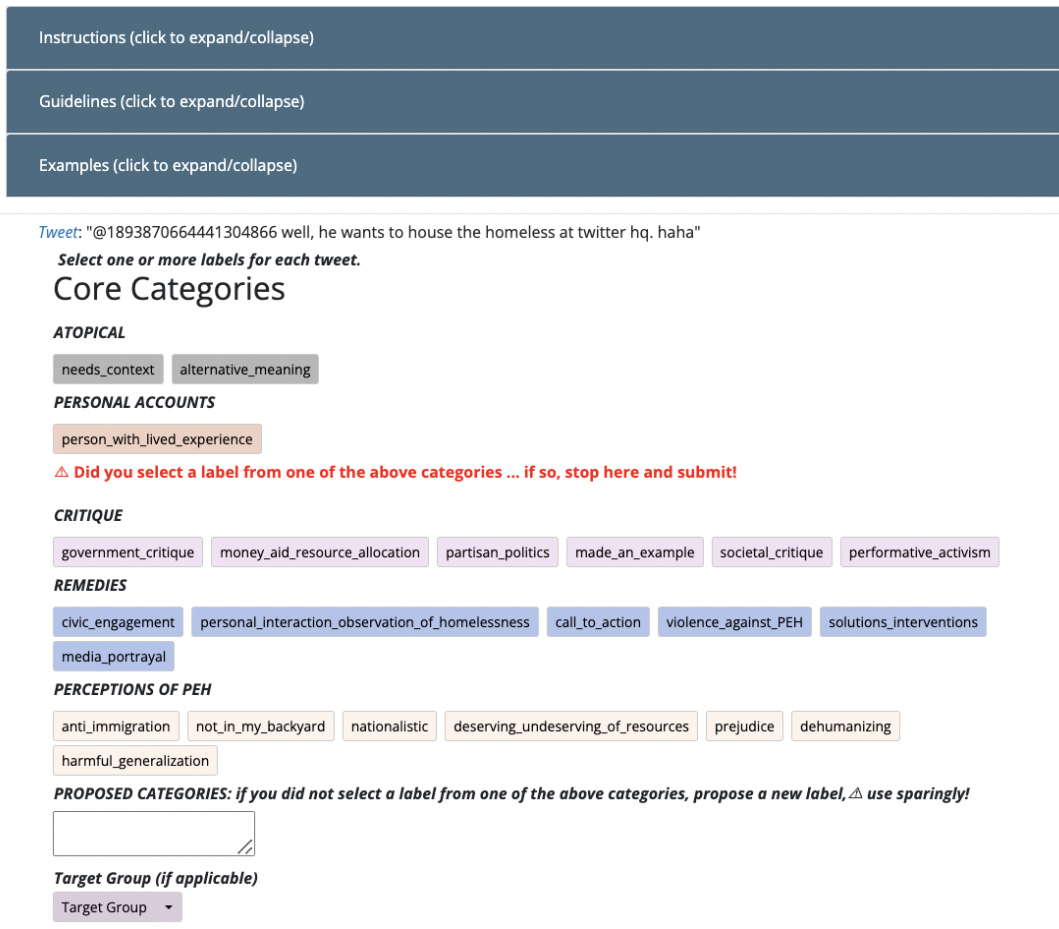


Figure 11: **Amazon Mechanical Turk Sandbox Interface for Domain / Trained Expert Annotation.** Each annotator is asked to select labels that apply the most to a given post. The time taken for each annotation is also recorded in the backend. This interface details an earlier iteration of OATH-Frames that included a larger number of Issue-specific frames.

- The remaining set of domain-expert annotated and Expert+GPT-4 annotated posts are then combined and split into a train and validation split.

We use the pre-trained Flan-T5 models for the conditional generation task, pooled from the HuggingFace¹⁴ repository. The task is structured as follows:

- **Input format:** The input format consists of the following instruction, along with the post to be labeled – Classify the given post into one or more of the following 10 labels: {list of labels} post: {}
- **Output format:** If a post needs to be filtered out using the Public Opinion filter, we generate a ‘0’ label, otherwise, we generate

a comma-separated list of Issue-specific frames. Note, this is slightly different from how GPT-4 predictions are generated where GPT-4 is prompted in a sequential manner to first filter out the posts according to the Public Opinion filter and then predict the Issue-specific frames on the remaining posts whereas we use Flan-T5 to predict the Public Opinion filter in conjunction with the 9 Issue-specific frames all at once.

We demonstrate the training details and results of different Flan-T5 variations in Table 9 and Table 12.

D Analyzing Posts with OATH-Frames

We include additional results from analyzing posts with OATH-Frames including predominant attitudes by state mentions in Figure 12 and corresponding examples of posts about states and their most

¹⁴https://huggingface.co/docs/transformers/model_doc/flan-t5

Public Opinion prompt

System Prompt: You are an AI model trained to classify tweets related to homelessness into 2 different labels. The labels are (attitude_towards_homelessness) and (other).

Label Descriptions: - (attitude_towards_homelessness): Includes tweets about homelessness to talk about another topic or are generally about the social issue of homelessness. They cover a range of topics related to homelessness, including criticism of government bodies, institutions, or political parties, discussions about the allocation and disbursement of money, aid, and resources, criticism of societal attitudes towards homelessness, debates about who is more deserving of resources, harmful generalizations or stereotypes about homeless people, opposition to having homeless people in local areas or neighborhoods, references to media portrayals of homelessness, anecdotes about interactions with homeless people, and suggestions or ideas for solving the homelessness crisis.

- (other): Includes personal anecdotes from people experiencing homelessness who are sharing their personal experience while being homeless or asking for assistance and aid. This category does NOT include tweets about fictional characters and personal interactions with other homeless people. Includes statements that are nonsensical or difficult to decipher and require access to additional resources like links, media, images, etc in order to properly interpret the tweet or references to homeless animals or being politically homeless.

Prompt: Classify the following tweet into one of the provided labels:
"post"

In concise points, please provide the relevant label that best characterizes the content of the tweet. Do not "read into" the text with interpretations, stick to the definitions of the categories strictly. The format should be the predicted label, followed by "because", followed by reason. Do not add any additional text.

Feel free to reference the label descriptions to support your classification. Provide any relevant context that influenced your classification.

OATH-Frames prompt

System Prompt: You are an AI model trained to classify tweets related to homelessness into 9 different labels. The labels include (government_critique), (money_aid_resource), (public_critique), (deserving_undeserving_of_resources), (harmful_statements_against_homelessness), (not_in_my_backyard), (media_portrayal), (interaction_with_homeless_person), and (solutions_interventions).

Label Descriptions:

- (government_critique): criticism about the government body, government institutions or political parties including critique of specific politicians, policies about homelessness, critique of programs that are being funded or considered by the government such as welfare programs, and the policing of homelessness. Also includes statements where homelessness is used as a vehicle or stand-in to talk about a broader issue portraying homelessness amongst other negative social and government problems in a list-like manner in a tweet like "murder rates, homelessness, immigration and inflation. all suck". Also includes statements that mention names of politicians.
- (money_aid_resource): Primarily includes discussion of money, for long term relief of homelessness. Includes aid or resource disbursement and allocation by government, institutions, organizations or wealthy individuals (not regular public) and also includes discussion or critique and suggestions on how the government decides to spend money and resources. Also includes discussions of giving or providing money, aid and resources to homeless people.
- (public_critique): Criticism of society in general or social norms that includes discussion of society at large instead of specific people, often pointing out hypocrisy and critiquing society's general attitudes towards homelessness. Also includes critiquing someone helping homelessness in order to gain some personal benefit where someone is being explicitly called out for doing charitable acts while filming a video or for recognition.
- (deserving_undeserving_of_resources): Discussion of competing priorities where homelessness is compared to other issues that more or less deserve aid and resources. Includes statements that express anti-immigration and support for policies, political initiatives and actions that restrict immigration often comparing and prioritizing aid to people experiencing homelessness over immigrants. Also includes nationalistic statements that prioritize one's own nation over others including discussion about prioritizing aid and relief for veterans and the nation's citizens over non-citizens.
- (harmful_statements_against_homelessness): Blanket statements that generalize a negative, harmful or undesirable attribute to all people experiencing homelessness and invoke stereotypes and make assumptions about people experiencing homelessness as a whole. Examples include statements that say all people experiencing homelessness are violent, addicts, thieves, mentally ill, unkempt, dirty, and poor at managing finances and also comparing dirty, disheveled clothing to 'looking homeless'. Includes statements that express prejudice against homelessness such as sexism, homophobia, racism, anti-semitism and transphobia or dehumanize people experiencing homelessness depriving them of positive human qualities and viewing them as sub-human or as trash. Includes statements that portray homelessness as the lowest point in one's life where homelessness is used as an example of something wrong or bad. This also includes metaphors to describe objects like anti-homeless. Could also include statements that express the desire to be violent strictly against people experiencing homelessness including threats against homelessness. Includes statements that portray homelessness as the lowest point in one's life where homelessness is used as an example of something wrong or bad. This includes listing homelessness in conjunction with other issues that are viewed as problematic or negative.
- (not_in_my_backyard): Opposition by residents to proposed developments in their local area, as well as support for strict land use regulations against wanting to see homelessness in their local area and neighborhood. Also includes displacement sweeps to remove PEH from certain areas and neighborhoods.
- (media_portrayal): Reference to a fictional character that is portraying homelessness and includes tweets and links about local news media.
- (interaction_with_homeless_person): Only includes anecdotes describing a real-life interaction with a homeless person.
- (solutions_interventions): Suggestions, remedies, problem solving and ideas for alleviating the homelessness crisis including support for policy reform, existing policies and welfare programs. Includes individual people giving money, food and help for immediate relief of homelessness. Also includes charitable acts, non-profit work, providing help and emergency aid relief, and defending people experiencing homelessness from harmful stereotypes and generalizations and advocating for positive qualities for people experiencing homelessness. Also includes call to action statements that invoke a sense of urgency in taking action towards helping the homelessness crisis.

Prompt: Classify the following tweet into one or more of the provided labels:
"post"

In concise points, carefully assess the relevant label(s) that best characterize the content of the tweet; try to list all the labels that are applicable for the tweet. Do not "read into" the text with interpretations or indications or make any assumptions, and stick to the definitions of the labels strictly. Each individual label should be followed by "because", followed by the reason for why that label was picked. Do not add any additional text. You have to select atleast one label, you cannot leave it out.

Feel free to reference the label descriptions to support your classification. Provide any relevant context that influenced your classification.

Table 6: **GPT-4 prompts:** We use the above Public Opinion and OATH-Frames prompt in sequence, to obtain GPT-4 labels.

salient attitudes in Table 10. We also include the co-occurrence of PEH and vulnerable populations in Figure 13 and corresponding examples of the most significantly occurring attitudes and examples in posts containing co-occurrences of (PEH, Vulnerable Population) pairs in Table 11 where we find marked differences amongst PEH when

they are compared to Ukrainians vs. when they are compared to immigrants in posts that reflect (Un)Deserv. . Our complete set of log-odds results from our analysis on SolnInt. in CA is detailed in Table 13 where we compute the log-odds ratio metric across CA posts comparing posts

Issue-specific Frames	Examples
(Un)Deserv. (Highest False Discovery)	<p>Post: mention we have 600,000 homeless and you sound self-centered. typical capitalist mindset. Predicted CoT: Because the post mentions the number of homeless people, implying that resources should be allocated to them, and criticizes the individual for not prioritizing this issue. Observation: The post uses homelessness as an example to critique an individual, but the CoT makes implications of PEH being in competition for resources when there is not enough information in the post to make that claim.</p> <hr/> <p>Post: mention furthermore this city has at least three homeless shelters available for these people - they just have to stay sober. Predicted CoT: Because the post implies that homeless people are undeserving of shelter if they can't stay sober, suggesting a condition for them to receive resources. Observation: The post does not imply that PEH are in competition for resources, rather it is making a harmful generalization about substance use.</p>
Interact. (Highest False Discovery)	<p>Post: mention are you helping homeless british people including veterans? Predicted CoT: Because the post is addressing someone (presumably a person or organization) about their actions towards homeless people, which can be seen as a form of interaction. Observation: It is not clear if this post is talking about a one on one interaction with PEH, the CoT is generalizing and reading into the post.</p> <hr/> <p>Post: mention wonder if any one here has offered up a room to the many thousands of irish homeless Predicted CoT: Because the post is suggesting that the individuals addressed could have a direct interaction with homeless people by offering them a room Observation: The post does not mention a direct interaction with PEH and the CoT implies that GPT-4 is reading into the post.</p>
HarmGen. (Highest Precision)	<p>Post: mention how are they a nuisance? i'd rather have that than the homeless; their garbage littering public spaces. Predicted CoT: Because the post generalizes all homeless people as a nuisance and associates them with littering public spaces, which is a negative stereotype. Observation: The CoT correctly identifies that PEH are all being portrayed as a nuisance which can be considered a negative stereotype.</p> <hr/> <p>Post: mention furthermore this city has at least three homeless shelters available for these people - they just have to stay sober. they're out there by their f*cken choice Predicted CoT: Because the post generalizes that all homeless people are choosing to stay on the streets due to their inability to stay sober, which is a harmful stereotype. Observation: The CoT reasoning correctly identifies the harmful generalization about PEH and substance use, and that all PEH are responsible for their own plight.</p>
GovCrit. (Highest Precision)	<p>Post: mention mention mention mention ur obviously not from the city cause bratton is a notorious bigot who hates the homeless, blacks and the poor. he said what he said and he meant what he meant... and what he meant was "homeless people are a nuisance." his record speaks for itself. but go on... Predicted CoT: because the post criticizes a specific government figure, Bratton, for his alleged negative attitudes towards the homeless, blacks, and the poor. The user is critiquing his policies and statements about homelessness. Observation: The CoT reasoning correctly identifies critique against William Bratton, the NYC law enforcement officer.</p> <hr/> <p>Post: mention mention big ron is the one who destroyed the institutions that held our mentally ill. he put them on the streets that added thousands of homeless and mentally ill people to live like animals. and people think the republicans are there to help all the conservatives, wrong! just the rich. Predicted CoT: because the post criticizes a specific politician ("big ron") for his policies that allegedly led to an increase in homelessness and mental illness on the streets. The post also criticizes the Republican party, implying they only help the rich. Observation: The CoT reasoning correctly identifies critique against a politician for his policies and political party values.</p>

Table 7: **GPT-4 Error Analysis:** We compare instances that are annotated using GPT-4 and provide the generated Chain-of-Thought explanations, along with expert observations of how GPT-4 is getting these examples incorrect/correct.

labeled with (+) **SoInInt.** and not labeled with (-) **SoInInt.**. We also compute the log-odds ratio metric between posts labeled with positive sentiment $p(+ve)$ and $p(-ve)$. Lastly, we provide a complete set of toxicity and sentiment distributions across all our Issue-specific frames in Figure 14 and Figure 15.

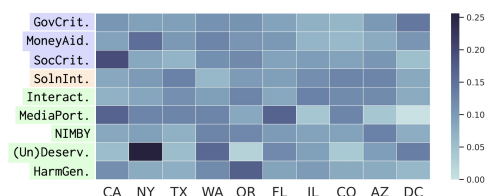


Figure 12: **Attitudes across state mentions:** Predominant attitudes in posts containing mentions of states labeled with Issue-specific frames as a fraction of total posts labeled with a specific frame (normalizing each cell by the row).

Post	GPT-4 CoT and Expert Observation	Prompt Editing: Frame Definition Clarification
what a shi*show, can't wait til your all gone, go volunteer at homeless shelters do some good	GPT-4 CoT: SocCrit. because the tweet criticizes a group of people and suggests they should do something beneficial like volunteering at homeless shelters. Expert Observation: The generated CoT is not incorrect, however, GPT-4 was not able to capture the call-to-action in the last part of the post. According to OATH-Frames-Guidelines #14: SolnInt. includes calls to action which are more directed and pointed and often more declarative in nature. We use this guideline to expand the definition of SolnInt.	Original Definition + Addition: SolnInt. : Suggestions, remedies, problem solving and ideas for alleviating the homelessness crisis including support for policy reform, existing policies and welfare programs. Includes defending people experiencing homelessness from harmful stereotypes and generalizations and advocating for positive qualities for people experiencing homelessness. Also includes individual people giving money, food and help for immediate relief of homelessness, charitable acts, non-profit work, providing help and emergency aid relief, call to action statements that invoke a sense of urgency in taking action towards helping the homelessness crisis Updated CoT: SocCrit. because the tweet criticizes unspecified individuals or groups, suggesting they should be doing more to help the homeless, such as volunteering at shelters. SolnInt. because the tweet suggests volunteering at homeless shelters as a way to help alleviate the homelessness crisis.
... hey @mention would it be able to get a better @mention package? you got it set to homeless edition and it's terrible to work with....how am i supposed to use all my crayons?	GPT-4 CoT: GovCrit. because the tweet is addressing a specific entity (presumably a government body or institution, given the context) with a critique about a package that is described as homeless edition, implying dissatisfaction with a program or policy related to homelessness. Expert Observation: It is unclear whether the post is addressing the government; thus we cannot reliably assume it is GovCrit. According to OATH-Frames-Guidelines #8: metaphors/similes ascribing homelessness with negative qualities should be labeled with HarmGen. as being in done in this post with 'homeless edition'.	Original Definition + Addition: HarmGen. : Blanket statements that generalize a negative, harmful or undesirable attribute to all people experiencing homelessness and invoke stereotypes and make assumptions about people experiencing homelessness as a whole. Examples include statements that say all people experiencing homelessness are violent, addicts, thieves, mentally ill, unkempt, dirty, and poor at managing finances and also comparing dirty, disheveled clothing to 'looking homeless'. Includes statements that express prejudice against homelessness such as sexism, homophobia, racism, anti-semitism and transphobia or dehumanize people experiencing homelessness depriving them of positive human qualities and viewing them as sub-human or as trash. Includes statements that portray homelessness as the lowest point in one's life where homelessness is used as an example of something wrong or bad. This also includes metaphors to describe objects like anti-homeless. Could also include statements that express the desire to be violent strictly against people experiencing homelessness including threats against homelessness. Updated CoT: HarmGen. because the tweet uses the term "homeless edition" in a negative context, implying that something associated with homelessness is "terrible to work with". This could be seen as a generalization that negatively portrays homelessness.
Tragedy after tragedy, we come into work and go throughout our day as conditions worsen before our eyes. Kids are murdered. Physical/mental health issues go untreated. People are unable to afford food & housing and many are dying in the streets homeless, sick and in debt.	GPT-4 CoT: SocCrit. because the tweet criticizes society's indifference towards worsening conditions, including homelessness and poverty. Expert Observation: All the problems being listed are social problems that can be considered the faults of the government. According to OATH-Frames-Guidelines #1: References to homelessness as lowest point, or linguistic structure of "listing": 'country is in shambles, immigration sucks, inflation is up, and homeless crisis is up' are GovCrit. We modify the definition using this guideline.	Original Definition + Addition: GovCrit. : Criticism about the government body, government institutions or political parties including critique of specific politicians, policies about homelessness, critique of programs that are being funded or considered by the government such as welfare programs, and the policing of homelessness. Also includes statements where homelessness is used as a vehicle or stand-in to talk about a broader issue portraying homelessness amongst other negative social and government problems in a list-like manner in a tweet like "murder rates, homelessness, immigration and inflation. everything is terrible!" Updated CoT: GovCrit. because the tweet implies a critique of the government's inability to address worsening social conditions, including homelessness.

Table 8: Examples of prompt refinement using insights from domain experts, CoT and OATH-Guidelines.

Config	Assignment
	Flan-T5-XL
	Number of parameters: 3B
	Flan-T5-Large
	Number of parameters: 770M
models	Flan-T5-Base
	Number of parameters: 250M
	Flan-T5-Small
	Number of parameters: 80M
train batch size	32
eval batch size	16
seed	0
max epochs	25
learning rate	3e-5
learning scheduler	fixed
GPU	Quadro RTX 8000
# of GPUs	8 (training), 20 (inference)
Training time	8 hours
Stopping Criteria	Early Stopping on Validation Loss

Table 9: **Model Training Details:** We provide details about our models that were finetuned on our Domain + Expert data and Expert+GPT-4 data where Flan-T5-Large achieved the highest Macro F_1 across all 9 Issue-specific frames.

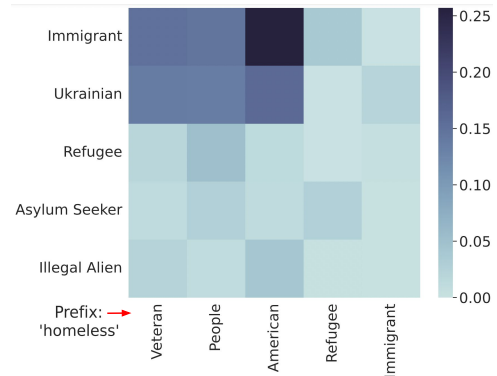


Figure 13: **Co-occurrence of PEH and vulnerable populations:** Number of posts containing mentions of PEH: 'homeless'-('veteran', 'people', 'american', 'refugee', 'immigrant') and vulnerable populations ('Immigrant', 'Ukrainian', 'Refugee', 'Asylum Seeker', 'Illegal Alien') as a fraction of the total number of posts mentioning PEH (normalizing each cell by the column).

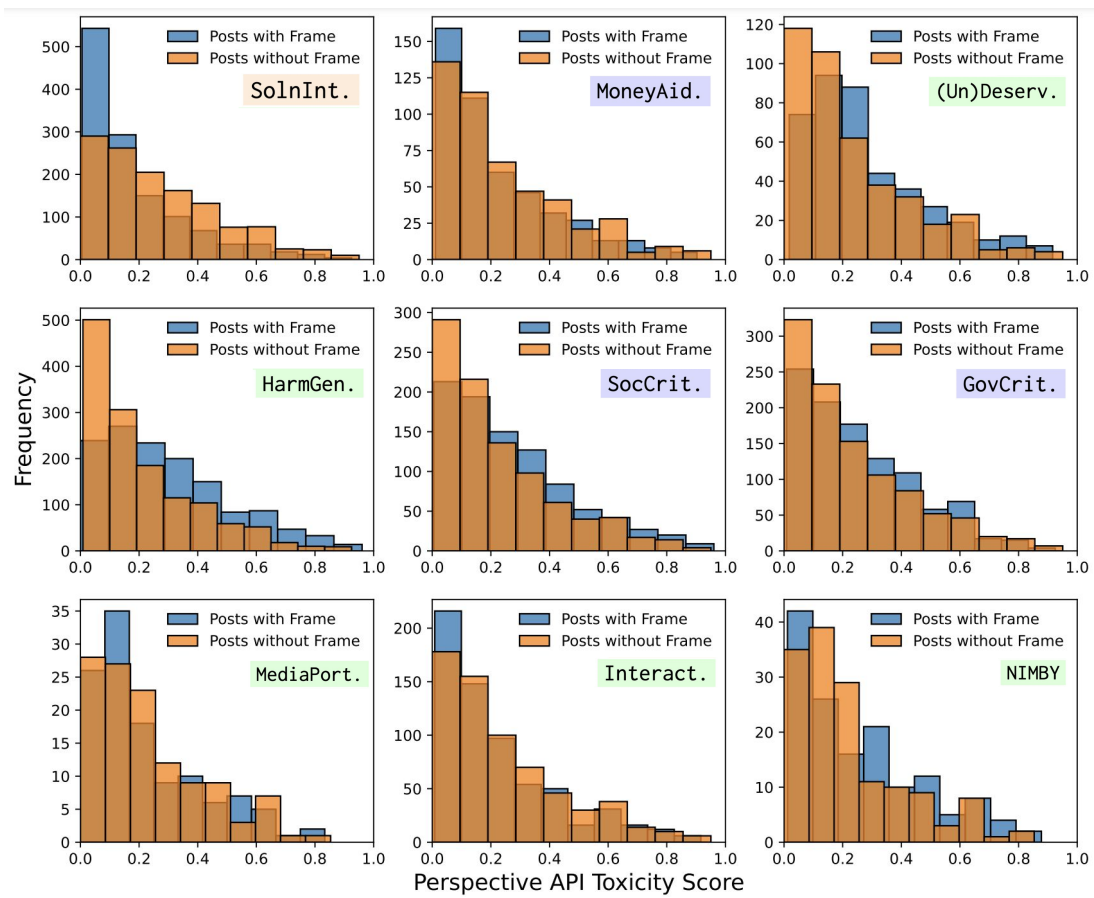


Figure 14: Distribution of toxicity scores from Perspective API toxicity classifier for posts labeled with each Issue-specific frame in our subset of 4.1k human-annotated posts.

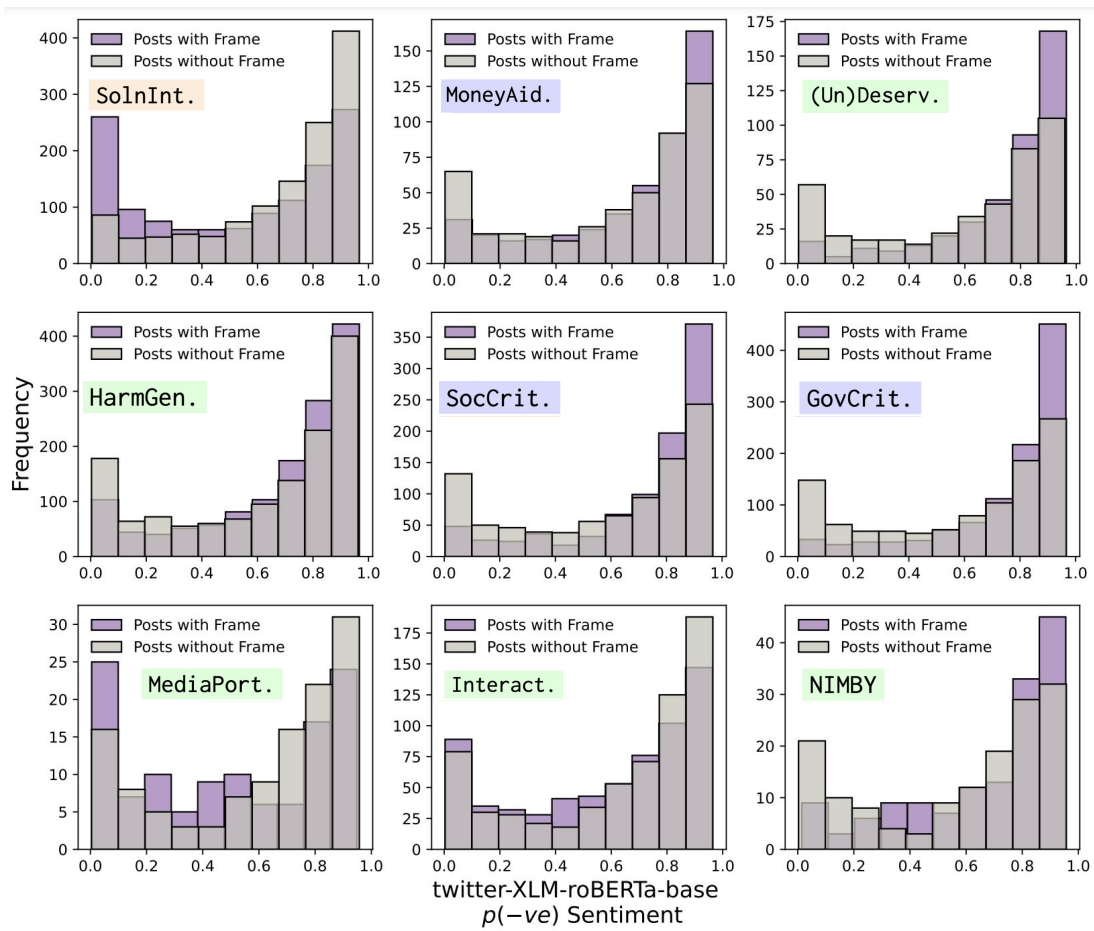


Figure 15: Distribution of twitter-XLM-roBERTa-base $p(-ve)$ sentiment probabilities (Rosenthal et al., 2017) for posts labeled with each Issue-specific frame in our subset of 4.1k human-annotated posts.

(State, Attitude)	Post
(New York, MoneyAid.)	how much money is allotted to illegal migrants? do you know? they sure are trashing out new york. everywhere they congregate there's piles of trash. how dare them through food away when we have so many homeless americans! funds need to help our people first!
(New York, (Un)Deserv.)	you're right. we should be feeding and housing our homeless veterans with our tax dollars instead of housing ungrateful illiberal immigrants in nyc high rises. also - pay your own debt.
(California, SocCrit.)	you idiots , help our homeless people first , help people in california that can't even find food or clothes.
(California, HarmGen.)	talk about how california has been ruined under you! los angeles and san francisco used to be wonderful, safe, clean cities! now they're packed with homeless, drug addicts, violent lunatics, car thieves, murderers, and more. you're a vile pig!
(Florida, MediaPort.)	google alexia son peter. he absolutely beat the shit out of a homeless man on season 3. he's also beat up his girlfriend. also he never even went to university of miami and is not allowed to be there.
(Oregon, HarmGen.)	portland street response is a joke, i'm sorry. the majority of the homeless who come to portland aren't from here and likely to be criminals. have you missed the amount of fentanyl being dealt and the amount of guns within the homeless population?

Table 10: Examples of posts containing mentions of states and their corresponding salient attitudes: Based on Figure 12, we provide example posts for the most salient attitudes across posts containing state mentions.

	Vulnerable Population	Significant OATH-Frames	Example
'Homeless' Veteran	Ukrainian	GovCrit. , MoneyAid.	@mention they have been given not 1 billion dollars, not 80 billion dollars (russian military budgets, but 180 billion dollars. do you really think that money is being used to help the ukrainian people-or is it something else? go help a homeless vet today! Post labeled with: MoneyAid. , GovCrit. , (Un)Deserv.
	Immigrant	NIMBY	@mention oh please they are dropping off bus loads of immigrants to hotels 6 blocks from my house that we the tax payers get to pay for. we have homeless vets what about them !! wake up Post labeled with: NIMBY , GovCrit. , (Un)Deserv.
'Homeless' People	Ukrainian	GovCrit. , MoneyAid.	ukrainian refugees can get ssi, housing benefits, and free health care. meanwhile, we have a lot of homeless people who have no access to such benefits - we have the highest number of the unemployed black population. Post labeled with: MoneyAid. , GovCrit. , (Un)Deserv.
	Immigrant	Interact. , NIMBY , HarmGen.	@mention new york city's pulling the rug out from underneath the illegal immigrants kicking them out of the \$700 a day hotel rooms. was it the bedbugs or the thousands of empty bottles of beer and liquor. we used to have a problem with homeless people now we have a catastrophe of homeless Post labeled with: NIMBY , GovCrit. , (Un)Deserv.
'Homeless' American	Ukrainian	GovCrit. , MoneyAid.	@mention grifter journalist got to grift . how about reporting on how close we are from being in a nuclearwar with russia. or how corrupted our politicians have become. write about how homeless americans live across this country and the us thinks sending money to ukraine is saving anyone Post labeled with: MoneyAid. , GovCrit. , (Un)Deserv.
	Immigrant	SocCrit. , SolnInt. , Interact. , MediaPort. , NIMBY , HarmGen.	@mention help homeless americans that are being discriminated against by the huge number of foreign aliens in america or give money to another county. i say americans are f***ed until we start saying no to everything immigrant. more foreign aliens = more poor americans = communism grows Post labeled with: SocCrit. , SolnInt. , (Un)Deserv.

Table 11: **Statistically significant frames for posts containing most frequent (PEH, Vulnerable Population) co-occurrence pairs.** For each subset of posts containing a (PEH, Vulnerable Population) co-occurrence pair, we find the most statistically significant frames ($p < 0.0001$). We observe that posts containing comparisons between PEH and Ukrainians elicit more Critiques whereas posts comparing PEH to immigrants elicit more NIMBY and HarmGen. attitudes.

Issue-specific	Metrics	Experts	GPT-4	Expert+GPT-4	Flan-T5-Small	Flan-T5-Base	Flan-T5-Large	Flan-T5-XL
GovCrit.	Prec.	0.69 ± 0.17	0.64 ± 0.16	0.68 ± 0.16	0.59 ± 0.15	0.66 ± 0.14	0.61 ± 0.13	0.65 ± 0.14
	Rec.	0.69 ± 0.17	0.82 ± 0.09	0.81 ± 0.09	0.79 ± 0.07	0.67 ± 0.09	0.74 ± 0.09	0.70 ± 0.10
	F1	0.65 ± 0.06	0.7 ± 0.08	0.72 ± 0.07	0.66 ± 0.09	0.65 ± 0.06	0.65 ± 0.07	0.65 ± 0.05
MoneyAid.	Prec.	0.66 ± 0.22	0.45 ± 0.14	0.5 ± 0.16	0.66 ± 0.16	0.75 ± 0.16	0.60 ± 0.23	0.54 ± 0.17
	Rec.	0.66 ± 0.22	0.87 ± 0.06	0.86 ± 0.07	0.70 ± 0.17	0.72 ± 0.16	0.61 ± 0.07	0.85 ± 0.09
	F1	0.60 ± 0.16	0.58 ± 0.14	0.61 ± 0.15	0.65 ± 0.10	0.70 ± 0.09	0.59 ± 0.18	0.63 ± 0.15
SocCrit.	Prec.	0.46 ± 0.15	0.36 ± 0.09	0.47 ± 0.09	0.28 ± 0.05	0.35 ± 0.11	0.44 ± 0.15	0.50 ± 0.15
	Rec.	0.46 ± 0.15	0.71 ± 0.12	0.56 ± 0.13	0.36 ± 0.12	0.43 ± 0.15	0.43 ± 0.11	0.42 ± 0.13
	F1	0.43 ± 0.09	0.46 ± 0.06	0.5 ± 0.05	0.31 ± 0.07	0.37 ± 0.10	0.42 ± 0.09	0.44 ± 0.10
SolnInt.	Prec.	0.70 ± 0.11	0.54 ± 0.07	0.65 ± 0.07	0.53 ± 0.06	0.54 ± 0.07	0.67 ± 0.04	0.59 ± 0.06
	Rec.	0.70 ± 0.11	0.72 ± 0.09	0.66 ± 0.1	0.66 ± 0.02	0.69 ± 0.05	0.70 ± 0.06	0.72 ± 0.03
	F1	0.69 ± 0.06	0.62 ± 0.06	0.65 ± 0.06	0.59 ± 0.04	0.60 ± 0.03	0.68 ± 0.03	0.64 ± 0.03
Interact.	Prec.	0.48 ± 0.18	0.34 ± 0.07	0.45 ± 0.1	0.54 ± 0.10	0.41 ± 0.11	0.53 ± 0.11	0.46 ± 0.13
	Rec.	0.48 ± 0.18	0.73 ± 0.13	0.66 ± 0.15	0.34 ± 0.12	0.57 ± 0.13	0.56 ± 0.14	0.58 ± 0.22
	F1	0.45 ± 0.13	0.45 ± 0.07	0.52 ± 0.07	0.41 ± 0.10	0.46 ± 0.09	0.54 ± 0.11	0.50 ± 0.15
MediaPort.	Prec.	0.73 ± 0.18	0.41 ± 0.09	0.6 ± 0.11	0.67 ± 0.47	0.47 ± 0.14	0.75 ± 0.16	0.42 ± 0.33
	Rec.	0.73 ± 0.18	0.69 ± 0.14	0.67 ± 0.15	0.17 ± 0.13	0.57 ± 0.10	0.64 ± 0.08	0.35 ± 0.26
	F1	0.70 ± 0.10	0.5 ± 0.08	0.63 ± 0.11	0.26 ± 0.20	0.51 ± 0.11	0.69 ± 0.10	0.37 ± 0.28
NIMBY	Prec.	0.20 ± 0.26	0.27 ± 0.19	0.48 ± 0.39	0.21 ± 0.37	0.30 ± 0.22	0.50 ± 0.42	0.42 ± 0.35
	Rec.	0.20 ± 0.26	0.47 ± 0.29	0.47 ± 0.29	0.05 ± 0.07	0.21 ± 0.18	0.28 ± 0.21	0.61 ± 0.34
	F1	0.17 ± 0.21	0.3 ± 0.16	0.39 ± 0.21	0.08 ± 0.12	0.23 ± 0.17	0.34 ± 0.26	0.45 ± 0.32
(Un)Deserv.	Prec.	0.61 ± 0.22	0.32 ± 0.06	0.56 ± 0.07	0.57 ± 0.25	0.64 ± 0.17	0.64 ± 0.14	0.60 ± 0.15
	Rec.	0.61 ± 0.22	0.87 ± 0.11	0.75 ± 0.17	0.33 ± 0.09	0.59 ± 0.13	0.73 ± 0.17	0.73 ± 0.17
	F1	0.58 ± 0.19	0.46 ± 0.06	0.63 ± 0.08	0.41 ± 0.13	0.60 ± 0.11	0.67 ± 0.10	0.64 ± 0.11
HarmGen.	Prec.	0.54 ± 0.11	0.53 ± 0.13	0.58 ± 0.15	0.50 ± 0.12	0.59 ± 0.19	0.61 ± 0.13	0.57 ± 0.15
	Rec.	0.54 ± 0.11	0.48 ± 0.07	0.48 ± 0.07	0.49 ± 0.10	0.30 ± 0.05	0.51 ± 0.05	0.58 ± 0.08
	F1	0.53 ± 0.07	0.49 ± 0.06	0.51 ± 0.06	0.48 ± 0.06	0.39 ± 0.07	0.55 ± 0.05	0.56 ± 0.05
Public Opinion	Prec.	0.76 ± 0.10	0.14 ± 0.2	0.14 ± 0.2	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	Rec.	0.76 ± 0.07	0.02 ± 0.03	0.02 ± 0.03	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	F1	0.75 ± 0.02	0.04 ± 0.05	0.02 ± 0.03	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Aggregated ($ D_{\text{agreement}} = 280$)	Prec.	0.58 ± 0.08	0.39 ± 0.05	0.51 ± 0.07	0.45 ± 0.08	0.47 ± 0.06	0.54 ± 0.07	0.47 ± 0.06
	Rec.	0.58 ± 0.08	0.64 ± 0.05	0.6 ± 0.05	0.39 ± 0.03	0.47 ± 0.05	0.52 ± 0.05	0.55 ± 0.03
	F1	0.55 ± 0.04	0.46 ± 0.04	0.52 ± 0.04	0.38 ± 0.04	0.45 ± 0.04	0.51 ± 0.03	0.49 ± 0.03
	Fleiss' κ	0.46 ± 0.04	0.31 ± 0.04	0.37 ± 0.04	0.24 ± 0.01	0.31 ± 0.02	0.34 ± 0.03	0.38 ± 0.04
Aggregated ($ D_{\text{test}} = 1280$)	Prec.	-	-	-	0.46	0.48	0.50	0.47
	Rec.	-	-	-	0.41	0.47	0.48	0.51
	F1	-	-	-	0.42	0.47	0.50	0.48

Table 12: Precision, Recall and F_1 for each of the Issue-specific frames and overall macro F_1 (Aggregated) on ($|D_{\text{agreement}}| = 280$). Due to compute restrictions, we do not generate GPT-4 and Expert+GPT-4 predictions on ($|D_{\text{test}}| = 1280$).

(+) SolnInt.	Score	(-) SolnInt.	Score	$p(+ve)$	Score	$p(-ve)$	Score
rent	3.72	crime	-7.30	great	4.68	drug	-3.84
voucher	2.59	drug	-6.73	love	4.22	crime	-2.74
abandon	2.48	illegal	-5.76	new	3.26	addict	-2.61
bass	4.68	shit	-5.62	thank	3.16	democrat	-2.24
healthcare	2.03	tax	-4.73	help	2.32	highest	-2.12

Table 13: Log-odds ratios for terms that are statistically more likely to appear in posts mentioning California and labeled with (+) SolnInt., posts not labeled with (-) SolnInt., posts with positive sentiment ($p(+ve) > 0.5$) and negative sentiment ($p(-ve) > 0.5$). Scores outside the range $[1.96, -1.96]$ are considered statistically significant.

E Further Related Work

E.1 Characterizing Attitudes Towards Homelessness

Prior work characterizes public opinion and attitudes towards homelessness through ethnographic studies, surveys, content analyses and psychometric analyses on small samples of human participants (Gowan, 2010; Barnett et al., 1997; Toro and McDonnell, 1992; Agans and Liu, 2015; Kingree and Daves, 1997). Prior studies have found that diversity in public attitudes towards homelessness are heavily influenced by personal experiences and socio-demographic backgrounds. More specifically, Tsai et al. (2019) conducted a survey of 541 U.S. adults on public attitudes about homelessness with the goal of determining socio-demographic characteristics of survey respondents that are associated with particular attitudes. They find that a majority of the participants who identified as female, lower income, Democrat and were personally exposed to homelessness expressed attitudes of compassion towards PEH and believed that the federal government should take more constructive measures in funding and reform policies for PEH. Similarly, Smith et al. (2023) and Mullenbach et al. (2023) study people’s prior personal experiences with homelessness and its effect on attitudes and perceptions towards PEH. Prior studies also analyze how sociodemographic backgrounds of participants influence an individual’s belief behind the cause of homelessness through surveys (Lee et al., 1990; Dittmeier et al., 2018; Agans and Liu, 2015; Hocking and Lawrence, 2000). Such studies on attitudes and perceptions towards PEH play an important role in influencing public policies that address homelessness and restrictive housing policies (Mullenbach et al., 2023; Link et al., 1995). Public opinion is heavily influenced by the social and political climate. To this end, prior work have also looked at changing attitudes towards homelessness from 1990 to 2016 (Tsai et al., 2017) and 1993 to 2001 (Tompsett et al., 2006) by comparing survey responses conducted in different years and they found increasingly more liberal and compassionate responses towards homelessness. Beyond analyzing survey responses, Baiocchi and Arguello (2019) identifies a basic topology from ethnographic studies to identify key narratives in discourse about homelessness: sin-talk, sick-talk, and system-talk which characterize the causes of homelessness enlisting criminal behavior and mental illness as pri-

mary contributing factors. Although such studies provide invaluable insights into public opinion, it is extremely challenging to collect and analyze survey responses at a larger scale. Furthermore, prior work has also identified shortcomings in existing datasets and toxicity classifiers in recognizing degrading rhetoric against poor and disadvantaged populations. This is because harmful language against marginalized groups is often masked by humor and sarcasm and includes more subtle accounts of stereotyping that may be specific to a particular vulnerable population and thus harder to recognize (Kiritchenko et al., 2023; Curto et al., 2023). As a result, there exist few specialized datasets to study language targeted specifically towards PEH (Kiritchenko et al., 2023). To this end, we plan to expand on these studies by performing large-scale analyses on social media posts annotated with OATH-Frames and providing a data resource that analyzes complex, and at times, harmful attitudes against PEH.

E.2 Quantifying Framing in Social Media Discourse.

Framing (Entman, 1993) offers one mechanism to understand the communicative power of language by characterizing what people find most important in a text. Framing has been used to analyze the politics of homelessness and crime in New York City to characterize responses to homelessness (Upton, 2016). Furthermore, framing has also been used to study discourse on social media for topics such as immigration and politics (Mendelsohn et al., 2021; Boydstun and Gross, 2013). Our line of work is most similar to prior work that studies stigma and stereotypes against PEH in the media (Kim et al., 2023; Link et al., 1991; Leibowitz and Krueger, 2005; Belcher and DeForge, 2012). More specifically, Kim et al. (2023) characterize stigma against PEH using social media posts by manually annotating 6.4k posts and defining a set of codes that characterize different stigmatizing traits attributed to PEH. We build on prior work that use framing theory to surface Issue-specific frames that characterize attitudes towards homelessness on social media and apply OATH-Frames to study attitudes at scale.

E.3 LLMs as an Annotation Assistant.

Collecting high quality annotations for socially complex, and sensitive tasks is not only a laborious and time consuming process, but also mentally

taxing. Prior work have explored using LLMs as an assistant in the annotation process to resolve ambiguities and reduce the burden off of human annotators (Li et al., 2023; Rytting et al., 2023). However, Pangakis et al. (2023) test the performance of GPT-4 across 27 annotation tasks and find that LLM performance varies widely across tasks and is largely dependent on the difficulty and subjectivity of the annotation task. They suggest evaluating generated annotations with a set of high quality expert annotations during prompt tuning. Furthermore, Rytting et al. (2023) experiment with GPT-3 as an annotator on its own and find that across four different tasks, GPT-3 performs more efficiently and as effectively as human annotators. For socially complex tasks such as homelessness, we use GPT-4 as an assistant in the annotation process while validating its predictions with human experts to leverage both the efficiency while also maintaining of quality of LLM generated annotations. It is important for us to validate GPT-4 predictions as LLMs have been shown to misrepresent identity groups which can cause downstream harms when reasoning about problems with sensitive social implications such as homelessness (Wang et al., 2024).