

Secondary Publication



Wührl, Amelie; Klinger, Roman

Recovering Patient Journeys : A Corpus of Biomedical Entities and Relations on Twitter (BEAR)

Date of secondary publication: 15.07.2024

Version of Record (Published Version), Conferenceobject

Persistent identifier: urn:nbn:de:bvb:473-irb-964630

Primary publication

Wührl, A.; Klinger, R.(2022): „Recovering Patient Journeys : A Corpus of Biomedical Entities and Relations on Twitter (BEAR)“. In: N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, S. Goggi, H. Isahara, B. Maegaard, J. Mariani (Ed., et al.), Proceedings of the Thirteenth Language Resources and Evaluation Conference, Marseille: European Language Resources Association, pp. 4439–4450, <https://aclanthology.org/2022.lrec-1.472>.

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by-nc/4.0/legalcode>

Recovering Patient Journeys: A Corpus of Biomedical Entities and Relations on Twitter (BEAR)

Amelie Wühl, Roman Klinger

Institut für Maschinelle Sprachverarbeitung, University of Stuttgart

Pfaffenwaldring 5b, 70569 Stuttgart, Germany

{amelie.wuehl, roman.klinger}@ims.uni-stuttgart.de

Abstract

For a long time, text mining and information extraction for the medical domain has focused on scientific text generated by researchers. However, their direct access to individual patient experiences or patient-doctor interactions is sometimes limited. Information provided on social media, e.g., by patients and their relatives, complements the knowledge available in scientific text. It reflects the patient’s journey and their subjective perspective on the process of developing symptoms, being diagnosed and offered a treatment, being cured or learning to live with a medical condition. The value of this type of data is therefore twofold: Firstly, it offers direct access to people’s perspectives. Secondly, it might cover information that is not available elsewhere, including self-treatment or self-diagnoses. Named entity recognition and relation extraction are methods to structure information that is available in unstructured text. However, existing medical social media corpora focused on a comparably small set of entities and relations and were focused on particular domains, rather than putting the patient into the center of analyses. With this paper we contribute a corpus with a rich set of annotation layers following the motivation to uncover and model patients’ journeys and experiences in more detail. We label 14 entity classes (incl. environmental factors, diagnostics, biochemical processes, patients’ quality-of-life descriptions, pathogens, medical conditions, and treatments) and 20 relation classes (e.g., prevents, influences, interactions, causes) most of which have not been considered before for social media data. The publicly available dataset consists of 2,100 tweets with $\approx 6,000$ entity and $\approx 3,000$ relation annotations. In a corpus analysis we find that over 80 % of documents contain relevant entities. Over 50 % of tweets express relations which we consider essential for uncovering patients’ narratives about their journeys.

Keywords: social media health mining, biomedical information extraction, BioNLP, relation extraction

1. Introduction

On social media, doctors, patients, concerned relatives or other laypeople frequently discuss medical information. Twitter posts for example contain opinions and recommendations about treatments, recounts of medical experiences, or hypotheses and assumptions about medical issues like in Figure 1. This information is by design centered around the patient. It is impacted by the patient’s journey and their subjective perspective on processes like developing symptoms, being diagnosed and offered a treatment, being cured or learning to live with a disease. This data offers direct access to people’s perspectives and covers information that is not available elsewhere, e.g., aspects that might not be considered important or difficult to assess in clinical settings. This includes, e.g., assessments of a patient’s quality of life (Table 1, Ex. 2 and 5), or which environmental factors people consider when talking about their health (Table 1, Ex. 3 and 4).

At the same time, established resources and systems for text mining and information extraction in the medical domain have mostly been centered around

scientific and biomedical text generated by researchers. Such texts seldomly focus on individual patient’s experiences or patient-doctor interactions which makes the information and knowledge contained in the text distant by nature. While scientific resources contain high quality information, many studies struggle with gender biases and population imbalance (Weber et al., 2021), which leads to blind spots in the literature. The time-consuming nature of clinical studies causes delays until information is available to practitioners. Both limitations can be mitigated by accessing social media data. Duh et al. (2016) find in fact that social media can lead to earlier detection of adverse drug reactions. While social media data has come more into focus recently, existing corpora are limited with respect to the types of entities and relations they cover. Most commonly, biomedical entity corpora focus on diseases, symptoms and drugs (Jimeno-Yepes et al., 2015; Alvaro et al., 2017, i.a.). With regards to relation detection, work on Twitter is limited to causal relations (Doan et al., 2019), or a very small number of relation classes (i.e. *reason-to-use*, *outcome-negative*, *outcome-positive*) (Alvaro et al., 2017). This leaves a gap for medical information needs. As described above, content from social media holds this type of information. Extracting it is required if we want to uncover more fine-grained aspects of patients’ medical journeys complementary to the knowledge in scientific text.

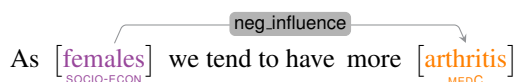


Figure 1: Example of our annotation scheme.

To facilitate research in this area, we contribute a corpus of medical tweets annotated with a fine-grained set of medical entities and relations between them. For the BEAR Corpus of *Biomedical Entities And Relations on Twitter*, we annotate 14 entity and 20 relation classes. Entities include environmental factors, diagnostics, biochemical processes, quality-of-life assessments, pathogens, as well as more established entity classes such as medical conditions, and treatments. Relation classes model how entities prevent, influence, interact with, cause or worsen other entities, or how they relate to each other as a symptom, side-effect, or diagnosis.

The dataset consists of 2,100 tweets with roughly 6,000 entities and 3,000 relations. To the best of our knowledge the majority of those classes which are centered around patient journeys have not been considered before. The dataset is available at <https://www.ims.uni-stuttgart.de/data/bioclaim>.

2. Related Work

Biomedical natural language processing (BioNLP) is an established field in computational linguistics, with a rich set of shared tasks including BioCreative and the competitions organized by the BioNLP workshop series (bio, 2021; Ben Abacha et al., 2021). Research topics include automatic information extraction from clinical reports, discharge summaries or life science articles, e.g., in the form of entity recognition for diseases, proteins, drug and gene names (Habibi et al., 2017; Giorgi and Bader, 2018; Lee et al., 2019, i.a.). A subsequent task to entity recognition is relation extraction which covers clinical relations (Uzuner et al., 2011; Wang and Fan, 2014; Sahu et al., 2016; Lin et al., 2019; Akkasi and Moens, 2021) or biomedical relations/interactions (e.g., drug-drug-interactions) between entities (Lamurias et al., 2019; Sousa et al., 2021, i.a.).

While scientific resources contain high quality information, studies might not be fully representative regarding population groups or gender (Weber et al., 2021), which leads to blind spots in the literature – the general population can barely be captured in such studies. In addition, clinical studies or reports are time-consuming which inevitably leads to delays, e.g., with regards to indications of adverse drug events. Both limitations can be mitigated by accessing social media data. Duh et al. (2016) find in fact that social media can lead to earlier detection of adverse drug reactions. This is why biomedical NLP also works with social media texts and online content (Wegrzyn-Wolska et al., 2011; Yang et al., 2016; Sullivan et al., 2016, i.a.), including established shared tasks (Magge et al., 2021a). A major focus has been to inform pharmacovigilance by identifying and extracting mentions of adverse drug reactions (Nikfarjam et al., 2015; Cocos et al., 2017; Magge et al., 2021b). Additionally, the community has

explored leveraging social media postings to monitor public health (Paul and Dredze, 2012; Choudhury et al., 2013; Sarker et al., 2016; Stefanidis et al., 2017), and detect personal health mentions (Yin et al., 2015; Klein et al., 2017; Karisani and Agichtein, 2018).

A few studies compare biomedical information in scientific documents with social media: Thorne and Klinger (2017) explore how disease names are referred to across both domains, while Seiffe et al. (2020) look into laypersons’ medical vocabulary. A related task is entity normalization which links a given mention of an entity to the respective concept in a formalized medical ontology. Limsopatham and Collier (2016) and later Basaldella et al. (2020) explore this task for medical entities on social media showcasing the difficulties in mapping laypeople’s health terminology to structured medical knowledge bases.

The ongoing COVID-19 pandemic has sparked bioNLP research to leverage or contextualize information about the disease and virus from social media. A number of studies explore detecting COVID-19-related misinformation and fact-checking (Hossain et al., 2020; Chen and Hasan, 2021; Mattern et al., 2021; Saakyan et al., 2021, i.a.). Others have looked into monitoring information surrounding the virus using social media (Cornelius et al., 2020; Hu et al., 2020).

2.1. BioNER on Social Media

Early contributions on biomedical information extraction from Twitter aimed at the extraction of adverse drug reactions from social media – a fundamentally different use case than scientific text analytics. The goal is to provide access to information even before it becomes available to doctors or researchers. This work includes corpus creation efforts on dedicated platforms like AskAPatient¹ (Karimi et al., 2015) and Twitter (Nikfarjam et al., 2015; Magge et al., 2021b)

With a similar motivation, Jimeno-Yepes et al. (2015) created Micromed, a Twitter corpus annotated with disease names, drug names, and symptom mentions. Further, TwiMed (Alvaro et al., 2017) is a dataset which combines social media and scientific text with annotations of diseases, symptoms and drug names to study drug reports across both sources. Annotated with the same entity classes, the MedRed dataset consists of Reddit posts (Scepanovic et al., 2020) labeled via crowdsourcing.

In addition to identifying entities, there has also been some work on linking them to existing databases. To facilitate this task for social media, Limsopatham and Collier (2016) contribute a Twitter corpus in which entities are linked to the SIDER 4 (Kuhn et al., 2016) database of drug profiles. Basaldella et al. (2020) subsequently introduce COMETA, a Reddit corpus in

¹<https://www.askapatient.com/>

id	Tweet
1	<p>[Prochlorperazine] <small>DRUG</small> is <small>is_type_of</small> [compazine] <small>DRUG</small>, just the generic name. [Ativan] <small>DRUG</small> also causes <small>side-effect-of</small> [drowsiness] <small>MEDC</small></p>
2	<p>[...] I was on [Lyrica] <small>DRUG</small> to help with the horrific [neuropathic pain] <small>MEDC</small> but cause <small>side-effect-of</small> [mind numbing] <small>MEDC</small> and [bowel problems] <small>MEDC</small> so I <small>cause_of</small> [stopped taking the Lyrica] <small>MEDC</small>. [I'm in more pain now but feel more like me.] <small>QOL</small></p>
3	<p>[Meditation] <small>HABITUAL</small>, [yoga] <small>HABITUAL</small> [...] are all effective at relieving [stress] <small>MEDC</small> and helping with [#IBS] <small>MEDC</small>.</p>
4	<p>[Alcohol] <small>DIETARY</small> disrupts [production of adenosine] <small>PROCESS</small> which results in [lighter sleep] <small>MEDC</small> [...]</p>
5	<p>[I'm awake just can't get going.] <small>QOL</small> Need cat food seriously only reason go out [...] [#SpoonieLife] <small>MEDC</small></p>
6	<p>[...] neighbour has been diagnosed with [c19] <small>MEDC</small> which means admin has to [self isolate] <small>THERAPY</small> and do a [test] <small>DIAGNOSTICS</small> [...]</p>
7	<p>[Support dogs] <small>OTHER</small> can improve the effectiveness of [dementia] <small>MEDC</small> therapy! Miracle creatures.</p>

Table 1: Annotated tweets from the dataset.

which entities are linked to SNOWMED-CT². With regards to the groups of entities considered (phenotype, disease, anatomy, molecule (incl. drugs, toxins, nutrients etc.), gene/DNA/RNA, device, procedure) this is similar to our contribution.

Existing resources do not cover enough entities to extract patient narratives from social media. They do not allow us yet to access the fine-grained information that social media content holds, and that would allow us to fill the information gap in scientific text.

2.2. Detection of Medical Relations on Social Media

Relation extraction contextualizes entities with each other. Medical relation extraction resources for social media are rare. Existing studies have focused on causal relations (Doan et al., 2019), or a small number of relation classes (i.e., *reason-to-use*, *outcome-negative*, *outcome-positive*) (Alvaro et al., 2017).

With regards to scientific text, and specifically clinical relation extraction, closest to our annotation scheme are approaches by Uzuner et al. (2011) and Wang and Fan (2014). Classes for both their work describe relations between treatments and medical conditions, relations between two treatments, medical conditions, or diagnoses (e.g., *treatment caused medical problem*, *treatment improved or cure medical problem*, *test reveal medical problem* in Uzuner et al. (2011), or *treats*, *prevents*, *has symptom*, *contraindicates* in Wang and Fan (2014)). However, both work with clinical

and scientific texts. Medical relation extraction on social media is understudied and missing resources that facilitate extracting patients' experiences and opinions towards entities of their medical history which would allow us to recover their medical narratives.

3. Corpus Creation

3.1. Data Collection

We collect English tweets between January 01 and November 02, 2021 using the official keyword-based Twitter API.

3.1.1. Corpus Subselection

The list of keywords to retrieve the data stems from three different sources. Refer to Table 6 for examples for each source.

1. DrugBank: DrugBank is a database for drugs which provides molecular information about drugs, their mechanisms, interactions and targets (Wishart et al., 2018). We use generic and brand/product names which allows us to collect tweets discussing treatments, or descriptions of off-label drug use.
2. MeSH: Medical Subject Headings is a controlled vocabulary thesaurus used for indexing articles in PubMed³. We use terms from the subcategories disease and therapeutics to collect tweets that address specific diseases and therapeutic measures. We use all terms that appear with a frequency ≥ 1000 in PubMed articles

²<https://www.snomed.org/>

³<https://pubmed.ncbi.nlm.nih.gov/>

hypothesizing that the distribution of those terms mirrors the usage on Twitter.

- Manual: MeSH and DrugBank mostly contain scientific terms (see Table 6), so we also query with a manually compiled list of medical terms. Partly, those relate to 10 medical conditions⁴. This is to collect tweets that either use Twitter specific hashtags, abbreviations, or community-based terms related to a condition, or mention terms generally related to the medical domain.

All terms combined result in a list of 22,874 keywords. From this list, 10,599 terms return results from Twitter during a test crawl. We remove unproductive terms and use a final list of 7,358 keywords from DrugBank, 3,120 from MeSH, and 121 from the manually compiled list.⁵ We acknowledge that by using this approach, we can not sample tweets with incorrectly spelled mentions of drug or disease names. We only keep non-duplicate tweets (based on the tweet ID) which do not contain a URL due to their increased probability of containing advertisements. Further, we only keep tweets which contain a relational term. Examples include words like *treats*, *prescribed*, or *diagnosed* (and variations thereof). From the resulting collection of tweets, we draw a sample balanced across the three keyword sources. We subsequently annotate 700 tweets per data source (350 per MeSH subcategory) which amounts in a total of 2,100 tweets.

3.2. Annotation

We label entity and relation classes that allow us to include individual aspects within people's disease-treatment cycles. Classes cover information concerning developing symptoms, being diagnosed and offered a treatment, being cured or learning to live with a medical condition. They allow us to model statements about how to self-diagnose, treat a particular condition by themselves, or capture how people perceive risk factors. For both annotation tasks, we therefore follow the central paradigm which tells annotators to label entities and relations the way a tweet's author intends or understands them. A mention like *UV radiation* could either be intended as an environmental factor (High UV radiation causes skin cancer.), or a treatment (UV radiation will help with my low vitamin D levels).

3.2.1. Entity classes

We label seven groups of entities. Each group contains a respective label or subset of labels which the annotators use to label the text. We visualize the entities in Figure 2 and depict which entity-pairs can be related. Each entity group will be briefly described in the following section. Table 1 additionally provides

⁴COVID-19, Alzheimer's disease, borderline personality disorder, cancer, depression, irritable bowel syndrome, measles, multiple sclerosis, post-traumatic stress disorder, stroke.

⁵Lists we used to collect and filter the data are available in the suppl. material together with the corpus.

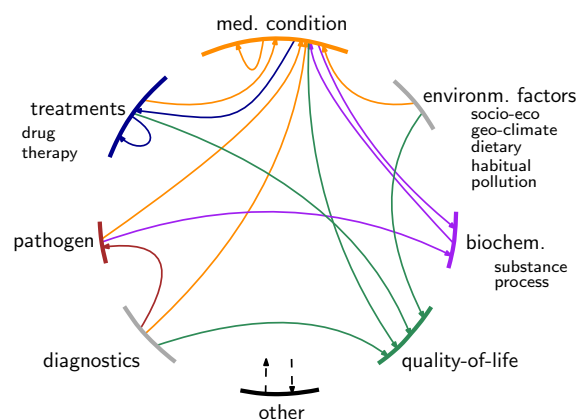


Figure 2: Visualization of entity classes and the relations between them.

fully annotated examples from the dataset, to which we will refer to in the following descriptions.

Medical Conditions. All mentions of diseases, symptoms, side effects, and medical events or descriptions thereof. See [*#IBS*] in Ex. 3 or [*drowsiness*] in Ex. 1. Phrases like [*stopped taking the Lyrica*] in Ex. 2 are considered relevant medical events, and labeled as *medC*, too.

Treatments. Mentions of any kind of treatment. That includes drug names, generic and brand names (see [*Prochlorperazine*] and [*compazine*] in Ex. 1) and all types of therapy or prevention methods (see [*self isolate*] in Ex. 6).

Environmental Factors. Entities that influence, cause or contribute to a medical condition. We annotate socio-economic (age, gender, ethnicity, social background etc.), geographic/climatic (geography, climate, weather etc.), dietary, habitual (exercise, stress etc.) or pollution-related (air/water pollution, UV or nuclear radiation) factors. See [*yoga*] in Ex. 2 or [*Alcohol*] in Ex. 4.

Pathogens. Pathogens are organisms that cause diseases. This includes mentions of bacteria, fungi, parasites, or viruses, e.g., [*coronavirus*].

Biochemical Entities. Biochemical substances such as proteins or hormones (e.g., [*Lactose*]). The class includes biochemical processes such as biological, pathogenic or chemical mechanisms (see [*production of adenosine*] in Ex. 4).

Diagnostics. Mentions of tests or other diagnostic instruments that are used to diagnose or test for a medical condition. Refer to [*test*] in Ex. 6.

Quality of Life Assessments. Descriptions of patients' quality of life, i.e. mentions of how a disease or its management impacts a patient's well-being. See [*I'm awake just can't get going*] in Ex. 5.

Other. Relevant entities that can not be covered by any of the other classes. See [Support dogs] in Ex. 7.

OTHER

3.2.2. Relation Classes

Each relation is directed and connects two entities (see Figure 2 for a depiction of which entities can be related and Table 1 for examples). We annotate the following entity pairs with relations. (\pm indicates that a relation has a positive and negative variant, e.g., (does not) treat.)

treat \rightarrow **medC** \pm treats, worsens, \pm prevents, \pm causes, contraindicates, prescribed, \pm influences

medC \rightarrow **treat** side effect of

env/pathogen/biochem \rightarrow **medC** \pm causes, \pm influences, \pm prevents

medC \rightarrow **medC/biochem** has symptom, \pm causes, is similar to

treat \rightarrow **treat** \pm interaction, is similar to

diag \rightarrow **medC/pathogen** \pm diagnoses

pathogen \rightarrow **biochem** \pm causes

medC/treat/env/diag \rightarrow **qol** \pm causes, \pm influences

general type of, other

3.2.3. Evaluation metrics

We measure the agreement between annotations by calculating the inter-annotator F_1 . Specifically, we treat one annotator's labels as the gold annotations and consider the other annotator's labels as predictions (Hripcsak and Rothschild, 2005).

We report the agreement for varying levels of strictness. We consider entity span (S) and type (T) as follows:

S1T1 The two spans and types of the entities are entirely identical.

S0T1 The two spans overlap by min. one token, entity type is identical.

S0T0 The two spans overlap by min. one token, entity type is ignored in the comparison.

When evaluating the annotated relation (R) between two entities, we consider two modes:

R1 Relation type and direction are identical.

R0 Relation type and direction are ignored in the comparison.

On the entity level, comparing S1T1 to S0T1 shows to which extent the span of an entity influences the annotation task. Comparing S0T1 to S0T0 indicates the impact of assigning a label on the difficulty of the task.

Analyzing the relation annotation follows the same objectives with respect to the entities, but adds the impact of the relation assignment. R1S1T1 is the strictest evaluation mode. The comparison to both R1S0T1 and R1S0T0 helps in understanding how the entity annotation influences the relation annotation task. R0S0T0 captures the most general level of agreement indicating how well the annotators can identify the fact that any two entities are somehow related. Comparing this to R1S0T0, we can conclude how difficult it is to identify relation types.

3.2.4. Guideline development & annotator training

We work with two in-house annotators (A1, A2) to label the tweets with entities and relations. Both annotators are female, ages 20 to 25, and 25 to 30, respectively. Their backgrounds are in linguistics and computational linguistics. They have no medical training. We iteratively train the annotators over the course of three months. In each training iteration, all annotators label a small set of instances independently following our annotation guidelines. Subsequently we discuss each set within the group. In addition, we calculate the inter-annotator F_1 for each round of training annotations (refer to Section 3.2.3 for an explanation of the eval. metrics used), and adapt the guidelines with findings from the discussions and analysis to clarify the annotation tasks further. The training instances are not part of the final corpus. The final version of the guideline document is available in the supplementary material.

Table 2 shows the development of the inter-annotator F_1 over the training iterations. For each round we report the macro F_1 score across all entity/relation classes in the different evaluation settings. We find that the agreement increases for the entities and the relation annotation over time. The agreement increases as we allow for less precise matches to be counted as true positive instances. By the end of the training period, annotators agreed with $.53F_1$ on exact entity types and boundaries (S1T1). Comparing the impact of each subtask in the last round, we observe that agreeing on the entity type is more challenging than identifying the entity span (decrease of $.25F_1$ between S0T0 and S0T1 vs. $.02F_1$ decrease between S0T1 and S1T1). Evaluating the relation type strictly (R1S0T0 vs. R0S0T0), the agreement drops by $.07F_1$ which indicates that the relation type is fairly ambiguous, and therefore hard to agree upon. The strictest evaluation measures (S1T1, R1S1T1) show that the task remains challenging even after substantial annotator training which we attribute to the diverse nature of text in tweets. Presumably, this is also why the agreement fluctuates over training rounds.

3.3. Aggregation

We provide an adjudicated version of the dataset which combines both annotators' results. In case of disagreements of entity spans between the annotators, we choose the longest overlapping sequence between two instances. We further prefer more frequent entity and relation classes over less frequent ones, and choose more general concepts over more specific ones. Generally, our aggregation strategy is motivated by a high recall approach to ensure that we lose as little of the nuances from the individual annotations as possible. We aggregate in two steps and first align the entity annotations, followed by aggregating the relations. Please refer to Section 6 for more details.

Round	Evaluation mode						
	S1T1	S0T1	S0T0	R1S1T1	R1S0T1	R1S0T0	R0S0T0
1	.44	.66	.73	.18	.4	.4	.4
2	.23	.34	.75	.18	.25	.35	.35
3	.37	.45	.79	.05	.23	.27	.59
4	.64	.76	.95	.44	.48	.51	.68
5	.39	.46	.76	.07	.1	.29	.43
6	.44	.49	.89	.12	.3	.44	.58
7	.62	.68	.96	.4	.4	.44	.47
8	.42	.61	.77	.15	.39	.41	.56
9	.69	.77	.82	.31	.33	.35	.54
10	.53	.55	.8	.28	.4	.59	.66

Table 2: Macro inter-annotator F_1 across all entity and relation classes throughout the training rounds. S1T1 through R0S0T0 indicate the evaluation mode.

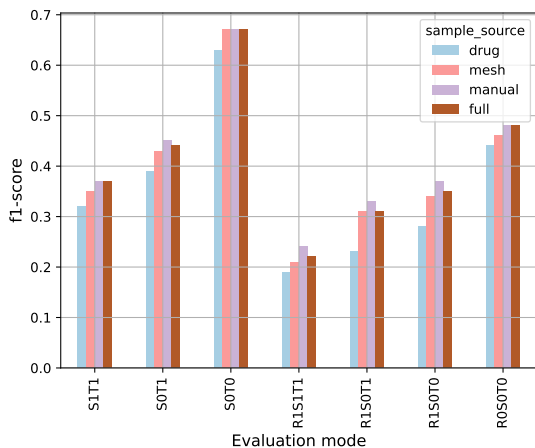


Figure 3: Inter-annotator macro F_1 -scores for each subsample of the corpus (DrugBank (drug), MeSH (mesh), manually researched keywords (manual)), and the full dataset (full) across evaluation modes.

4. Analysis

4.1. Agreement Between Annotators

The annotators labeled the final corpus over the course of four months. Since both sets of annotations provide unique perspectives on the data, we release the individual annotations along with an aggregated version. We evaluate the annotations using the inter-annotator F_1 -scores as described in Section 3.2.3 and provide scores for the full dataset as well as individual scores for each sampling method in the following. Figure 3 shows the inter-annotator F_1 -scores for each subsample of the corpus evaluated with descending strictness. For the final corpus we find that annotators are fairly synchronized in identifying entities in tweets (.67 F_1 S0T0). Agreeing on the entity type is more challenging than identifying the same entity span (.07 F_1 decrease between S0T1 and S1T1 vs. .23 F_1 decrease between S0T0 and S0T1). This is also the

case for the relation agreement. Labeling the relation type is by far the most difficult task. When we compare the agreement levels in R0S0T0 with R1S0T0, we report a difference of .13 F_1 which showcases how ambiguous the relations are.

We observe a slight decrease of the agreement compared to the last training round. We attribute this to the fact that annotators are continued to be faced with novel variations of entities and relations because of Twitter’s diverse nature.

Agreement across sources. Across all evaluation modes, tweets from the subsample Manual show the strongest agreement, followed by subsamples MeSH, and DrugBank. The results indicate that tweets from the Manual category are easier to annotate than the other documents, presumably because they mostly use laypeople’s vocabulary. Due to the nature of the DrugBank database, tweets from this set might be more scientific, making them more difficult to annotate.

Agreement across entities. Table 4 reports the inter-annotator F_1 -score (iaa) for each entity class (eval. mode: S1T1). A1 and A2 agree most strongly on instances of *medC* and *treat_drug* (.73 and .74 F_1 , respectively). We observe the lowest agreement for mentions of *biochem_process* (.05 F_1).

We observe that the agreement for highly frequent classes is stronger than the agreement in less frequent ones. Presumably, this is because these classes are also the most concrete, and therefore easier to detect. Less frequent classes (e.g., *env* or *gol*) could be considered more abstract or vague. At the same time, we presume that seeing a certain type of entity more often acts like a training effect for the annotators.

Agreement across relations. Table 5 reports the inter-annotator F_1 -score for each relation class (evaluation mode: R1S0T0⁶). Across all classes, we report a macro F_1 -score of .35. *has_symptom* and *does_not_prevent* are the classes with highest agreement (.59 F_1 respectively), followed by *treats* (.58 F_1), *may_diagnose* and *prevents* (.56 F_1 , respectively). We observe no agreement for *is_contraindicated*, *may_not_diagnose*, and *pos/neg_interaction*.

4.2. Corpus Statistics

The final corpus contains 2,100 tweets with labels for medical entities and the relations connecting them. Table 3 lists the number of documents with and without entities and relations. The majority of documents in the dataset contain entities. 86.2 % of all documents in the dataset are labeled with at least one entity. Slightly more than half of all documents containing entities also express a relevant relation (56.5 %).

The corpus consists of 93,258 words (17,559 words are unique). The longest tweet consists of 114 words, the two shortest tweets are made up of 4 words each

⁶We choose R1S0T0 to focus specifically on the relation while allowing a imprecise agreement on the entities.

	Number of documents			
	no ent	with ent	no rel	with rel
A1	330 (15.7)	1770 (84.3)	835 (47.2)	935 (52.8)
A2	378 (18.0)	1722 (82.0)	833 (48.4)	889 (51.6)
agg	289 (13.8)	1811 (86.2)	788 (43.5)	1023 (56.5)

Table 3: Number of documents with and without entities (ent) and relations (rel) for both annotators (A1, A2) and the aggregated dataset (agg). Values in parenthesis report the respective percentages. For relations this is w.r.t. all instances which contain entities.

(see Table 7). A tweet from our corpus has an average length of 44.41 words. There is no substantial difference between tweets from different sampling sources.

The following sections describe our dataset in more detail. We present corpus statistics regarding the entity and relation class distribution. Note that we describe the aggregated version of the dataset.

4.2.1. Entities

Table 4 shows the number of instances per entity class. We include the statistics for both annotators (A1, A2) and for the adjudicated dataset. Additionally, we report the statistics for the whole corpus (full), and divided by the method the documents were sampled with (DrugBank, MeSH terms, Manual).

The dataset contains 6,324 entities. The biggest entity class is medical conditions (3,553 instances), followed by mentions of *treat_drug* (1,240). The remaining entity classes are substantially less frequent. *env_pollution* has the smallest number of instances (5). Annotators label approx. 3.01 entities per document.

Entities across sources. Mentions of medical conditions are more frequent in tweets from the subsamples MeSH and Manual (1,458 and 1,367, respectively) than they are the DrugBank sample (728). Tweets from set DrugBank exhibit the majority of mentions of *treat_drug* as well as *biochem_substance* entities (1,035 and 163, respectively). Notably, mentions of the second treatment-related entity class, *treat_therapy*, are more frequent in tweets from the MeSH and Manual sample.

These results confirm that tweets in the DrugBank sample more frequently discuss treatments, and therefore exhibit a high number of drug and biochemical entities. *treat_therapy* captures more general treatment descriptions than specific mentions of drugs. Regarding the subsample Manual, we presume that the high frequency of therapy mentions indicates that laypeople speak in more general terms about treatments.

4.2.2. Relations

Table 5 reports the number of annotated relations for each class. We calculate the statistics for both annotators (A1, A2) and for the adjudicated data. We report the numbers of relations for the full corpus as well as for each of the three subsamples (DrugBank, MeSH, Manual).

In total, the corpus contains 2,959 relations. The *cause_of* relation is the most frequent (983), followed by *treats* (500), *is_type_of* (336), and *pos_influence* (263). *worsens* is the class with the lowest frequency (1 instance). For relations which can be either positive or negative, the negative relations are always less frequent. On average, a document in our dataset contains 1.41 relations.

Relations across sources. While documents from the subsample DrugBank and MeSH show relatively equal numbers of total relations (averages of 1,043 and 1,081, respectively), the Manual subsample has the least amount of relations (av. of 835). *cause_of* relations are most frequent in the subsamples MeSH (407) and Manual (331). In the DrugBank set, *treats* is the most prevalent relation class (277). Notably, for set Manual, we find that *cause_of* is by far more frequent than any other relation. All other classes count (mostly substantially) less than 100 instances each.

5. Conclusion and Future Work

We introduce and describe BEAR, a corpus of 2,100 medical tweets annotated with a detailed set of biomedical entities, and the relations connecting them. Both the entity and relation classes are motivated by the need to capture fine-grained aspects of patients' medical journeys. In our annotation study, we show that tweets hold this type of information, and that non-expert annotators can detect this reasonably well. With this dataset, we lay the groundwork to develop entity and relation extraction systems that give medical professionals access to patient narratives which are not covered in scientific texts. This includes quality-of-life assessments, perception of risk factors, unconventional treatments, or self-diagnoses that people might feel uncomfortable or irrelevant to share with their doctors. Such systems could help answer detailed questions like "How does chemotherapy affect the social life of breast cancer patients?" or "Which habits serve as coping mechanisms for people suffering from depression?".

Acknowledgments

This research has been conducted as part of the FIBISS project which is funded by the German Research Council (DFG, project number: KL 2869/5-1). We thank our annotators for their hard work and tireless attention to detail.

		Entities															
		medC	treat_drug	treat_therapy	env_geo-cli	env_diet	env_habit	env_pollution	env_socio-econ	pathogen	qol	diag	biochem_process	biochem_subst	other	total	av. #ents/doc
DB		674	1053	99	1	47	4	1	3	14	23	23	24	139	48	2153	3.08
		620	967	84	0	43	2	1	6	13	22	8	19	189	39	2013	2.88
		728	1035	114	0	34	2	1	5	7	28	13	6	163	41	2177	3.11
MeSH		1361	151	366	2	42	23	2	8	63	11	40	12	43	45	2169	3.1
		1288	108	324	3	18	23	1	20	45	15	25	8	39	58	1975	2.82
		1458	128	347	3	36	25	2	19	35	12	27	3	38	43	2176	3.11
Manual		1329	89	254	0	29	30	3	10	36	38	33	4	18	49	1922	2.75
		1276	60	226	3	29	27	0	40	41	64	21	14	31	45	1877	2.68
		1367	77	234	3	34	21	2	28	30	67	28	7	22	51	1971	2.82
full		3364	1293	719	3	118	57	6	21	113	72	96	40	200	142	6244	2.97
		3184	1135	634	6	90	52	2	66	99	101	54	41	259	142	5865	2.79
		3553	1240	695	6	104	48	5	52	72	107	68	16	223	135	6324	3.01
iaa		.73	.74	.66	.22	.36	.39	.25	.18	.43	.15	.44	.05	.42	.12		.37

Table 4: Number of annotated entities and inter-annotator F_1 (iaa) per entity class. We report the statistics across the whole corpus (full) as well as divided by the method the documents were sampled with (DB = DrugBank, MeSH = Medical subject headings, Manual = manually compiled medical keywords). Within each sampling method we report the statistics for annotator 1 and 2 and for the adjudicated dataset. Reported agreement scores (iaa, eval. mode S1T1) for all instances across the full corpus.

		Relations																					
		cause_of	does_not_prevent	does_not_treat	has_symptom	is_contraindicated	is_similar_to	is_type_of	may_diagnose	may_not_diagnose	neg_influence_on	neg_interaction	not_cause_of	pos_influence_on	pos_interaction	prescribed_for	prevents	side_effect_of	treats	worsens	other	total	av. #rels/doc
DB		215	1	13	3	3	8	138	8	0	49	12	7	67	29	9	27	35	197	1	30	852	1.22
		157	1	16	12	7	6	92	4	0	31	1	13	75	0	7	29	41	212	1	13	718	1.03
		245	1	15	11	4	1	163	7	0	50	11	12	88	28	8	31	52	277	0	39	1043	1.49
MeSH		324	0	10	81	5	3	81	14	1	45	2	10	69	3	1	30	23	132	1	10	845	1.21
		242	0	12	90	1	6	51	9	2	74	0	20	61	0	0	44	38	121	1	6	778	1.11
		407	0	11	108	5	5	102	16	3	48	1	18	92	3	0	40	39	166	1	16	1081	1.54
Manual		292	7	7	36	0	13	69	20	1	36	2	7	62	1	0	28	3	48	0	35	667	0.95
		257	8	8	62	0	8	24	17	0	91	0	17	80	0	0	31	9	40	0	8	660	0.94
		331	9	10	57	0	13	71	23	1	78	0	15	83	1	0	41	10	57	0	35	835	1.19
full		831	8	30	120	8	24	288	42	2	130	16	24	198	33	10	85	61	377	2	75	2364	1.13
		656	9	36	164	8	20	167	30	2	196	1	50	216	0	7	104	88	373	2	27	2156	1.03
		983	10	36	176	9	19	336	46	4	176	12	45	263	32	8	112	101	500	1	90	2959	1.41
iaa		.48	.59	.52	.59	.0	.55	.45	.56	.0	.25	.0	.24	.38	.0	.12	.56	.54	.58	.5	.08		.35

Table 5: Number of annotated relations and inter-annotator F_1 (iaa) per class. We report the statistics across the whole corpus (full) as well as divided by the method the documents were sampled with (DB = DrugBank, MeSH = Medical subject headings, Manual = manually researched medical keywords). Within each sampling method we report the statistics for annotator 1 and 2 and for the aggregated dataset. Reported agreement scores (iaa, eval. mode R1S0T0) for all instances across the full corpus.

6. Bibliographical References

- Akkasi, A. and Moens, M.-F. (2021). Causal relationship extraction from biomedical text using deep neural models: A comprehensive survey. *Journal of Biomedical Informatics*, 119:103820.
- Alvaro, N., Miyao, Y., and Collier, N. (2017). TwiMed: Twitter and PubMed comparable corpus of drugs, diseases, symptoms, and their relations. *JMIR Public Health Surveill*, 3(2):e24.
- Basaldella, M., Liu, F., Shareghi, E., and Collier, N. (2020). COMETA: A corpus for medical entity linking in the social media. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3122–3137, Online, November. Association for Computational Linguistics.
- Ben Abacha, A., Mrabet, Y., Zhang, Y., Shivade, C., Langlotz, C., and Demner-Fushman, D. (2021). Overview of the MEDIQA 2021 shared task on summarization in the medical domain. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 74–85, Online, June. Association for Computational Linguistics.
- (2021). *Proceedings of the BioCreative VII Challenge Evaluation Workshop*.
- Chen, Y. and Hasan, M. (2021). Navigating the kaleidoscope of COVID-19 misinformation using deep learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6000–6017, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Choudhury, M. D., Counts, S., and Horvitz, E. (2013). Social media as a measurement tool of depression in populations. In *In Proceedings of the 5th ACM International Conference on Web Science (Paris, France, May 2-May 4, 2013)*. *WebSci 2013.*, 05.
- Cocos, A., Fiks, A. G., and Masino, A. J. (2017). Deep learning for pharmacovigilance: recurrent neural network architectures for labeling adverse drug reactions in twitter posts. *Journal of the American Medical Informatics Association*, 24(4):813–821.
- Cornelius, J., Ellendorff, T., Furrer, L., and Rinaldi, F. (2020). COVID-19 Twitter monitor: Aggregating and visualizing COVID-19 related trends in social media. In *Proceedings of the Fifth Social Media Mining for Health Applications Workshop & Shared Task*, pages 1–10, Barcelona, Spain (Online), December. Association for Computational Linguistics.
- Doan, S., Yang, E. W., Tilak, S. S., Li, P. W., Zisook, D. S., and Torii, M. (2019). Extracting health-related causality from twitter messages using natural language processing. *BMC Medical Informatics and Decision Making*, 19(3):79.
- Duh, M. S., Cremieux, P., Audenrode, M. V., Vekeman, F., Karner, P., Zhang, H., and Greenberg, P. (2016). Can social media data lead to earlier detection of drug-related adverse events? *Pharmacoepidemiology and drug safety*, 25(12):1425–1433.
- Giorgi, J. M. and Bader, G. D. (2018). Transfer learning for biomedical named entity recognition with neural networks. *Bioinformatics*, 34(23):4087–4094, 06.
- Habibi, M., Weber, L., Neves, M., Wiegandt, D. L., and Leser, U. (2017). Deep learning with word embeddings improves biomedical named entity recognition. *Bioinformatics*, 33(14):i37–i48, 07.
- Hossain, T., Logan IV, R. L., Ugarte, A., Matsubara, Y., Young, S., and Singh, S. (2020). COVIDLies: Detecting COVID-19 misinformation on social media. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online, December. Association for Computational Linguistics.
- Hripcsak, G. and Rothschild, A. S. (2005). Agreement, the f-measure, and reliability in information retrieval. *Journal of the American Medical Informatics Association : JAMIA*, 12(3):296–298.
- Hu, Y., Huang, H., Chen, A., and Mao, X.-L. (2020). Weibo-COV: A large-scale COVID-19 social media dataset from Weibo. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online, December. Association for Computational Linguistics.
- Jimeno-Yepes, A., MacKinlay, A., Han, B., and Chen, Q. (2015). Identifying diseases, drugs, and symptoms in twitter. In *MEDINFO 2015:eHealth-enabled Health - Proceedings of the 15th World Congress on Health and Biomedical Informatics*, pages 643–647.
- Karimi, S., Metke-Jimenez, A., Kemp, M., and Wang, C. (2015). Cadec: A corpus of adverse drug event annotations. *Journal of Biomedical Informatics*, 55:73–81.
- Karisani, P. and Agichtein, E. (2018). Did you really just have a heart attack? Towards robust detection of personal health mentions in social media. In *Proceedings of the 2018 World Wide Web Conference*, page 137–146, Republic and Canton of Geneva, CHE.
- Klein, A., Sarker, A., Rouhizadeh, M., O’Connor, K., and Gonzalez, G. (2017). Detecting personal medication intake in Twitter: An annotated corpus and baseline classification system. In *BioNLP 2017*, pages 136–142, Vancouver, Canada., August. Association for Computational Linguistics.
- Kuhn, M., Letunic, I., Jensen, L. J., and Bork, P. (2016). The SIDER database of drugs and side effects. *Nucleic acids research*, 44:D1075–1079.
- Lamurias, A., Sousa, D., Clarke, L. A., and Couto, F. M. (2019). Bo-lstm: classifying relations via long short-term memory networks along biomedical ontologies. *BMC Bioinformatics*, 20(1):10, 01.

- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., and Kang, J. (2019). Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 09.
- Limsopatham, N. and Collier, N. (2016). Normalising medical concepts in social media texts by learning semantic representation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1014–1023, Berlin, Germany, August. Association for Computational Linguistics.
- Lin, C., Miller, T., Dligach, D., Bethard, S., and Savova, G. (2019). A BERT-based universal model for both within- and cross-sentence clinical temporal relation extraction. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 65–71, Minneapolis, Minnesota, USA, 06. Association for Computational Linguistics, Association for Computational Linguistics.
- Arjun Magge, et al., editors. (2021a). *Proceedings of the Sixth Social Media Mining for Health (#SMM4H) Workshop and Shared Task*, Mexico City, Mexico, June. Association for Computational Linguistics.
- Magge, A., Tutubalina, E., Miftahutdinov, Z., Alimova, I., Dirkson, A., Verberne, S., Weissenbacher, D., and Gonzalez-Hernandez, G. (2021b). DeepADEMiner: a deep learning pharmacovigilance pipeline for extraction and normalization of adverse drug event mentions on Twitter. *Journal of the American Medical Informatics Association*, 28(10):2184–2192, 07.
- Mattern, J., Qiao, Y., Kerz, E., Wiechmann, D., and Strohmaier, M. (2021). FANG-COVID: A new large-scale benchmark dataset for fake news detection in German. In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 78–91, Dominican Republic, November. Association for Computational Linguistics.
- Nikfarjam, A., Sarker, A., O'Connor, K., Ginn, R., and Gonzalez, G. (2015). Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *Journal of the American Medical Informatics Association*, 22(3):671–681, 03.
- Paul, M. J. and Dredze, M. (2012). A model for mining public health topics from Twitter. *Health*, 11(1).
- Saakyan, A., Chakrabarty, T., and Muresan, S. (2021). COVID-fact: Fact extraction and verification of real-world claims on COVID-19 pandemic. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2116–2129, Online, August. Association for Computational Linguistics.
- Sahu, S., Anand, A., Oruganty, K., and Gattu, M. (2016). Relation extraction from clinical texts using domain invariant convolutional neural network. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*, pages 206–215, Berlin, Germany, August. Association for Computational Linguistics.
- Sarker, A., O'Connor, K., Ginn, R., Scotch, M., Smith, K., Malone, D., and Gonzalez, G. (2016). Social media mining for toxicovigilance: automatic monitoring of prescription medication abuse from Twitter. *Drug safety*, 39(3):231–240.
- Scepanovic, S., Martin-Lopez, E., Quercia, D., and Baykaner, K. (2020). Extracting medical entities from social media. In *Proceedings of the ACM Conference on Health, Inference, and Learning, CHIL '20*, pages 170–181. Association for Computing Machinery. event-place: Toronto, Ontario, Canada.
- Seiffe, L., Marten, O., Mikhailov, M., Schmeier, S., Möller, S., and Roller, R. (2020). From witch's shot to music making bones - resources for medical laymen to technical language and vice versa. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6185–6192, Marseille, France, May. European Language Resources Association.
- Sousa, D., Lamurias, A., and Couto, F. M., (2021). *Using Neural Networks for Relation Extraction from Biomedical Literature*, pages 289–305. Springer US, New York, NY.
- Stefanidis, A., Vraga, E., Lamprianidis, G., Radzikowski, J., Delamater, P. L., Jacobsen, K. H., Pfoser, D., Croitoru, A., and Crooks, A. (2017). Zika in twitter: Temporal variations of locations, actors, and concepts. *JMIR Public Health Surveill*, 3(2):e22, Apr.
- Sullivan, R., Sarker, A., O'Connor, K., Goodin, A., Karlsrud, M., and Gonzalez, G. (2016). Finding potentially unsafe nutritional supplements from user reviews with topic modeling. In *Biocomputing 2016*, pages 528–539, Kohala Coast, Hawaii, USA, January.
- Thorne, C. and Klinger, R. (2017). Towards confidence estimation for typed protein-protein relation extraction. In *Proceedings of the Biomedical NLP Workshop associated with RANLP 2017*, pages 55–63, Varna, Bulgaria, September.
- Uzuner, O., South, B. R., Shen, S., and DuVall, S. L. (2011). 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556, 06.
- Wang, C. and Fan, J. (2014). Medical relation extraction with manifold models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 828–838, Baltimore, Maryland, June.

Association for Computational Linguistics.

Weber, A. M., Gupta, R., Abdalla, S., Cislighi, B., Meausoone, V., and Darmstadt, G. L. (2021). Gender-related data missingness, imbalance and bias in global health surveys. *BMJ global health*, 6(11).

Wegrzyn-Wolska, K., Bougueroua, L., and Dziczkowski, G. (2011). Social media analysis for e-health and medical purposes. In *2011 International Conference on Computational Aspects of Social Networks (CASoN)*, pages 278–283.

Wishart, D. S., Feunang, Y. D., Guo, A. C., Lo, E. J., Marcu, A., Grant, J. R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., Assempour, N., Iynkkaran, I., Liu, Y., Maciejewski, A., Gale, N., Wilson, A., Chin, L., Cummings, R., Le, D., Pon, A., Knox, C., and Wilson, M. (2018). DrugBank 5.0: a major update to the DrugBank database for 2018. *Nucleic acids research*, 46:D1074–D1082.

Yang, F.-C., Lee, A. J., and Kuo, S.-C. (2016). Mining health social media with sentiment analysis. *Journal of Medical Systems*, 40(11):236, Sep.

Yin, Z., Fabbri, D., Rosenbloom, S. T., and Malin, B. (2015). A scalable framework to detect personal health mentions on Twitter. *Journal of Medical Internet Research*, 17(6):e138, 06.

Appendix

Example Terms from the Sampling Methods

Table 6 shows terms from each sampling method.

Source	Example terms
DrugBank	advil, Benzylamine, Cobalt, S-Acetyl-Cysteine, Wellbutrin, zzzquil
MeSH	Anaphylaxis, Cough, Drainage, Hospitalization, Neoplasms, Self-Testing
Manual	#antivaxxer, #cancersucks, #depressionisreal, #mswarrior, #plantbasedhealing, #SocialDistancing

Table 6: Example terms from each sampling method.

Additional Examples from the Corpus

Table 7 shows the shortest and longest tweets in the dataset.

Annotation Aggregation Strategies

We provide an aggregated version of the dataset which adjudicates both annotators’ results. In general, our strategy is motivated by a high recall approach to ensure we do not lose any annotated perspectives on the data. When combining the annotations, we choose the longest overlapping sequence between two instances. We prefer more frequent entity and relation classes over less frequent ones, and choose more general concepts over more specific ones. We aggregate in two steps by first aligning the entity annotations, followed by aggregating the relations.

id	#words	tweet
1	4	bpd symptoms on 1000
2	4	Increasing pain unlocked #PTSDAwarenessDay
3	114	@username [...] @username I’ve been to every hospital in my region. ”Sorry, can’t help you” I don’t want drugs, I want my back fixed. I know for a fact the tech is there. They don’t want the liability. They should just Quit Medicine! I’m called inoperable with intractable pain, none will help. No Pain RX

Table 7: Longest and shortest tweet in the dataset.

Entities With regards to the entity span, we use the longest overlapping span between A1’s and A2’s annotation. In cases in which they disagree on the entity type, we chose the more frequent class. Exceptions are the entity classes *treatment* and *biochem*. For those classes, one subgroup is more general than the other. If both annotators agree on the major class (*treat*), but disagree on the subtype (*drug* vs. *therapy*) we aggregate to the more general one which are *treat_therapy* or *biochem_substance*.

For cases in which one annotator labeled an entity as *other* while the second annotator chose a different entity class, we aggregate to the more frequent entity class. However, if the annotator used *other* to model a relation, we keep the entity as *other* to keep the relation intact and valid.⁷

If one annotator labeled an entity, but the other one did not, we generally follow a high recall approach and add this entity to the aggregated document. However, we additionally check if the annotator who marked the entity used it to model a relation. If the relation is valid (i.e. the involved entities are allowed to be connected), we use the entity, otherwise it is dropped.

Relations To adjudicate the relation annotation, we identify cases in which both annotators agreed on the fact that there is any type of relation between a given entity pair. First, we check if the relation tags are valid (i.e. the involved entities are allowed to be connected). If one of them is invalid, we choose the valid one for the aggregated version. If both are invalid, the relation is dropped. If they are both valid, we choose the more frequent relation class. One exception to this rule concerns cases in which one annotator identified an *other* relation while the second annotator chose a different relation class. Here, the tag *other* indicates a vague relation which is not in line with our aim to adjudicate to the more specific class. Therefore,

⁷In the guidelines annotators are instructed to prioritize assigning an accurate relation over an accurate entity type. In some cases this means they may default to an *other* entity if the relation they want to model is not allowed for a particular entity pair.

we can not resolve this by simply assigning the more frequent label, because some of the small relation classes are less frequent than the class *other*. A1 and A2 consequently revisit those cases (11 instances) and decide jointly which relation type should be added to the aggregated version.

For annotations in which A1 and A2 only agreed on one of the involved entities, we follow a high recall approach and keep both relations for the adjudicated version of the data as long as the relations are valid. Finally, we consider cases in which one annotator did not label any relation while the other identified one. For those, we hypothesize that they are ambiguous and that the missing relation reflects that (i.e. that the relation marked by one of the annotators might be covering a political claim about a medical topic). In an effort not to lose these borderline cases, we add them to the aggregation as long as the relation is valid.