

From Witch’s Shot to Music Making Bones – Resources for Medical Laymen to Technical Language and Vice Versa

Laura Seiffe, Oliver Marten, Michael Mikhailov, Sven Schmeier,
Sebastian Möller, Roland Roller

German Research Center for Artificial Intelligence (DFKI),
Speech and Language Technology Lab, Berlin, Germany
firstname.secondname@dfki.de

Abstract

Many people share information in social media or forums, like food they eat, sports activities they do or events which have been visited. This also applies to information about a person’s health status. Information we share online unveils directly or indirectly information about our lifestyle and health situation and thus provides a valuable data resource. If we can make advantage of that data, applications can be created that enable e.g. the detection of possible risk factors of diseases or adverse drug reactions of medications. However, as most people are not medical experts, language used might be more descriptive rather than the precise medical expression as medics do. To detect and use those relevant information, laymen language has to be translated and/or linked to the corresponding medical concept. This work presents baseline data sources in order to address this challenge for German. We introduce a new data set which annotates medical laymen and technical expressions in a patient forum, along with a set of medical synonyms and definitions, and present first baseline results on the data.

Keywords: medical laymen to technical language, text simplification, concept normalization

1. Introduction

Every day people generate and share information online which sheds light on our lifestyle and also to a certain extent to the health situation. Provided information might include data about sports activities, food, alcohol and drug intake, but also indirectly about potential risk factors of diseases or possible adverse drug reactions, see e.g. Abbar et al. (2015) or Weissenbacher et al. (2018). Mining for instance adverse drug reactions has a high relevance for the general public as well as for pharmacological companies. As the level of medication intake is generally increasing all over the world, so does the risk of unwanted side effects (Karapetiantz et al., 2018).

In most cases, models to extract health related information from text are trained on large annotated data sets, mainly in English language, and on well formed sentences. Text in social media, forums, but also in emails, can differ in terms of sentence structure, writing style and word usage in comparison to news articles or scientific publications. Thinking particularly of health related information, the language used might be more casual and descriptive rather than the precise medical expression, as most people are not medical experts. This makes it difficult to identify the precise technical expression and to link it against a unique concept in a biomedical ontology, in order to e.g. gather further background knowledge. This makes it difficult to identify the precise technical expression and to link it to a unique concept in a biomedical ontology, in order to e.g. gather further background knowledge. For instance referring to the title of this work, patients might use laymen expressions such as ‘Hexenschuss’ (*lit.*: ‘a witch’s shot’, known as ‘*lumbago*’) or ‘Musizierknochen’ (*lit.*: ‘music making bone’, aka ‘funny bone’ or ‘ulnar nerve’) rather than their technical equivalent.

Conversely medical language might be difficult to under-

stand for non-experts. Technical terms and a special language use make it difficult to get an easy access to information that concerns the patient. The medical science is built on a vast amount of technical expressions that are not necessarily part of a patient’s everyday language. The majority of the clinical lexicon has its origin in Latin or Greek. Although the access to information is crucial for keeping track on personal conditions, for most patients the structure of the medical language remains obscure. Thus, understanding medical articles and most importantly understanding our own clinical reports written by our attending doctor may raise some challenges. In order to understand a possible serious health condition faster, automatic methods might help to simplify technical language. However, as most resources concern English language, a technical-laymen translation (and vice versa) for non-English raises further issues.

To address those challenges, this work introduces new data sets for German which support the linking of medical laymen language to technical language. Firstly we introduce a new corpus which annotates medical laymen language and technical language in a patient forum. Additionally we introduce two data sets which include different synonyms of medical concepts and sort them by complexity (rather technical to rather laymen). All data sets described in this paper will be made available¹. Our corpus in combination with the additional resources can serve as a baseline to train and to evaluate systems to map laymen into technical language and vice versa.

2. Related Work

In recent years, the biomedical domain has become an important field of research for natural language processing

¹<http://biomedical.dfki.de>

tasks. Enhancing the patient’s understanding of clinical texts is one major objective. The automatic processing of medical free text is one obstacle that is addressed by these research efforts. One step towards the processing is the mapping from free text-expressions to structured representations of domain knowledge. This includes the detection of technical terms and the normalization to an appropriate knowledge base. Synonymous expressions, terminological variants and paraphrases as well as spelling mistakes and abbreviations occur frequently in natural texts. By linking them to one unique concept, the lexical information in the text is structured and unified. In the context of medical language, different approaches face the normalization of medical concepts, such as in Leaman et al. (2013), Suominen et al. (2013) or Doğan et al. (2014).

Systems and methods that particularly address the transition from medical technical language to lay language often pursue similar approaches. Under these conditions, the linked knowledge base must provide lay language synonyms or simplified explanations for technical terms. In Zeng-Treitler et al. (2007), the Unified Medical Language System (UMLS) and especially the Consumer Health Vocabulary (CHV) are used as sources of lay vocabulary knowledge. Abrahamsson et al. (2014) conduct a synonym replacement for medical Swedish, using a system which assesses the difficulty of technical terms. If the technical term is considered as more difficult than the corresponding entry in the Swedish MeSH, the terms are replaced.

Apart from approaches that aim at simplifying the technical language, also the mapping of laymen language to medical technical expressions has gained attraction. Social media texts are a thriving resource for genuine lay language use. Recognizing meaningful elements and linking these expressions to technical counterparts allows structured insights into the health status or health related behaviour.

For example, O’Connor et al. (2014) create a data set of annotated tweets with potential adverse drug reactions. The authors test a lexicon-based approach to detect the concepts of interest. Limsopatham and Collier (2015) improve this baseline in order to normalize medical terms from social media messages using a phrase-based machine translation technique. The authors also present a system which learns the transition between lay language used in social media and the formal medical language used in descriptions of medical concepts in a standard ontology (Limsopatham and Collier, 2016).

Recently the Shared Task of *Social Media Mining for Health (SMM4H)* has gained much interest and targets this topic as well. Some of the tasks involve for instance classification of tweets presenting adverse drug reactions or vaccine behavior mentions, see Weissenbacher et al. (2019) for more information.

Now that we introduced work related to make technical expressions more comprehensible and methods to map laymen expressions to their precise equivalent and vice versa, something still remains unclear: What actually are laymen expressions and how are medical technical expressions defined?

Previous and related work does not provide a clear definition for both. Elhadad and Sutaria (2007) make use of the

contrast between a text written by a medical professional (scientific articles) and a text written by a journalist, addressing a lay audience. They consider a term as an appropriate lay expression if it is the most frequent candidate in the lay texts.

Chen et al. (2017) provide a method to rank medical terms extracted from electronic health records. The higher a term is ranked, the more urgently a lay translation is needed. Therefore they consider unithood, termhood, unfamiliarity and quality of compound term as relevant criteria for terms that must be translated for a lay audience. In contrast to these vague definitions, Grabar and Hamon (2014) concentrate on terms that show neoclassical compounding word formation. Consequently words with Latin or Greek roots are seen as technical terms.

Definition 1: (a) A medical technical term is that which is used by physicians whereas (b) a medical lay term can be easily understood by patients (medical non-experts).

Definition 2: (a) A medical term which includes (at least in parts) words with a Latin or Greek origin is defined as medical technical term. (b) All other terms belong to lay language. Lay terms are based on everyday words/language.

Table 1: Definitions used in this work of medical technical terms and laymen expressions

As there is no clear definition for technical and lay expressions, we decide to incorporate the mentioned aspects and use the definitions in Table 1. Both definitions are not entirely satisfactorily. The first definition is subjective, depends on the background of a person and requires potentially a manually generated gold standard data set. Moreover, there might be words which belong to both groups, as they are used by physicians and at the same time are understood by patients, such as *cancer*. The second definition makes it much easier to differ between both language types. However, also Latin or Greek rooted words can be very common in our daily language thus be easily understood by medical non experts, such as *hallucination*.

3. Technical-Laymen Corpus

This section introduces the Technical-Laymen Corpus (TLC) an annotated forum based on Med1.de². Med1 is a German patient forum that provides a large variety of health related topics. Users are non-professionals who seek for exchange, opinions and advice. Med1 is freely accessible and the discussions can be read without being registered. A registration is necessary to participate in the discussion. The operating team of Med1 does not provide medical consultation, however they guide the community in terms of netiquette. The users are anonymous and only their usernames are known to us. We would have been prepared to anonymize any personal data but we did not encounter data that could link to someone.

²<https://www.med1.de/forum/>

Forum	Example	Translation
Stomach-Intestines	Ja. Der Termin ist tatsächlich durch. Ich wurde an den Nieren geschallt die dort unauffällig aussehen. (Kp was das schon ausschließt) 24h Urin würde abgegeben und eine 24h Blutdruckuntersuchung angeordnet. Die haben mich komplett zerlegt: EKG Blut Spontanurin.	Hi, I am very unsure at the moment, my doctors have different opinions, some doctors say that my kidneys are not looking well, the others say that I should not be worried until GFR decreases, but what is right?
Kidney	Hallo, ich bin momentan sehr verunsichert, meine Ärzte sind nicht gleicher Meinung, die einen Ärzte sagen meine Nieren sehen nicht gut aus, die anderen sagen, solange der GFR nicht fällt muss ich mir keine Gedanken machen, was stimmt den nun?	Yes, the appointment is really over. The renal ultrasound showed no pathologies. (no idea what it can rule out) I gave 24 urine sample and a 24h blood pressure test was ordered. They have analyzed me completely: EKG, blood analysis, urine test.

Table 2: Excerpt of patient forum in German and (translated) English

Tag	Example	Annotation
L	Blut im Urin (<i>blood in urine</i>) Hexenschuss (lit.: <i>a witch's shot</i>) Eiweissverlust über die Nieren (<i>protein loss through kidneys</i>) Durchfall (lit.: <i>fall through</i>) Nierenstein-Zertrümmerung (<i>smashing of kidney stones</i>)	Hämaturie (<i>haematuria</i>) Lumbago (<i>lumbago</i>) Proteinurie (<i>proteinuria</i>) Diarrhö (<i>diarrhea</i>) Extrakorporale Stoßwellenlithotripsie (<i>extracorporeal shockwave therapy</i>)
T	Aerophagie (<i>aerophagy</i>) Appendizitis (<i>appendicitis</i>)	Luftschlucken (<i>air swallowing</i>) Blinddarmenzündung (<i>appendix infection</i>)

Table 3: Annotated examples of both tags (**L**ay, **T**echnical) from the Technical-Laymen Corpus, including translations

We are mainly interested in the medical language that is used by patients and medical laymen. A non-professional forum is likely to show the biggest source of lay language use. A corpus consisting of this kind of data should give the most realistic impression of the medical lay language. The annotation of technical and lay expressions should provide valuable insights into the relationship of technical and lay language.

For this work we selected two subforums, namely *kidney diseases* and *stomach and intestines* as text source. Each subforum provides a variety of user questions (“threads”), each containing a varying number of corresponding answers. We crawled posts of the two subforums, including the time of posting, the author’s nickname and the thread title. As the forum continuously grows, the corpus only represents the forum’s status of the crawling date. Table 2 shows two exemplary sentences from the patient forum. The examples show characteristic entries in the forum, including a specific syntax and spelling errors.

3.1. Annotation Schema

Mainly we are interested in terms and expressions that are used by medical non-professionals as those provide a large variety which cannot be entirely covered in medical dictionaries. However, as people might undergo a lifelong treatment (kidney diseases are chronic diseases) patients are well informed and also use frequently technical terms and abbreviations. For a newbie this might be difficult to understand. Thus, we target also the other direction – the detection of technical terms in order to simplify them. Our

annotation involves two different concepts: (1) *lay expressions* and (2) *technical expressions*. Regarding those information we mainly focus on *symptoms*, *diseases*, as well as *treatments and examinations*. However annotators were free to also label information that goes beyond the focus information (e.g. body parts, medication).

Annotators were asked in case of a lay expression to include the corresponding technical counterpart as well, and in case of a technical expression, the most common lay expression. We opt for a single word counterpart. If this is not possible, we choose a paraphrase or a short, appropriate explanation. In case of abbreviations we treat them accordingly: If the abbreviation is presumably known to a layman or even typical layman use (e.g. *KKH* for “Krankenhaus”, *hospital*), we annotate it as a lay expression. If the abbreviation is untypical or unlikely to be known to a patient (e.g. *NBE* for “Nierenbeckenentzündung”, *Inflammation of the Renal pelvis*) we treat it as technical term. In both cases we add the expanded version. Table 3 presents examples of the categories including their English translation.

3.2. Annotation Setup and Process

The annotation has been then carried out by two medical students within various iterations using the brat³ annotator tool (Stenetorp et al., 2012). The first annotation cycle concentrated on medically obvious cases. This means that we focused on medically clear translations from lay to technical language or vice versa. For example, the term “Normo-

³<http://brat.nlplab.org/>

tonie” (*normotonia*) is assigned the tag *technical* and the corresponding lay expression “normaler Blutdruck” (*normal blood pressure*) is given as free text.

However the results of the first cycle were not satisfying yet, as most translations were already well documented in existing vocabularies. Therefore we extended the annotations by including cases in which a non-professional describes a medical concept in such way that a definite technical translation is difficult. For example, if a user describes *problems with passing water* (“Probleme beim Wasserlassen”), a possible technical equivalent could be *dysuria*.

From the medical point of view, this procedure is difficult because it includes to some extent interpretation work: While *problems with passing water* is only a rough symptom description, a dysuria is a pathological state. The transfer from a symptom description to a disease can be seen as kind of diagnostic process which must be avoided at that point. As the annotation was carried out by medical students we trusted their expertise to decide at which point the annotation would exceed a reasonable interpretation. Thus we do not opt for a diagnostic interpretation of symptoms. In order to retrace such cases, the annotators highlighted annotated terms that came close to a critical interpretation level.

Within a final iteration one of the authors examined the annotations and highlighted potential errors (wrong labels, missing information etc.). Those highlighted information were then again manually examined, in order to provide a corpus with an appropriate quality.

3.3. Corpus Analysis

Table 4 provides an overview about TLC. The table lists for each forum topic the number of included files, number of tokens, as well as the average number of tokens per file and the average number of annotations per file. Note that not all files included relevant information to be annotated. A more detailed overview about the annotated information itself is presented in Table 5. The table lists the the number of overall and number of unique annotations for each label. As the table shows, the most annotated labels are laymen expressions. Moreover those expressions also have the largest variety in terms of different unique terms. This makes sense and highlights the importance detecting laymen expressions.

	Kidney	Stomach-Intestines
Number of files	2000	2000
Number of tokens	203,553	234,914
Avg. tokens /file	101.78	117.46
Avg. annotations /file	2.52	1.41

Table 4: General overview about Med1 Corpus

4. Additional Resources and Methods to Process Technical-Laymen Language

In addition to the Technical-Laymen Corpus we extract data from two additional resources: UMLS and Wiktionary. We

Label	#Annotations	#Unique
Lay Expression	4727	1246
Technical Term	1745	376

Table 5: Overview about number of annotated and unique concepts of each category label.

aim at providing assorted data sets which incorporate a matching of technical and laymen language in the biomedical domain. Both resources are processed and can be used to support the linking from laymen to technical terms and vice versa. However as both resources do not systematically differ between lay and technical terms, we additionally propose a simple method to identify technical (and less technical) terms.

4.1. UMLS Synonym Subset

The Unified Medical Language System (UMLS) is a biomedical ontology and knowledge source. The Metathesaurus of UMLS provides a vocabulary database for the biomedical and health domain. Synonymous expressions are linked by the same concept unique identifier (CUI). The same CUI also links equivalent expressions in different languages. The Semantic Network of UMLS categorises all terms into broad subject categories, providing a categorization into 127 semantic types (STY) and 54 relation types (RL). Overall UMLS includes concepts of over 34 million concepts in English language, whereas only approximately 100,000 in German. Roughly half of those concepts include at least two mentions. While the German UMLS subset is relevant for concept normalization in general, particularly concepts including synonyms are interesting, as they might include technical and laymen expressions.

4.2. Wiktionary Synonym Subset

Our second resource is build from the German version of Wiktionary⁴. Wiktionary provides 741,260 (Jan 2019) entries in German. Although biomedical information is not a special focus of Wiktionary, there is a large range of related subcategories. In order to create our technical/laymen language resource the (in November 2019 newest) German Wiktionary dump has been downloaded and further processed and filtered to our needs. In order to build a technical/laymen language resource from Wiktionary, we parsed the provided dump and automatically gathered for each entry the term, its explanation and, if available, synonyms. Our focus is the biomedical domain, thus we limited the data by selecting medical related entries only. These entries come from the categories *Medicine*, *Pharmacy*, *Pharmacology*, *Anatomy*, *Psychiatry*, *Psychology*, *Physiology*, *Ophthalmology*, *Pathology*, *Dentistry*, *Gynaecology* and *Dermatology*. Additionally, we included every entry that contains at any place the regular expression *krank (sick)* which should relate to mentions of diseases. By doing so, the resulting resource is larger than necessary (e.g. some veterinary entries are included). However we ensure to make

⁴<https://de.wiktionary.org/wiki/Wiktionary:Hauptseite>

Term	Explanation	Synonym
Dialyse	Anwendung der Dialyse, vor allem zur Reinigung von Blut	Blutreinigung; Blutwäsche
Diabetes	Stoffwechselerkrankung, bei der eine gesteigerte Unempfindlichkeit gegenüber Insulin besteht (sogenannter Diabetes mellitus Typ 2 oder Typ-2-Diabetes oder Altersdiabetes)	Zuckerkrankheit; Zucker
Delirium tremens	Ernste und potentiell lebensbedrohende Komplikation im Alkoholzug bei einer schon länger bestehenden Alkoholkrankheit	Alkoholdelir; Önomanie; Säufervahn; Säufervahnsinn

Table 6: Example of extracted information from Wiktionary

CUI	English	German	Spanish	French	Swedish	Russian
C0007097	carcinoma	Karzinom	carcinoma	carcinome	Karzinom	KARTSINOMA
C0012503	Dioxins	Dioxine	Dioxinas	Dioxines	Dioxiner	DIOKSINY
C0023531	Leukoplakia	Leukoplakie	Leucoplaquia	Leucoplasie	Leukoplaki	LEUKOPLAKIJA
C0027804	Neurasthenia	Neurasthenie	neurastenia	Neurasthénie	Neurasteni	NEVRASTENIA

Table 7: Similar mentions of different languages in UMLS linked by the same concept unique identifier (CUI).

use of all entries that could be relevant. Only entries of the mentioned categories were used for our resource. The final biomedical Wiktionary subset comprises 4468 concepts and nearly all including a definition. 2155 of the entries include at least one synonym. Overall this subset includes 8657 different entries.

Even though the data set appears to be small in comparison to UMLS, an interesting aspect about Wiktionary is the variety of laymen synonyms. It includes lay expressions which are often not covered by UMLS. Table 6 shows some examples: Diabetes for instance is characterized by recurrent or persistent high blood sugar. A non-professional German term for diabetes is “Zuckerkrankheit” (lit.: *sugar disease*) or simply “Zucker” (*sugar*). These terms, even though frequently used, are not listed in UMLS. The large variety of lay expressions includes not only lay expressions to the respective technical term but also colloquial or even vulgar terms. For example, the entry of “Diarrhoe” (*diarrhea*) lists as synonyms “Schnelle Katharina” (*fast Katharina*) and “Flotter Otto” (*quick Otto*).

4.3. Aligning data sets

UMLS is frequently used for concept normalization and it comprises much more concepts than the Wiktionary subset. Conversely, Wiktionary appears to be a highly useful resource as it contains more casual expressions in medical context. For this reason we try to combine both data sets. For this, we identify expressions from Wiktionary which also occur in UMLS. If a term from Wiktionary also occurs within exactly one CUI in UMLS, we can simply align the Wiktionary concept with all its synonyms to this CUI. For instance if the Wiktionary term ‘pain’ (and all its synonyms) would occur only in context of one single UMLS-CUI, we can map the Wiktionary term ‘pain’ and all its synonyms to this corresponding CUI. However, this is not possible in all cases, as terms in UMLS might be assigned to various CUIs.

In this way, 768 CUIs can be extended by overall 3082 additional mentions. We refer to the resulting data set as

Wiktionary-UMLS (WUMLS).

4.4. Sorting Synonyms

The mapping from technical to laymen language is one of the aspects of this work. However, the largest of our supporting resources, UMLS, does not provide any information about technical or laymen language for German. For this reason we provide a simple technique to identify technical and less technical terms according to definition 2 (see Table 1). According to this, technical terms have their origin in Latin or Greek language. Moreover, we know that those technical terms are very common in many (particular European) languages. Table 7 shows examples of similar expressions across various languages. Using this characteristic we propose the following method to identify medical technical expressions:

For each German target mention (G_t) we identify the English (E_j) and French (F_k) synonym with the lowest Levenshtein distance ($lev(a, b)$) for each of both languages. Next we calculate the average between both minimum distance scores. Note, we chose two languages rather than one to have a more robust distance score. Finally we harmonize this score, dividing it by the length of the target mention ($len(a)$). This should avoid that short strings are favoured over longer strings with similar edits. We refer to this score as the *harmonized distance* (h_dist). The harmonized distance can be formulated as follows:

$$h_dist_{G_t} = \frac{\min(lev(G_t, E_j)) + \min(lev(G_t, F_k))}{2 * len(G_t)} \quad (1)$$

Sorted Synonym data set (SSD): Following the assumption from above, we assume that a German mention with a low *harmonized distance* might likely to have a Greek or Latin origin, thus tends to be a technical term. Thus we calculate the *harmonized distance* of all German mentions of UMLS (and WUMLS) and sort all synonyms of each concept according to this score. Starting with the term with the lowest distance score and finishing with the one with the largest score.

distance (\geq)	0	5	10	15	20	25	30	35	40	45	50
#instance	300	237	193	161	144	124	97	87	74	56	49
%is-easier	50	59	65	71	74	74	75	74	70	70	71
%is-easier-or-equal	88	89	91	92	92	93	93	92	91	89	88

Table 8: Manual Evaluation of 300 selected examples to explore if the term ranked as easiest term is in fact easier than the term ranked as most technical. Considering only pairs with a larger edit distance, the results show that precision increases for both *is-easier* (checking whether the term is in fact simplified) or *is-easier-or-equal* (checking whether the term is at least not more complicated).

As we are interested in particular concepts we select only those which belong to one of these semantic types (STY): ‘Anatomical Abnormality’, ‘Anatomical Structure’, ‘Body Location or Region’, ‘Body Part, Organ, or Organ Component’, ‘Body Space or Junction’, ‘Disease or Syndrome’, ‘Injury or Poisoning’, ‘Mental or Behavioral Dysfunction’, ‘Sign or Symptom’. Using the technique from above and including English and French as reference language, we can generate sorted synonym sequences of 28,495 different concepts with overall 47,996 different mentions.

Evaluation 1 – Are synonyms with a low harmonic distance technical terms? In order to examine this question we randomly select 300 concepts and their lowest *h.dist* mention from UMLS-SSD. All selected mentions had a different harmonic score, whereas the largest score of the subset was 120. The selected mentions have been manually evaluated according to our two definitions by one of the authors. The analysis shows that 75% of all terms are technical expressions according to definition 1 and 90% according to definition 2. Table 9 shows an analysis considering only concepts below a certain harmonic distance threshold. In this way we can see that a harmonic distance below 60 leads to a high accuracy, which supports our assumption. The larger the distance the more the accuracy decreases. However the score decreases faster using definition 1.

distance (\leq)	20	40	60	80	100	120
#instances	59	105	174	277	297	299
%definition-1	93	93	91	79	75	75
%definition-2	98	99	99	94	90	90

Table 9: Manual examination of 300 randomly selected expressions of a concept with the lowest harmonic distance score.

Evaluation 2 – Are synonym mentions with a larger harmonic distance less technical and possibly laymen expressions? In order to examine this question we examine whether the term with the lowest score in UMLS-SSD is more or at least similarly technical as the term with the largest score of all synonyms. Thus, we selected randomly 300 German concept mention pairs, this time with the lowest and the largest *harmonic distance score* and examined whether the first term is a) more technical, b) similar technical or c) less technical than the second term. As we do not know whether there is always a simplified term within the synonym set, we evaluate according to *is-easier* ($a/(a + b + c)$), as well as *is-easier-or-equal*

$$((a + b)/(a + b + c)).$$

The results in Table 8 show that in only 50% of the cases the expression with the highest *harmonic distance* is less technical than the expression with the lowest *harmonic distance*. This does not look very promising at first. However we can make the following analyses: First considering all synonym pairs, in 88% of the cases the expression with the highest *harmonic distance* is easier or at least similarly technical as the expression with the lowest score. Moreover the table shows that the absolute distance between both scores has a strong influence on the outcome. Increasing the absolute distance between both scores quickly increases also the accuracy (%). In case of examining whether the expression with the higher score is in fact less technical, we can see a constant increase from 50%, using all pairs, to 75% considering a minimum absolute distance of 30. Increasing the distance, decreases obviously the number of synonym pairs. However, after reaching a maximum of 75%, the scores drop slightly, but never undergo 70. A similar effect can be observed for *is-easier-or-equal*. After a maximum of 93% with a distance of 30, the values slightly decrease but remain always above 88.

Overall these results are very promising. Considering a certain distance (e.g. of 15 or more), we can ensure that in more than 70% of the cases the synonym with the larger harmonic distance is less technical and in 92% of the cases the term is at least not more complicated.

5. Baseline Experiments

In the previous sections we presented the TLC corpus and in addition two further resources to support the mapping between German medical laymen to technical language and vice versa. The main focus of our work is the presentation of new resources in this domain. In this section, however, we present in addition some baseline results on TLC which can be used as benchmark for future work.

Regarding baseline results, we carry out two different experiments: 1) the normalization of medical technical terms including a term simplification and the 2) normalization of medical laymen expressions. For our experiment we indexed the mentions (and its stemmed version) from UMLS/WUMLS in Solr.

5.1. Experiment 1 – Normalization and Simplification

For experiment 1 we extract all technical terms and examine whether we can align it to a corresponding concept unique identifier. Using UMLS in 72.10% of the cases we can find the corresponding medical concept. However only

in 31.11% of those cases we find an easier synonym. The usage of WUMLS does not increase the performance much. However if we analyse the terms found in UMLS in more detail, we can see that the average harmonic distance score of those expressions is 39.93. As we know from Evaluation 1 in Section 4.4. that a low score is an indicator for a technical term, this score is no surprise. We can also see that a large number of expressions include a larger harmonic distance, for instance 143 expressions have a score of 70 or above.

5.2. Experiment 2 – Normalization Laymen Language

For experiment 2 we extract all laymen terms and examine whether the corresponding technical term can be found. In case of using UMLS terms for only 57.37% of the mentions a corresponding CUI can be detected. As laymen expressions provide much more variations in comparison to technical terms, this outcome was expected. If we again examine the expressions found in UMLS in more detail we can see that the average *harmonic distance* is at 82.05. However also here we can find a large number of expressions supposed to be non-technical, but have a low *harmonic distance*. For instance 137 expressions have a score below 170.

Finally, using WUMLS data for the normalization the score can be increased to 64.08%. This shows clearly the advantage of including additional information of Wiktionary.

5.3. Discussion

Overall the results of our baseline experiments show that laymen language concept normalization is much more difficult in comparison to the normalization of medical technical expressions. This highlights the importance of creating further resources of laymen synonyms but also methods being able to map between those language types.

Methods trained on definitions such as in Limsopatham and Collier (2016) might be helpful to tackle this challenge. However, in comparison to English UMLS and also Wiktionary do not contain as many German definitions as for English language. This again highlights the aspect that German, in comparison to English, is a low resourced language considering existing and freely available structured resources. As mentioned above, the German UMLS subset covers only 3.2% of all English concepts and involves only 2.3% of all existing English synonyms. Thus, it is obvious that concept normalization even for technical terms is much more challenging. Cross-lingual methods such as in Roller et al. (2018) might help to increase the coverage of technical terms.

6. Conclusion

In this work we presented a new corpus based upon a patient forum for kidney disease and stomach-intestines. The data set labels medical laymen language and technical terms and assigns a corresponding description or expression. This resource might be valuable resource to map and translate between both types of language styles in the medical domain. In addition to that we also provided two resources which can support this translation process. Finally

we also tested a simple baseline on our corpus which can be used as reference for more complex methods.

Acknowledgments

This project was funded by the European Union's Horizon 2020 research and innovation program under grant agreement No 780495 (BigMedilytics) and by the German Federal Ministry of Economics and Energy through the project MACSS (01MD16011F).

7. Bibliographical References

- Abbar, S., Mejova, Y., and Weber, I. (2015). You tweet what you eat: Studying food consumption through twitter. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 3197–3206, New York, NY, USA. ACM.
- Abrahamsson, E., Forni, T., Skeppstedt, M., and Kvist, M. (2014). Medical text simplification using synonym replacement: Adapting assessment of word difficulty to a compounding language. In *Proceedings of the 3rd Workshop on Predicting and Improving Text Readability for Target Reader Populations (PITR)*, pages 57–65, Gothenburg, Sweden, April. Association for Computational Linguistics.
- Chen, J., Jagannatha, A. N., Fodeh, S. J., and Yu, H. (2017). Ranking medical terms to support expansion of lay language resources for patient comprehension of electronic health record notes: adapted distant supervision approach. *JMIR medical informatics*, 5(4):e42.
- Doğan, R. I., Leaman, R., and Lu, Z. (2014). Ncbi disease corpus: a resource for disease name recognition and concept normalization. *Journal of biomedical informatics*, 47:1–10.
- Elhadad, N. and Sutaria, K. (2007). Mining a lexicon of technical terms and lay equivalents. In *Proceedings of the Workshop on BioNLP 2007: Biological, Translational, and Clinical Language Processing*, pages 49–56. Association for Computational Linguistics.
- Grabar, N. and Hamon, T. (2014). Automatic extraction of layman names for technical medical terms. In *2014 IEEE International Conference on Healthcare Informatics*, pages 310–319. IEEE.
- Karapetiantz, P., Audeh, B., Lillo-Le Louët, A., and Bousquet, C. (2018). Signal Detection for Baclofen in Web Forums: A Preliminary Study. In *MIE*, pages 421–425.
- Leaman, R., Islamaj Doğan, R., and Lu, Z. (2013). Dnorm: disease name normalization with pairwise learning to rank. *Bioinformatics*, 29(22):2909–2917.
- Limsopatham, N. and Collier, N. (2015). Adapting Phrase-based Machine Translation to Normalise Medical Terms in Social Media Messages. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1675–1680, Lisbon, Portugal, September. Association for Computational Linguistics.
- 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.
- O'Connor, K., Pimpalkhute, P., Nikfarjam, A., Ginn, R., Smith, K. L., and Gonzalez, G. (2014). Pharmacovigilance on twitter? Mining tweets for adverse drug reactions. In *AMIA annual symposium proceedings*, volume 2014, page 924. American Medical Informatics Association.
- Roller, R., Kittner, M., Weissenborn, D., and Leser, U. (2018). Cross-lingual Candidate Search for Biomedical Concept Normalization. In *Proceedings of Multilingual BIO*, Miyazaki, Japan, May.
- Stenetorp, P., Pyysalo, S., Topić, G., Ohta, T., Ananiadou, S., and Tsujii, J. (2012). brat: a Web-based Tool for NLP-Assisted Text Annotation. In *Proceedings of the Demonstrations Session at EACL 2012*, Avignon, France, April. Association for Computational Linguistics.
- Suominen, H., Salanterä, S., Velupillai, S., Chapman, W. W., Savova, G., Elhadad, N., Pradhan, S., South, B. R., Mowery, D. L., Jones, G. J., et al. (2013). Overview of the share/clef ehealth evaluation lab 2013. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 212–231. Springer.
- Weissenbacher, D., Sarker, A., Paul, M. J., and Gonzalez-Hernandez, G. (2018). Overview of the Third Social Media Mining for Health (SMM4H) Shared Tasks at EMNLP 2018. In *Proceedings of the 2018 EMNLP Workshop SMM4H: The 3rd Social Media Mining for Health Applications Workshop & Shared Task*, pages 13–16, Brussels, Belgium, October. Association for Computational Linguistics.
- Weissenbacher, D., Sarker, A., Magge, A., Daughton, A., O'Connor, K., Paul, M. J., and Gonzalez-Hernandez, G. (2019). Overview of the Fourth Social Media Mining for Health (SMM4H) Shared Tasks at ACL 2019. In *Proceedings of the Fourth Social Media Mining for Health Applications (#SMM4H) Workshop & Shared Task*, pages 21–30, Florence, Italy, August. Association for Computational Linguistics.
- Zeng-Treitler, Q., Goryachev, S., Kim, H., Keselman, A., and Rosendale, D. (2007). Making texts in electronic health records comprehensible to consumers: a prototype translator. In *AMIA Annual Symposium Proceedings*, volume 2007, page 846. American Medical Informatics Association.