Bioinformatics, YYYY, 0–0
doi: 10.1093/bioinformatics/xxxxx
Advance Access Publication Date: DD Month YYYY
Manuscript Category

Subject Section

Lifestyle factors in the biomedical literature: comprehensive resources for named entity recognition

Esmaeil Nourani[1,2], Mikaela Koutrouli[1], Yijia Xie[1,#], Danai Vagiaki[1,$], Sampo Pyysalo[3], Katerina Nastou[1,*], Søren Brunak[1,*] and Lars Juhl Jensen[1,*]

[1]Novo Nordisk Foundation Center for Protein Research, Faculty of Health and Medical Sciences, University of Copenhagen

[2]Faculty of Information Technology and Computer Engineering, Azarbaijan Shahid Madani University, Tabriz, Iran

[3]TurkuNLP group, Department of Computing, Faculty of Technology, University of Turku, Turku, Finland

[#]Current address: Zhejiang Ponshine Information Technology Co., Ltd., Hangzhou 311100, China, [$]Current address: a. Genome Biology Unit, European Molecular Biology Laboratory (EMBL), Heidelberg, Germany, b. Collaboration for joint PhD degree between EMBL and Heidelberg University, Faculty of Biosciences, Heidelberg, Germany, c. Division of Computational Genomics and Systems Genetics, German Cancer Research Center (DKFZ), Heidelberg, Germany

*To whom correspondence should be addressed.

Associate Editor: XXXXXXX

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Abstract

Motivation: Despite lifestyle factors (LSFs) being increasingly acknowledged in shaping individual health trajectories, particularly in chronic diseases, they have still not been systematically described in the biomedical literature. This is in part because no named entity recognition (NER) system exists, which can comprehensively detect all types of LSFs in text. The task is challenging due to their inherent diversity, lack of a comprehensive LSF classification for dictionary-based NER, and lack of a corpus for deep learning-based NER.

Results: We present a novel hierarchical Lifestyle Factor Classification (LSFC), which we used to develop a dictionary-based system for recognition and normalization of LSFs. Additionally, we introduce a manually annotated corpus for LSFs (LSF200) suitable for training and evaluation of NER systems, and use it to train a transformer-based system. Evaluating the performance of both NER systems on the corpus revealed an F-score of 64% for the dictionary-based system and 76% for the transformer-based system. Large-scale application of these systems on PubMed abstracts and PMC Open Access articles identified over 300 million mentions of LSF in the biomedical literature.

Availability: LSFC, the annotated LSF200 corpus, and the detected LSFs in PubMed and PMC-OA articles using both NER systems, are available under open licenses via the following GitHub repository: https://github.com/EsmaeilNourani/LSFC-expansion. This repository contains links to two associated GitHub repositories and a Zenodo project related to the study.

Contact: katerina.nastou@cpr.ku.dk, soren.brunak@cpr.ku.dk, lars.juhl.jensen@cpr.ku.dk

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 Introduction

Lifestyle is a complex and multifaceted concept that involves nutrition, behaviors, habits, and activities that individuals engage in, which can impact their health. These lifestyle factors (LSFs) are widely recognized as important in shaping individual health trajectories and influencing the onset and progression of diseases (WHO, 2023). Numerous studies illustrate the association between an unhealthy lifestyle and an elevated risk of chronic disease as well as the potential of reducing the risk by adopting a healthy lifestyle (Subramanian et al., 2020; Nyberg et al., 2020; Tobias et al., 2023). LSFs can thus play a pivotal role together with genetic factors in precision medicine (Yurkovich et al., 2023; Gabbert et al., 2023; Gray et al., 2020; Jeon et al., 2018), and a crucial first step is to structure the existing information on LSFs.

Specific types of LSFs such as environments, exposure to risk factors, smoking, and food are covered by existing resources like the Environment Ontology (Buttigieg et al., 2013), the Exposome Explorer (Neveu et al., 2020), the Cigarette Smoke Exposure Ontology (Younesi et al., 2014), and FoodOn (Dooley et al., 2018). However, these fall short in comprehensively covering other crucial disease-associated LSFs, namely physical and leisure time activities, socioeconomic factors, personal care products and cosmetic procedures, sleep, mental health practices, and substance use. This severely hampers the development of text-mining solutions to help structure the LSF information in the literature.

The first critical step is to develop a named entity recognition (NER) method to identify LSF mentions within scientific publications. Many such systems exist for other biomedical entities, reviewed in (Song et al., 2021; Huang et al., 2020; Perera et al., 2020). Developing an NER method for LSFs requires either a comprehensive LSF dictionary that can be matched against text to find them, or text corpus with manually annotated LSFs that can be used to train a deep learning-based method.

In this paper, we address the task of recognizing LSFs within biomedical text by presenting four major contributions. Firstly, we introduce a novel LSF classification (LSFC), featuring a multilevel hierarchical structure. It includes the main LSF categories at the top level and extends to specific subcategories and low-level concepts. Secondly, we introduce the first annotated corpus for LSFs, LSF200, a valuable resource for the BioNLP community as it can be used for training and evaluating NER systems. Thirdly, we introduce a dictionary-based NER system enabling for the first time the detection of a diverse set of lifestyles in the literature. This NER system leverages the LSFC for dictionary creation, thus enabling both recognition and normalization for the matches to LSFC concepts. Lastly, we introduce a transformer-based NER system (Vaswani et al., 2017) for LSF detection, leveraging LSF200 for training and evaluation.

2 Materials and Methods

Currently, there is a lack of resources to capture the vast diversity of LSFs, under a single umbrella. Below, we have made an effort to generate a comprehensive categorization of the different aspects of LSFs accompanied by a small description of concepts that fit in each category. The creation of this categorization is the first crucial step that allowed us to annotate a text corpus and create a dictionary of LSFs, for the purposes of deep learning-based NER and dictionary-based NER, respectively. In the next sections we provide details in the methodology used to achieve our goals.

2.1 The LSF categorization

Within the context of lifestyle, we have identified nine categories that can collectively describe all LSFs, namely:

(1) **Nutrition**, a category that covers different branches such as dietary habits, food groups, food processing and preparation, macronutrients, and micronutrients, among others.

(2) **Socioeconomic factors**, which includes social and economic conditions such as income, wealth, education, and socioeconomic status.

(3) **Environmental exposures**, includes exposure to various environmental factors, such as air pollution, water quality, and workplace hazards, that can impact an individual's health.

(4) **Substance use**, covers concepts such as smoking, as well as illicit drug use.

(5) **Physical activities**, includes regular exercise and physically demanding activities like leisure time, occupational, and household physical activities.

(6) **Non-physical leisure time activities**, describes any activity an individual might engage in during their free time that is not a physical activity.

(7) **Personal care products and cosmetic procedures**, covers activities related to hygiene, use of cosmetic and cleaning products as well as invasive procedures that people undergo to improve their appearance, such as cosmetic surgery.

(8) **Sleep**, covers sleep quality, stages, and habits.

(9) **Mental health practices**, includes the behaviors and habits related to maintaining good mental health and emotional well-being such as meditation, and psychotherapy.

2.2 Annotation of the LSF200 text corpus

To create an LSF text corpus, we selected three of the most relevant journals within each of the nine LSF categories introduced above. We selected 200 abstracts for our corpus, evenly distributed across all categories. We aimed for a balanced selection of documents among LSF categories, considering the chosen journals' scope and focus. Selected top three journals per category are available in Supplementary Table 2. To ensure high diversity, we opted to annotate abstracts instead of full-text documents, as this would allow us to annotate a larger number of documents and thus obtain a wider selection of different entities with the same curation effort.

The annotation process started by creating an initial set of annotation guidelines, which we improved through two rounds of refinement. For each of the three versions of the guidelines, two curators annotated a new set of 15 abstracts, based on which we calculated the inter-annotator agreement. We subsequently annotated the entire LSF200 according to the final guidelines (https://esmaeilnourani.github.io/lifestylefactors-annotation-docs/entities). We used the BRAT rapid annotation tool for document annotation (Stenetorp et al., 2012).

2.3 Dictionary construction

One common approach to NER is to develop a dictionary-based system (Cook and Jensen, 2019), leveraging existing databases or ontologies to create the dictionary. For instance, a disease NER system can be established by constructing a dictionary derived from existing resources like the Disease Ontology (Baron et al., 2024). Since such resources were not

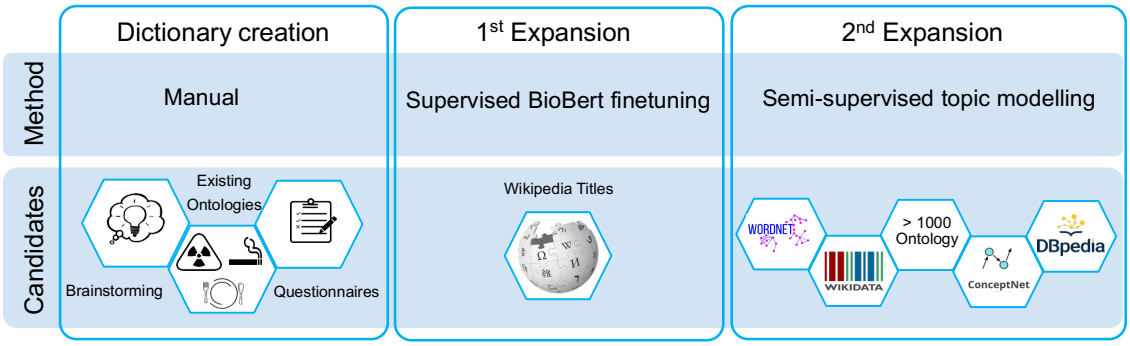

Fig. 1. LSF Dictionary creation and expansion stage

available for the purposes of this work, we manually constructed an initial dictionary inspired by existing biomedical literature, ontologies, and questionnaires. The aim was to come up with a wide selection of names belonging to each of the nine categories of LSFs described above, which would serve as a good starting point for semi-automatic expansion.

First, we wanted to introduce names from a resource that would contain a diverse set of candidates. As 45% of the names in the initial dictionary exist as Wikipedia pages, we decided to use Wikipedia page titles as a source of potentially missing LSFs. To predict which page titles are good LSF candidates, we used a transformer-based approach that scores names based on their contexts in biomedical literature (Nastou et al., 2023), and manually assessed all high-scoring candidates before adding them to the dictionary. For more details on the training, evaluation and prediction process please refer to Supplementary Section 1.

Second, we incorporated names from resources such as WordNet (Brown, 2005), Wikidata, DBpedia (Lehmann et al., 2015), ConceptNet (Speer et al., 2018), and 1085 existing ontologies registered in BioPortal (Whetzel et al., 2011). LSF-related candidates were extracted from these resources, and scored based on both semantic similarity to existing names and textual context using BERTopic (Grootendorst, 2022). Figure 1 displays the details of different stages of LSF dictionary creation and expansion, along with methods used and the resources considered in each stage as LSF candidates.

2.4 Creation of an LSF classification and referencing to external resources

In this study, we defined nine distinct categories, encompassing diverse LSF aspects, to establish a unified LSF classification across the various categories. For practical uses of the dictionary, such as concept indexing, it is important to know which names are synonyms for the same LSF and how the LSFs relate to each other. We thus created an LSF classification (LSFC) in which each LSF is a concept with a unique identifier associated with relevant synonyms, and the concept can be traced back to the root of the LSFC through *is_a* relationships. We selected the Open Biomedical Ontologies (OBO) format to enhance linking and integration within OBO-based workflows and ontologies. The files are hosted on Github/Zenodo (https://github.com/EsmaeilNourani/Lifestyle-factors-classification), allowing users to easily fork them for specific use-cases or contribute extensions.

We utilized the BioPortal annotator to match the existing LSFs in our LSFC with terms from over 1000 ontologies. To add cross-references (Xrefs) only to relevant ontologies, we relied on two metrics: *Coverage*, which determined the percentage of LSFs that exist in the target ontology, and *Overlap*, which assessed the percentage of terms from the target ontology found within our LSFC. Based on these criteria, we manually narrowed down our selection to 50 ontologies to which we added Xrefs. This enhances interoperability and enables integration of existing ontologies, for example, by importing all child fine-grained concepts under the matched target term from domain-specific ontologies, furthering the development of a more comprehensive and interconnected LSFC.

2.5 Dictionary-based NER

The JensenLab tagger (tagger hereafter) (https://github.com/larsjuhljensen/tagger) is a fast dictionary-based NER system that recognizes a wide variety of biomedical entities based on underlying dictionaries (Jensen, 2016). To provide a dictionary-based NER system for LSF recognition, we integrated the tagger into our workflow, enhancing its functionality by supplying a dedicated LSF dictionary. The LSF dictionary was constructed using names from LSFC, alongside orthographic variant generation and a block list to enhance recall and precision, respectively. The tagger assigns a single unique identifier to synonyms and automatically generates name variations (e.g. plural and adjective forms) of the same LSF, enabling effective NEN for matched names. To improve precision, we use a block list to exclude problematic names that would cause many false positives during tagging. Specifically, we manually inspected all names that gave more than 2000 matches in 36.1 million PubMed abstracts (as of August 2023) and 4.5 million articles from the PMC open access subset (as of April 2022) to identify those that should be added to the block list.

2.6 Transformer-based NER

To explore the potential of using transformers for NER, we adapted an existing NER system (Luoma et al., 2023). Specifically, we built upon the RoBERTa-large-PM-M3-Voc model (RoBERTa-bio hereafter), which has demonstrated the best performance in several NER tasks (Lewis et al., 2020; Miranda-Escalada et al., 2023). We trained the model for multi-class classification of the nine categories of LSFs using LSF200 without OOC annotations. The 200 abstracts were divided into a 40-document holdout test set and a 160-document combined train and development set. We made sure this split was balanced across the nine categories of LSFC.

Hyperparameter selection was done through a grid search using a 5-fold cross-validation setup. This was done instead of a simple train–development split to avoid the risk of overfitting on a small development set. To determine the best hyperparameters for the model's final evaluation, we computed the micro F1-score based on the total True Positives (TP), False Positives (FP), and False Negatives (FN) across all classes and folds. Finally, we trained a model on all available training and development data with the selected optimal hyperparameters and evaluated it on the holdout test set.

3 Results and Discussion

3.1 The LSF categorization

Being able to identify and categorize concepts as diverse as LSFs was the main challenge of this work and it was achieved through the creation of a comprehensive and inclusive LSF categorization (LSFC). Table 1 presents an overview of LSFC, with the last column displaying the number of Xrefs per LSF category. One of the most important and well-studied categories of LSFs is *Nutrition*. To get more fine-grained concepts for this category, we incorporated food groups from the widely used FoodOn ontology. Table 1 presents both statistics with and without expansions from FoodOn.

Table 1. LSFC statistics. In parentheses, the numbers including expansion with terms from FoodOn are shown for the *Nutrition* category and the total number of entries in LSFC.

Category	Unique LSFs	Total names	Xrefs
Nutrition (expanded with FoodOn)	766 (4710)	987 (5126)	2093 (4875)
Socioeconomic factors	302	604	994
Environmental exposures	175	694	643
Substance use	104	283	596
Physical activities	84	171	283
Personal care products and cosmetic procedures	78	139	361
Non-physical leisure time activities	58	124	68
Sleep	45	104	129
Mental health practices	40	78	135
Total (expanded with FoodOn)	1652 (5597)	3184 (7324)	5302 (8097)

3.2 The LSF200 Corpus

LSF200 comprises 200 abstracts, with a total of 39,416 tokens based on BERT basic tokenization (https://github.com/spyysalo/bert-vocab-eval). It contains 1876 manually annotated mentions of LSFs, which are broken down into categories in Table 2. The majority of the LSF mentions fall into the categories *Nutrition, Socioeconomic factors, and Physical activities.* In contrast, the categories *Non-physical leisure time activities, Mental health practices, and Personal care products and cosmetic procedures* are less prevalent in LSF200. The distribution of names is consistent with the statistics of names in the different LSF categories of LSFC (Table 1), showcasing that despite selecting journals to represent all nine categories of LSFs, names from the largest categories tend to appear in most abstracts. Evaluating the quality of the manual annotations in terms of inter-annotator agreement gave an F1-score of 83%. This score emphasizes the challenging nature of annotating concepts as diverse as LSFs, making it difficult to formulate a comprehensive definition and achieve perfect annotation even at a human level.

We also introduced a special category of annotations called Out of context (OOC), which we use to highlight mentions that appear in a context that falls outside the scope of LSFs. For example, the word "tobacco" is normally an LSF but not when it appears in the context of "tobacco company".

Table 2. LSF200 corpus statistics

Category	Total mentions	Unique names	% of names per category
Nutrition (expanded with FoodOn)	445	270	27.36%
Socioeconomic factors	275	174	17.63%
Environmental exposures	199	132	13.37%
Substance use	187	75	7.60%
Physical activities	222	88	8.92%
Personal care products and cosmetic procedures	56	35	3.55%
Non-physical leisure time activities	61	27	2.74%
Sleep	110	57	5.78%
Mental health practices	74	29	2.94%
OOC	247	100	10.13%
Total	1876	987	100%

3.3 Evaluation of dictionary-based NER

To assess the performance of our dictionary-based NER system, a crucial component involves evaluating the dictionary generated from the LSFC. It is important to note that this dictionary is not constructed using the text corpus, and, as such, it is appropriate to employ the entire LSF200 for evaluation of the dictionary and NER system.

We present performance results in two versions: one with OOC mentions included in the annotated corpus and one without. This evaluation assesses the impact of removing OOC mentions on the performance of dictionary-based NER. In the initial version, OOCs are treated as typical LSF mentions, resulting in an F1-score of 65.2% (precision: 96.0%, recall: 49.4%). In the second version, we removed OOC mentions entirely from annotations. The resulting performance was a slightly lower F1-score of 63.6% (precision: 85.0%, recall: 50.8%). When OOCs are not treated as LSFs, the number of false positives increases considerably, due to the presence of the OOC names in the dictionary, which results in a large drop in precision from 96.0% to 85.0%. Conversely, how OOC mentions are counted has very little impact on recall.

Figure 2 shows the performance for each of the nine categories, revealing that the recall varies much more than the precision and that OOC has a similar impact on precision across categories. For most categories, OOC mentions have only a minor impact on recall, with "Non-physical time activities" being the only notable exception. For this category, many OOC mentions go undetected by the dictionary-based NER system, and the recall thus improves substantially when OOC mentions are excluded.

The wide range of recall values is explained by fundamental differences between the categories. For example, the dictionary by design does not have concepts for all possible exposures, for which reason the *Environmental exposures* category has very low recall. This implies that the low recall can likely be improved by using the existing Xrefs to integrate fine-grained concepts from, e.g., the Exposome database.

In LSF200, this issue was not immediately evident, as it is selected from journals where all these entities are represented as LSFs. However, in a real-world scenario where the entire literature is tagged, this approach could result in detrimental effects on precision. Detailed performance results for various categories are available in Supplementary Table 3.

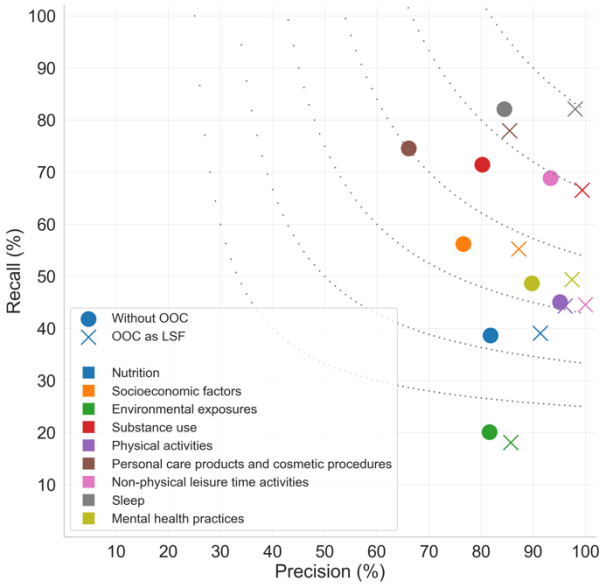

Fig. 2. Performance of dictionary-based NER system in two OOC annotation variants per LSF category.

3.3.1 Manual error analysis

As anticipated, the dictionary-based NER exhibits impressive precision but lower recall due to its limitations in recognizing out-of-dictionary mentions. In Table 3, we have categorized the missing LSF mentions (false negatives) by the dictionary-based NER system.

The majority of missing LSF mentions are due to either abbreviations or synonyms related to existing LSF concepts in the dictionary. This is promising, as it suggests LSFC covers most of the relevant concepts, and that the recall of the NER system can thus be improved by enriching our dictionary with more synonyms for existing concepts. Other false negatives are due to ambiguous names that have been intentionally excluded; for example, "therapy" can denote both psychotherapy, which falls under mental health practices, or medical treatments, which are not considered LSFs. This approach prioritizes precision by avoiding matches with too many irrelevant hits, although it results in missing some relevant mentions in LSF200.

Table 3. Error analysis for dictionary-based NER (False Negatives)

Error category	Total mentions	Unique names
Abbreviation & Synonym	332	155
Ambiguous name	213	97
Fine-grained concept	211	128
Discontinuous entity	20	14
Annotation error	17	15
Total	793	409

The last major group of missing LSF mentions is "Fine-grained LSF concepts", which refers to mentions that cannot be classified as major missing LSF concepts because the parent LSFs of these mentions already exist in the dictionary. The *Environmental exposures* is, as already mentioned, a prominent example of this. The final few false negatives are either due to discontinuous entities, which cannot be identified by a dictionary-based method (e.g. "Vitamin C" in the expression "Vitamins E and C") or simply annotation errors made by the human curators and are thus not actual errors of the NER system.

Table 4. Error analysis for dictionary-based NER (False Positives)

Error category	Total mentions	Unique names
OOC	120	31
Ambiguous name	13	4
Annotation error	11	7
Dictionary error	3	2
Total	147	44

As the dictionary-based system has high precision, it produces much fewer false positives than false negatives, the majority of which are due to OOC mentions as already described (Table 4). The remaining few false positives are either due to dictionary errors (names that should not have been in the dictionary), ambiguous names (which could be blocked at the price of more false negatives), or annotation errors in the corpus.

3.4 Comparison with transformer-based NER

To compare the performance of dictionary-based NER to state-of-the-art methods, we trained a transformer-based system on LSF200 without OOC mentions (see Methods for details). We evaluated the system on the 40-document test set of LSF200. The overall NER performance was 70.1% precision and 85.3% recall, corresponding to an F1-score of 77.0%. To allow direct comparison, we also reevaluated the dictionary-based NER system on the test set only, yielding 76.6% precision, 53.8% recall, and 63.2% F1-score. The results clearly highlight that the dictionary-based NER maximizes the precision by exclusively tagging predefined names from the dictionary and utilizing a block list to avoid tagging problematic names. In contrast, the transformer-based NER achieves better recall due to its ability to detect LSF mentions based on the context.

Figure 3 shows the performance of dictionary-based NER and transformer-based NER for the test data within each LSF category.

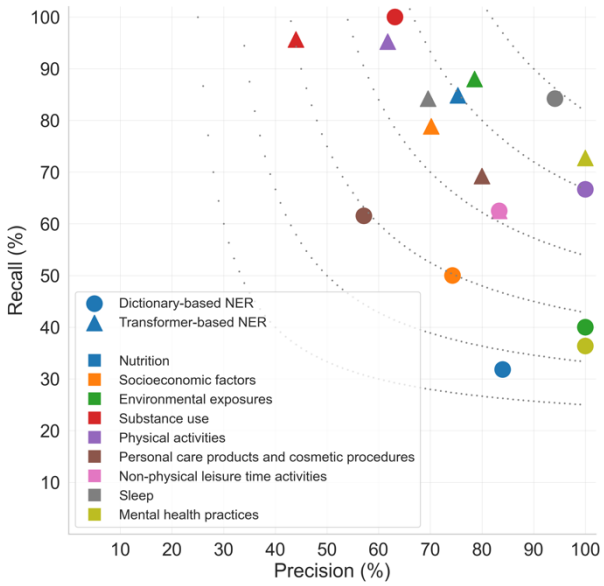

Fig. 3. Performance of dictionary-based NER and Transformer-based NER only for test data per LSF category.

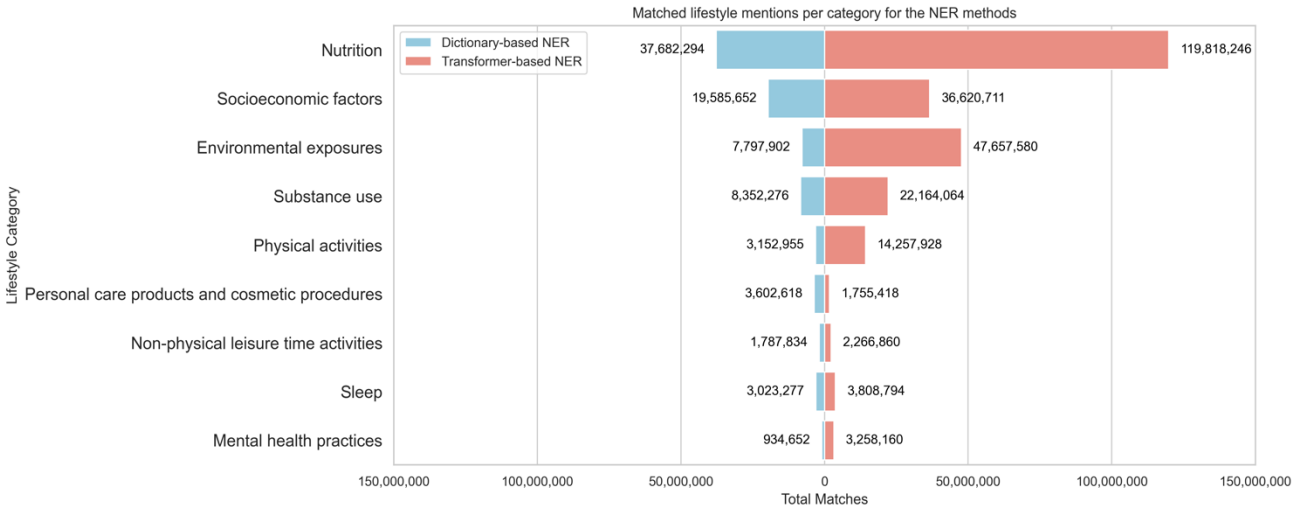

Fig. 4. Matched lifestyle mentions from large-scale tagging of scientific literature using the NER methods

The transformer-based NER demonstrates promising recall across most categories, with *Nutrition*, *Mental health practices*, and *Socioeconomic factors* showing particularly large improvements in recall with only small drops in precision. For *Personal care products and cosmetic procedures*, transformer-based NER surprisingly improved mainly the precision. Finally, the transformer-based NER system performs worse than the dictionary-based one in terms of both precision and recall for categories *Substance use* and *Sleep*.

3.4.1 Manual error analysis

In Table 5 we categorize the errors produced by the Transformer-based NER, encompassing both names missed by the system (FNs) and names incorrectly detected as LSFs (FPs).

Table 5. Error analysis for Transformer-based NER

	FP		FN	
Error category	Total mentions	Unique names	Total mentions	Unique names
Model error	40	23	31	25
OOC	25	6	0	0
Discontinuous entity	4	2	2	2
Annotation error	20	13	4	4
Total	89	44	37	31

Even for the transformer-based system, OOC mentions are frequently wrongly predicted as LSFs, showing that some OOC mentions are hard to distinguish, even for a method that looks at the context. The system also makes a few mistakes on discontinuous entities, although much fewer than the dictionary-based system. The majority of the errors, however, were labeled as "Model error", encompassing the broad class of cases where the model, for no obvious specific reason, fails to provide accurate predictions. Lastly, we have again some errors, which upon inspection turn out to be mistakes in the manual annotation rather than the model being wrong.

3.5 Large-scale tagging of the scientific literature

Results for the tagging of 36.1 million PubMed abstracts (as of August 2023) and 4.5 million articles from the PMC open access subset (as of April 2022) using both the dictionary-based NER and the transformer-based NER have been made available via Zenodo (https://zenodo.org/records/10450308). There are in total 85,919,460 LSF matches for dictionary-based NER, with approximately half of them (44%) corresponding to *Nutrition* terms and a quarter (23%) to *Socioeconomic factors*. Tagging with the transformer-based system yielded 251,607,761 total matches. *Nutrition* terms once again constitute almost half of the matches of the system (48%), and *Environmental exposures* come second at 19%. Figure 4 shows the matches per LSF category for the two methods. A key difference between the results of the two systems is that the dictionary-based system inherently provides matches that are normalized to LSFC identifiers whereas the transformer-based system does not. This makes the former a better starting point for many other text-mining tasks such as relation extraction.

4 Conclusions

In this paper, we introduce novel resources to address the challenge of recognizing LSFs within biomedical text. We present dictionary-based NER and transformer-based NER systems, both demonstrating promising performance in identifying LSFs. The LSF Classification stands out as a diverse and hierarchical classification of LSFs that we used as a backbone for the dictionary-based NER system, but which can serve as a resource for standardizing LSF information in general. Furthermore, the manually annotated LSF200 corpus proved to be sufficient for training a transformer-based NER and for evaluating both types of NER systems. The presented NER systems, LSFC, LSF200 corpus, and matched LSFs from large-scale runs of both NER systems on PubMed and PMC-OA articles, are made publicly available under open licenses to facilitate further research.

Supplementary data

Supplementary data are available at Bioinformatics online.

Funding

This work was supported by the Novo Nordisk Foundation [NNF14CC0001, NFF17OC0027594]. K.N. has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie [101023676]. M.K. has received funding from Novo Nordisk Foundation [NNF20SA0035590].

Conflict of Interest: none declared.

References

Baron,J.A. et al. (2024) The DO-KB Knowledgebase: a 20-year journey developing the disease open science ecosystem. Nucleic Acids Res, 52, D1305–D1314.

Brown,K. (2005) Encyclopedia of Language and Linguistics Elsevier.

Buttigieg,P.L. et al. (2013) The environment ontology: contextualising biological and biomedical entities. Journal of Biomedical Semantics, 4, 43.

Cook,H.V. and Jensen,L.J. (2019) A Guide to Dictionary-Based Text Mining. In, Larson,R.S. and Oprea,T.I. (eds), Bioinformatics and Drug Discovery, Methods in Molecular Biology. Springer, New York, NY, pp. 73–89.

Dooley,D.M. et al. (2018) FoodOn: a harmonized food ontology to increase global food traceability, quality control and data integration. npj Sci Food, 2, 23.

Gabbert,C. et al. (2023) Lifestyle factors and clinical severity of Parkinson's disease. Sci Rep, 13, 9537.

Gray,I.D. et al. (2020) Precision Medicine in Lifestyle Medicine: The Way of the Future? Am J Lifestyle Med, 14, 169–186.

Grootendorst,M. (2022) BERTopic: Neural topic modeling with a class-based TF-IDF procedure.

Huang,M.-S. et al. (2020) Biomedical named entity recognition and linking datasets: survey and our recent development. Briefings in Bioinformatics, 21, 2219–2238.

Jensen,L.J. (2016) One tagger, many uses: Illustrating the power of ontologies in dictionary-based named entity recognition. 067132.

Jeon,J. et al. (2018) Determining Risk of Colorectal Cancer and Starting Age of Screening Based on Lifestyle, Environmental, and Genetic Factors. Gastroenterology, 154, 2152-2164.e19.

Lehmann,J. et al. (2015) DBpedia – A large-scale, multilingual knowledge base extracted from Wikipedia. Semantic Web, 6, 167–195.

Lewis,P. et al. (2020) Pretrained Language Models for Biomedical and Clinical Tasks: Understanding and Extending the State-of-the-Art. In, Rumshisky,A. et al. (eds), Proceedings of the 3rd Clinical Natural Language Processing Workshop. Association for Computational Linguistics, Online, pp. 146–157.

Luoma,J. et al. (2023) S1000: a better taxonomic name corpus for biomedical information extraction. Bioinformatics, 39, btad369.

Miranda-Escalada,A. et al. (2023) Overview of DrugProt task at BioCreative VII: data and methods for large-scale text mining and knowledge graph generation of heterogenous chemical–protein relations. Database, 2023, baad080.

Nastou,K. et al. (2023) Improving dictionary-based named entity recognition with deep learning. 2023.12.10.570777.

Neveu,V. et al. (2020) Exposome-Explorer 2.0: an update incorporating candidate dietary biomarkers and dietary associations with cancer risk. Nucleic Acids Research, 48, D908–D912.

Nyberg,S.T. et al. (2020) Association of Healthy Lifestyle With Years Lived Without Major Chronic Diseases. JAMA Internal Medicine, 180, 760–768.

Perera,N. et al. (2020) Named Entity Recognition and Relation Detection for Biomedical Information Extraction. Front Cell Dev Biol, 8, 673.

Song,B. et al. (2021) Deep learning methods for biomedical named entity recognition: a survey and qualitative comparison. Briefings in Bioinformatics, 22, bbab282.

Speer,R. et al. (2018) ConceptNet 5.5: An Open Multilingual Graph of General Knowledge.

Stenetorp,P. et al. (2012) brat: a Web-based Tool for NLP-Assisted Text Annotation. In, Segond,F. (ed), Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics. Association for Computational Linguistics, Avignon, France, pp. 102–107.

Subramanian,M. et al. (2020) Precision medicine in the era of artificial intelligence: implications in chronic disease management. Journal of Translational Medicine, 18, 472.

Tobias,D.K. et al. (2023) Second international consensus report on gaps and opportunities for the clinical translation of precision diabetes medicine. Nat Med, 29, 2438–2457.

Vaswani,A. et al. (2017) Attention Is All You Need.

Whetzel,P.L. et al. (2011) BioPortal: enhanced functionality via new Web services from the National Center for Biomedical Ontology to access and use ontologies in software applications. Nucleic Acids Res, 39, W541-545.

WHO (2023) Non communicable diseases [Fact sheet].

Younesi,E. et al. (2014) CSEO – the Cigarette Smoke Exposure Ontology. J Biomed Semantics, 5, 31.

Yurkovich,J.T. et al. (2023) The transition from genomics to phenomics in personalized population health. Nat Rev Genet, 1–17.