

Large-scale investigation of weakly-supervised deep learning for the fine-grained semantic indexing of biomedical literature

Anastasios Nentidis^{a,b,*}, Thomas Chatzopoulos^{a,c}, Anastasia Krithara^a,
Grigorios Tsoumakas^b, Georgios Paliouras^a

^a*Institute of Informatics and Telecommunications, NCSR Demokritos, Athens, Greece*

^b*School of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece*

^c*Department of Computer Engineering and Informatics, University of Patras, Patras, Greece*

Abstract

Objective: Semantic indexing of biomedical literature is usually done at the level of MeSH descriptors with several related but distinct biomedical concepts often grouped together and treated as a single topic. This study proposes a new method for the automated refinement of subject annotations at the level of MeSH concepts. **Methods:** Lacking labelled data, we rely on weak supervision based on concept occurrence in the abstract of an article, which is also enhanced by dictionary-based heuristics. In addition, we investigate deep learning approaches, making design choices to tackle the particular challenges of this task. The new method is evaluated on a large-scale retrospective scenario, based on concepts that have been promoted to descriptors. **Results:** In our experiments concept occurrence was the strongest heuristic achieving a macro-F1 score of about 0.63 across several labels. The proposed method improved it further by more than 4pp. **Conclusion:** The results suggest that concept occurrence is a strong heuristic for refining the coarse-grained labels at the level of MeSH concepts and the proposed method improves it further.

*Corresponding author

Email address: nentidis@csd.auth.gr, tasosnent@iit.demokritos.gr (Anastasios Nentidis)

¹This is a preprint of an article published in the Journal of Biomedical Informatics. The final authenticated version is available online at: <https://doi.org/10.1016/j.jbi.2023.104499>

Keywords: Semantic Indexing, Medical Subject Headings (MeSH), Biomedical Literature, Weak Supervision, Deep Learning

2010 MSC: 00-01, 99-00

1. Introduction

Biomedical literature is a rich source of biomedical knowledge that is constantly expanding. Identification of relevant articles for specific topics in such a rich collection is a real challenge. In this direction, the National Library
5 of Medicine (NLM) indexes citations in the MEDLINE database with topic descriptors from the Medical Subject Headings (MeSH) thesaurus². This semantic indexing process allows MEDLINE/PubMed³ to offer advanced semantic search strategies, that can retrieve citations relevant to a topic of interest, addressing issues such as synonymy and polysemy of biomedical concepts.

10 MeSH provides more than thirty thousand descriptors, such as diseases and chemicals, hierarchically organized so that most are subcategories of at least one broader descriptor. For example, “Epilepsy, Reflex” is a subcategory of the broader descriptor “Epilepsy” as shown in Figure 1. Some of these descriptors cover several related but distinct subordinate *MeSH concepts*⁴, which are defined as sets of synonymous terms. These concepts are usually narrower⁵ than
15 the main topic of the descriptor. For example, “Epilepsy, Reflex” contains a *preferred concept*, providing the name of the descriptor and the synonymous term “Reflex Epilepsy”, but includes also additional terms for narrower concepts such as “Reflex Epilepsy, Audiogenic” and “Tactile Reflex Epilepsy”, that
20 are not MeSH descriptors themselves.

Such narrower concepts are not distinguishable by the semantic indexing process, which is done at the level of descriptors. Therefore, semantic search

²<https://meshb.nlm.nih.gov/>

³<https://pubmed.ncbi.nlm.nih.gov/>

⁴https://www.nlm.nih.gov/mesh/concept_structure.html

⁵In some cases subordinate concepts can be broader than the main topic.

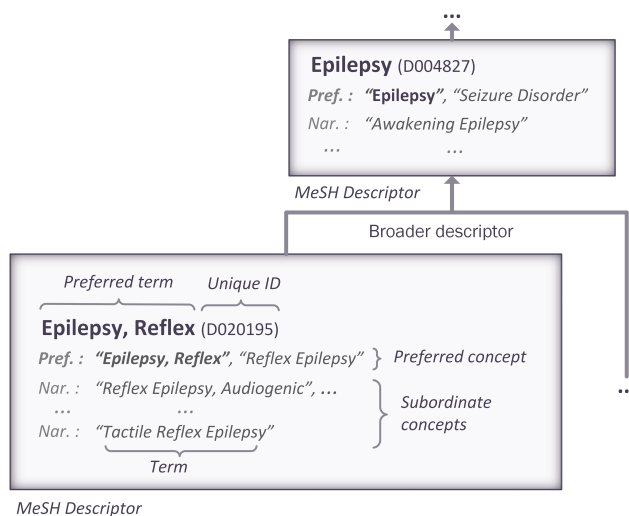


Figure 1: The structure and hierarchical position of the MeSH descriptor “Epilepsy, Reflex”.

cannot be used at this level of detail. The goal of this work is to enable such fine-grained semantic search, providing a method for the Fine-Grained Semantic
 25 Indexing (*FGSI*) of the biomedical literature, that is the automated refinement of MeSH labels from the descriptor level to the level of each narrower concept. For example, we would like to distinguish which of the articles already annotated with the “Epilepsy, Reflex” are for “Audiogenic Epilepsy” and which for “Tactile Reflex Epilepsy” or other subcategories of “Epilepsy, Reflex”.

30 Indexing at the level of MeSH concepts is beyond the current practice in MEDLINE/PubMed and, if enforced, it would almost double the number of distinct topics used for indexing. Doing this additional work manually is not practical, given the growing volumes of biomedical literature. Furthermore, state-of-the-art systems that provide automated topic suggestions focus on known MeSH
 35 topics and usually require several labelled training examples for each topic [1]. Therefore, obtaining reliable topic suggestions for new fine-grained topics at the level of concepts is an open research issue. The lack of labelled data for the development and evaluation of methods is a major challenge of this *FGSI* task.

In this work, we propose *Deep Beyond MeSH (DBM)*, a new method for

40 the refinement of coarse-grained MeSH annotations based on deep pre-trained
language models, namely *PubMedBERT* [2]. To overcome the lack of labelled
data for supervised training, *DBM* relies on weak supervision, investigating
whether such Deep Learning (DL) models can be resilient to heuristic labels.
For weak supervision, the method is based on *concept occurrence (CO)*⁶, which
45 is defined as the information extraction task of recognizing a biomedical named
entity in a text and then mapping it to a specific concept from a normalized
semantic system, such as the *UMLS*⁷.

Previous experiments with this idea [4], using small ground-truth datasets
for refining two specific MeSH descriptors, suggested that *CO* can provide a
50 strong basis for the *FGSI* task. However, this heuristic needed to be tested
on larger datasets of ground-truth data that cover several MeSH descriptors.
In this direction, we devised a retrospective scenario, based on the evolution of
MeSH indexing in MEDLINE, and introduce here *Retrospective Beyond MeSH*
(*RetroBM*). This new method derives large-scale datasets for several individual
55 concepts that got promoted over time to new fine-grained MeSH descriptors,
that is descriptors covering a single MeSH concept.

In brief, the main questions driving this work are the following:

- Is *CO* a good heuristic for concept-level *FGSI* at a large scale? Could
CO be enhanced when combined with other heuristics?
- 60 • Is it possible to train DL models for *FGSI* without ground-truth training
data, by using the above heuristics as weak supervision? Could these
models exceed the predictive performance of the heuristics themselves?
- Can the natural extension of MeSH towards fine-grained labels be ex-
ploited for the retrospective evaluation of automated *FGSI* approaches
65 at a large scale?

The rest of this paper is organized as follows. In Section 2, we provide some

⁶This is based on concept recognition by the *MetaMap* tool [3].

⁷<https://www.nlm.nih.gov/research/umls/index.html>

background and previous efforts towards *FGSI* of biomedical literature. In Section 3, we introduce two new methods: *RetroBM* for the development of evaluation datasets and *DBM* for indexing documents with fine-grained labels without ground-truth data. In Section 4, we present the experiments and the results of this work. Finally, in Section 5 we summarise the work and discuss its contribution towards addressing the motivating research questions.

2. Background and related work

In contrast to the broader problem of hierarchical text classification, coarse-grained class labels are already available in the *FGSI* task, which is defined as the refinement of existing coarse-grained labels in this work. Recently, Mekala *et al.* [5] focused on this special case proposing the *Coarse2Fine* (*C2F*) method. *C2F* relies on the literal occurrence of fine-grained labels as initial weak fine-grained labels, which are used together with the ground-truth coarse-grained ones to train a label-conditioned GPT2 [6] model that generates relevant documents for any given label. This model is used to generate new documents for each fine-grained label, which are, in turn, used to train *BERT* classifiers [7].

Semantic indexing of biomedical literature in MEDLINE/PubMed has been done for several decades by human expert indexers of NLM. Since 2002 some automated topic suggestions are offered to human experts by the *Medical Text Indexer* (*MTI*) Tool⁸, and since 2012 a dedicated shared task has been organized for this purpose, in the context of the BioASQ challenge [8]. Despite the particular challenges of this task, such as the huge number of classes, the availability of millions of manually indexed articles and advanced DL architectures allowed the gradual development of machine-learning models that achieve satisfactory results [1, 9, 10]. As a result, NLM gradually moved to fully automated indexing for all MEDLINE articles⁹. However, indexing at the level of MeSH concepts still remains beyond reach, partly due to the lack of adequate data.

⁸<https://lhncbc.nlm.nih.gov/ii/tools/MTI.html>

⁹https://www.nlm.nih.gov/pubs/techbull1/nd21/nd21_medline_2022.html

Early work towards *FGSI* of biomedical literature focused on enriching the
95 descriptor-based labels with MeSH qualifiers and MeSH Supplementary Con-
cept Records (SCRs)¹⁰ [3], which however do not address the coarseness of the
descriptors that aggregate distinct concepts. The importance of indexing the
biomedical literature at the level of MeSH concepts was first noted in[11]. More
recently, we proposed the *Beyond MeSH* method [12, 4] for the automated re-
100 finement of MeSH topic annotations at the level of MeSH concepts. In that
work, we introduced the *concept-occurrence (CO)* heuristic for assigning weak
labels for *FGSI*. This heuristic extends simpler dictionary-based ones using
the *MetaMap* tool, which handles information extraction challenges such as in-
flected and multi-word terms, and concept disambiguation [13]. Experiments
105 with common machine-learning algorithms on small datasets for refining two
specific descriptors showed that the *CO* heuristic is a strong baseline for these
two use cases and only some trained classifiers managed to improve upon it.

Building upon this earlier work, in this paper we perform an extensive eval-
uation of the promising *CO* heuristic in a large-scale scenario for several MeSH
110 descriptors, while also assessing the performance of DL approaches for the first
time in this task, revealing the particular challenges of the *FGSI* task and
proposing respective remedies. In order to form the large-scale experimental
scenario, we look retrospectively into the evolution of MeSH descriptors, focus-
ing on MeSH concepts that over time become descriptors themselves.

115 3. Material and methods

In this section, we introduce a) the *Retrospective Beyond MeSH (RetroBM)*
method for the development of ground-truth datasets for large-scale evaluation
on Fine-Grained Semantic Indexing (*FGSI*), and b) the *Deep Beyond MeSH*
(*DBM*) method for refining coarse-grained MeSH labels into fine-grained ones,
120 for which no ground-truth data are available.

¹⁰https://www.nlm.nih.gov/mesh/intro_record_types.html

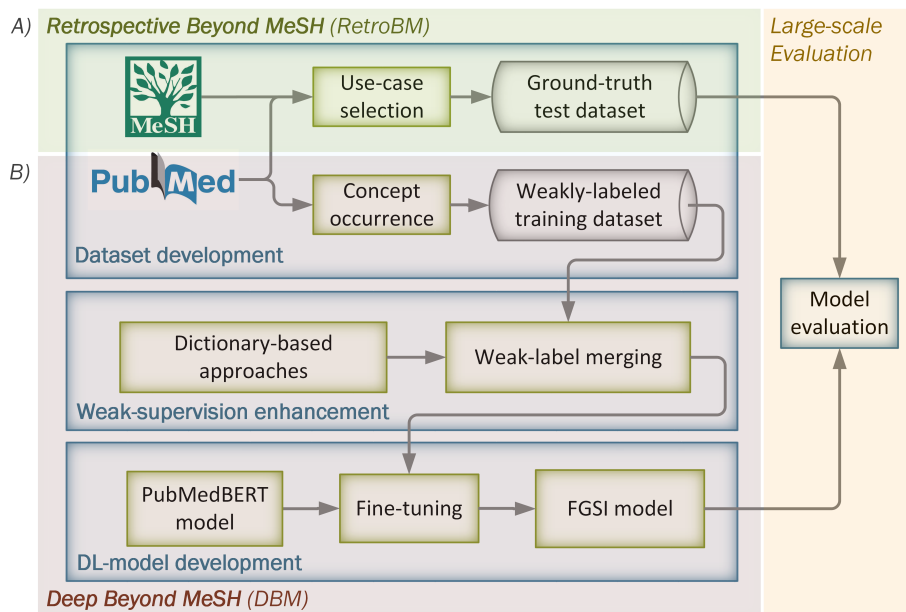


Figure 2: A) The *RetroBM* method for the development of large-scale ground-truth *FGSI* datasets and B) the *DBM* method for weakly-supervised *FGSI* model development. The *RetroBM* datasets are used for the large-scale evaluation of the *DBM* models.

As shown in Figure 2, part A, *RetroBM* first selects some adequate use cases of MeSH concepts promoted to descriptors, by examining retrospectively the evolution of MeSH. Then it develops ground-truth *FGSI* datasets for the task of refining these coarse-grained descriptors, based on the manual annotation already available in MEDLINE/PubMed.

DBM, on the other hand (Figure 2, B), provides *FGSI* for any subordinate MeSH concept, regardless of its promotion as a descriptor, without access to ground-truth labels. *DBM* is structured into three parts: a) the development of weakly-labelled training datasets based on *concept occurrence* (*CO*), b) the enhancement of weak supervision using dictionary-based approaches, and c) the development of *FGSI* models, by fine-tuning pre-trained *PubMedBERT* models.

Finally, the ground-truth *RetroBM* datasets are used both for the validation of the *DBM* method and the evaluation of weakly-supervised *DBM* models for

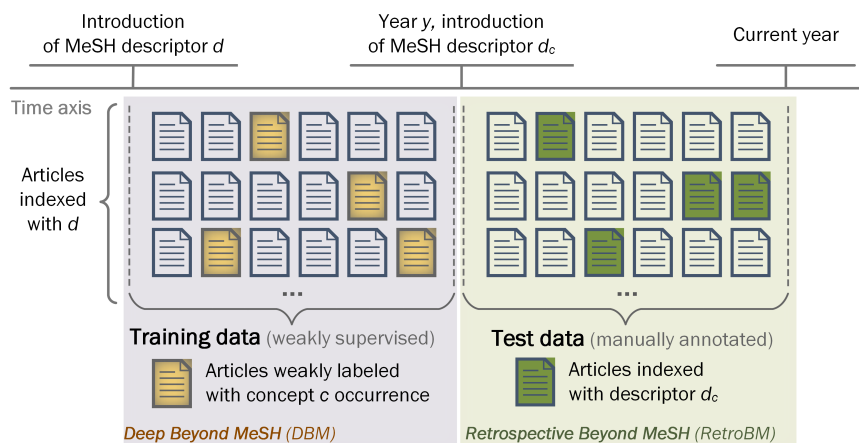


Figure 3: The retrospective evaluation scenario adopted in this work for the development of *FGSI* datasets is based on the promotion of a MeSH concept c into a descriptor d_c .

the selected use cases. These methods are openly available in GitHub¹¹.

135 3.1. Dataset development

Although the proposed *DBM* method develops weakly-labelled datasets to sidestep the need for ground-truth labels in model development, some ground-truth data are still needed for the validation and evaluation of such *DBM* models. Therefore, we also introduce *RetroBM*, an approach based on the
 140 evolution of MeSH indexing already done in MEDLINE/PubMed, to develop some ground-truth *FGSI* datasets.

The two methods complement each other in a retrospective evaluation scenario based on the promotion of a MeSH concept c to a dedicated descriptor d_c in a year y (Figure 3). In this scenario, *RetroBM* uses articles indexed with
 145 d after the year y to develop a ground-truth testset dataset, and *DBM* uses articles indexed with d prior to year y , to develop a respective weakly-labelled training dataset¹². This is done for simplicity and alignment with a prospective

¹¹<https://github.com/tasosnent/DBM>

¹²As articles indexed with d we consider both articles directly annotated with the descriptor d and articles annotated with any descriptor narrower to d in the MeSH hierarchy.

scenario of use, where we develop a model on existing articles indexed with d and use it in new ones. However, new articles could be weakly labelled and used
150 by *DBM* for model training as well.

3.1.1. Retrospective Beyond MeSH (*RetroBM*)

MEDLINE/PubMed provides annotated data for several parent-child pairs of MeSH descriptors. However, not all such pairs represent realistic use cases for evaluating the *FGSI* task where existing coarse-grained annotations need to
155 be refined into fine-grained ones. For example, there are cases where the child descriptor is older than its parent, leading to a situation where annotations for the child are already available and no refinement of parent annotations is actually needed. For this reason, *RetroBM* relies on the evolution of MeSH to identify suitable and realistic concept-promotion cases for *FGSI* evaluation.

160 During the annual update of the MeSH thesaurus for a year y , a subordinate concept c of a coarse-grained descriptor d that becomes increasingly important in the literature, can be promoted to a new MeSH descriptor d_c . In the example of Figure 4, “Niemann-Pick Disease, Type A” was a narrower subordinate concept (c) of the “Niemann-Pick Diseases” descriptor (d) until 2007 (y), when
165 it was promoted to a descriptor (d_c) itself. As a result, since 2007, articles relevant to “Niemann-Pick Disease, Type A” are indexed with a new dedicated descriptor (d_c). In such cases, if c is the only concept of the new descriptor d_c , indexing with d_c is equivalent to indexing with c . Therefore, articles manually indexed after the promotion of c into d_c have ground-truth labels for *FGSI* with
170 c and can be used by *RetroBM* to develop a testset, as illustrated in Figure 3.

In particular, the task of *FGSI* with c is to identify for each article already indexed with d (e.g. “Niemann-Pick Diseases”), whether it is about c (e.g. “Niemann-Pick Disease, Type A”) or not. Therefore, we consider only articles indexed with d and we refer to these articles as *valid* articles for *FGSI* with
175 c . A ground-truth dataset generated by *RetroBM* consists of all these *valid* articles annotated after year y , where the ones indexed with d_c are the positive instances, and the remaining ones are the negative instances for *FGSI* with c .

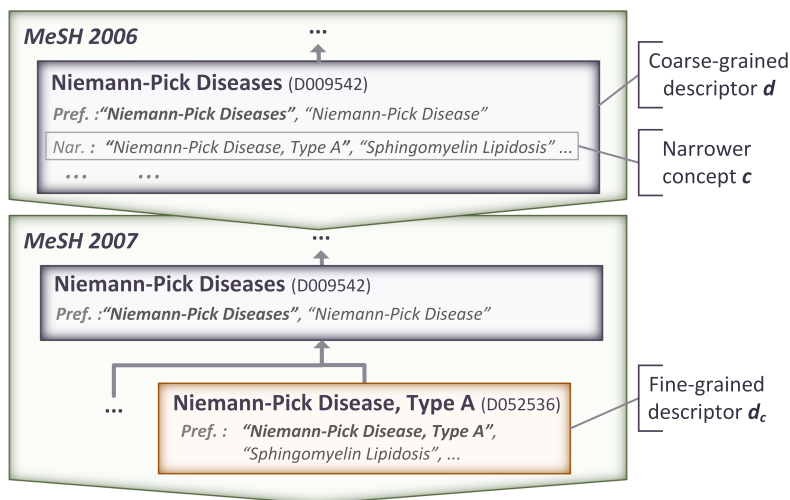


Figure 4: The promotion of “Niemann-Pick Disease, Type A”, from a subordinate concept of the descriptor “Niemann-Pick Diseases”, to a new dedicated descriptor in 2007.

Of course, indexing with d , is not a sufficient condition for indexing with d_c , it is, however, a necessary condition, as d_c is a subcategory of d . On the other hand, articles not indexed with d (non *valid* for c) are beyond the scope of the *FGSI* task, as they can’t be relevant to its subcategory d_c either.

Based on the above idea, *RetroBM* looks retrospectively to identify several such concept-promotion events and develop corresponding datasets, introducing two criteria. First, *RetroBM* relies on a typology that we introduced recently about the provenance of new MeSH descriptors [14]. Out of the types introduced in this work, the one named *subdivision*¹³ represents subordinate concepts c promoted to subcategories d_c of the coarse-gained descriptor d that used to host them. This type satisfies the scenario of Figure 4 and was used by *RetroBM* to identify adequate cases (*criterion 1*). The presence of additional parents for d_c does not affect this scenario, as c always belongs to a single descriptor d . However, by considering other types of evolution, *RetroBM* can be extended to similar problems, such as indexing with emerging concepts, concepts with

¹³We refer to cases with provenance code 1.2 as introduced in Nentidis *et al.* [14].

multiple parents, or even coarse-grained indexing with groups of concepts.

Beyond provenance, each new descriptor d_c in *FGSI* also needs to be fine-grained (*criterion 2*), which breaks into the following: a) d_c should be dedicated
195 to the corresponding promoted concept c , not covering other subordinate concepts as well, ensuring that all articles annotated with d_c are actually relevant to c (*criterion 2.1*). b) d_c should be a leaf in the MeSH hierarchy. Having narrower descriptors representing subcategories indicates that the descriptor covers
200 several concepts, hence it is not fine-grained (*criterion 2.2*).

3.1.2. Weakly-labelled training datasets

The first step of the *DBM* method uses articles annotated with any coarse-grained d to develop weakly-labelled training data for *FGSI* with any subordinate concept c , without any need for manual d_c annotations. In order to generate
205 the weak labels *DBM* uses the heuristic introduced in *Beyond MeSH* [12, 4], which relies on the occurrence of concept c in the title or abstract of the article, as identified by the *MetaMap* tool [13]. *DBM* is applicable to any coarse-grained descriptor d' having at least one narrower subordinate concept c' , in order to predict labels for *FGSI* with c' , regardless of whether c' has been promoted
210 to a descriptor or not. In MeSH 2020, for example, there are about 7,000 such narrower MeSH concepts with sufficient weak supervision for applying *DBM*¹⁴. However, ground-truth c' labels will be needed to verify the accuracy of this annotation process. For this reason, in this work, we employ *DBM* for use cases selected by *RetroBM* as illustrated in the evaluation scenario of Figure 3.

215 3.1.3. Multi-label setup

The *FGSI* task is inherently multi-label as a single article can be relevant to several distinct coarse-grained and fine-grained labels. In addition, in *FGSI* we would like any method to refine a single descriptor (e.g. “Epilepsy”) into any of its sibling subordinate concepts (e.g. “Audiogenic Epilepsy”, “Tactile Reflex

¹⁴For this estimation, we consider any narrower concept c' of a descriptor d' with at least 10 articles having a) d' annotation, and b) c' occurrence.

220 Epilepsy” etc.), which are not usually mutually exclusive. In this work, aligned with the retrospective evaluation scenario, we organize the *FGSI* datasets developed by *RetroBM* and *DBM* based on the promotion year y , gathering all the *valid* articles for any of the use cases corresponding in a given year y together and assigning multi-label annotations accordingly.

225 3.2. Weak-supervision enhancement

The second part of the *DBM* method is the enhancement of the original weak supervision provided by *CO*. In this section, we investigate a range of dictionary-based variants and whether their combination could enhance the quality of the weak labels. The results of this study on a few validation cases, 230 presented in Subsection 4.2, allow the selection of the strategy to be adopted in the second step of the *DBM* method for the enhancement of weak supervision. In particular, adopting the terminology introduced in data programming [15], we refer to these dictionary-based variants and *CO* as labeling functions (LFs) for the *FGSI* task.

235 Each dictionary-based labeling function (LF) assigns a label for concept c to an article if any dictionary element associated with c literally occurs in the title or abstract of the article¹⁵. In the *name exact* LF only the name of c is matched (e.g. “Niemann-Pick Disease, Type A”), and in the *synonyms exact* the synonyms of it (e.g. “Classical Niemann-Pick Disease”, etc). Then, six variants 240 of these two LFs are introduced¹⁶, in order to improve recall (see Table 1). In particular, the dictionary elements are first lower-cased (*name lowercase*, *synonyms lowercase*), then the punctuation marks are removed (*name no punct*, *synonyms no punct*), and finally, phrases are split into single-token elements (*name tokens*, *synonyms tokens*).

245 These dictionary-based LFs rely on the same basic information that is used by *CO*, that is the terms of c and the article text. Still, they may result in

¹⁵An LF with n dictionary elements (e_i) is implemented with the query: e_1 or e_2 ... or e_n .

¹⁶For applying the latter six LFs, the title and abstract of the article are also lower-cased.

Table 1: The dictionary elements of each LF for the label “Niemann-Pick Disease, Type A”.

| Labeling function (LF) | Abbr. | Dictionary elements |
|---------------------------|------------|---|
| <i>name exact</i> | <i>NE</i> | “Niemann-Pick Disease, Type A” |
| <i>synonyms exact</i> | <i>SE</i> | “Classical Niemann-Pick Disease”, ... |
| <i>name lowercase</i> | <i>NL</i> | “niemann-pick disease, type a” |
| <i>synonyms lowercase</i> | <i>SL</i> | “classical niemann-pick disease”, ... |
| <i>name no punct</i> | <i>NNP</i> | “niemann pick disease type a” |
| <i>synonyms no punct</i> | <i>SNP</i> | “classical niemann pick disease”, ... |
| <i>name tokens</i> | <i>NT</i> | “niemann”, “pick”, “disease”, “type”, “a” |
| <i>synonyms tokens</i> | <i>ST</i> | “classical”, “niemann”, “pick”, “disease” ... |

different weak labels. This is because *CO*, calculated by *MetaMap*, is more elaborate than exact term matching, considering additional information such as part-of-speech tagging and word-sense disambiguation based on context [13].

250 Therefore, combining these LFs with *CO* may bring improvements. The correlation between these LFs, on the other hand, may harm the accuracy of the ensemble. For this reason, the optimal LF subset is sought for adoption in *DBM*, by considering all the combinations of two or more of these LFs in the validation experiments (Subsection 4.2).

255 For combining the LFs into an ensemble we examine two voting classifiers based on *majority voting (MV)* and *at-least-one voting (ALO)* [16], as well as a state-of-the-art probabilistic *label modeling (LM)* approach for combining noisy heuristics, introduced in the context of Snorkel [17]. The *MV* schema assigns a label to an article if most of the LFs (voters) assign this label to this article, favoring the precision of the ensemble, as labels predicted by only a few LFs are disregarded as untrustworthy. In the *ALO* schema, on the other hand, a label is assigned to an article if any of the LFs assign this label to this article. This is a special case of the *at-least-N* schema, proposed by Kambhatla [16], where $N = 1$. In this approach, even if a single labeling function predicts a label, this will be included in the prediction of the ensemble, favoring its recall.

260

265

Finally, the Snorkel *LM* approach considers the true label for an article as a latent variable in a probabilistic generative model, which is estimated based on a weighted combination of the LFs [15, 17]. Eventually, a label is assigned to an article if the estimated probability of the label is greater than 0.5.

270 *3.3. Model development*

The final part of the *DBM* method is the development of weakly-supervised Deep Learning (DL) models for the *FGSI* task. In this direction, motivated by experiments on other tasks suggesting that deep pre-trained models can have an inherent resilience to label noise [18, 19], we investigate the adequacy of a state-of-the-art approach in text classification, that is fine-tuning a pre-trained *BERT*-based model, in our case a *PubMedBERT* model [2], with one additional task-specific output layer [7]. In particular, we experimented on a few validation cases, in order to identify the particular challenges of the *FGSI* task and adopt remedies to tackle them. These include preliminary experiments for basic design choices discussed in this section, as well as the validation experiments presented in Subsection-4.2.

In *DBM*, we opt for a multi-label scenario, where a distinct multi-label model is developed for each year y , as depicted in Figure 5¹⁷. In this architecture, the title and abstract of each article are concatenated, tokenized, and provided as input to a *PubMedBERT* model. The embedding of the classification token *CLS* is then provided as input in a fully-connected output layer that produces a vector of size equal to the number of labels as in [20]. This way, knowledge from different labels is integrated into shared deep representations for all the use cases of a single year. This can be useful, particularly for sibling concepts of a single coarse-grained descriptor, such as “Audiogenic Epilepsy” and “Tactile Reflex Epilepsy” in the example of Figure 1. Preliminary experiments confirmed that multi-label models are at least as good as single-label models in

¹⁷A common model for many years is possible but would pose experimental complications as a single article could end up in the test set of one use case and the training set of another.

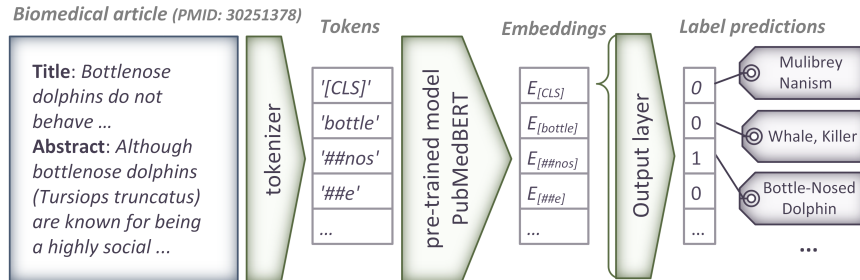


Figure 5: A multi-label model is developed by *DBM* for all the labels of a single year.

a One-Versus-Rest approach. Therefore, the additional cost of developing and storing multiple single-label models was not justified.

295 Regarding the specific *BERT*-based model to fine-tune, we opted for *Pub-MedBERT* [2], a domain-specific variant of *BERT* adopting a domain-specific vocabulary to better capture the semantics of the text. Although differences in the preliminary results were small, they confirmed the good performance of *PubMedBERT* and it was thus chosen as the basis for all *DBM* models. The
 300 preliminary experiments also indicated that the number of training epochs can affect the performance of the model. For this purpose, we split the weakly-labelled training data into a *training part* and a *validation part* (90-10%) in order to train only on the first and check the loss in the unseen *validation part*, identifying in which epoch this loss is minimised (*best_ep*).

305 For a higher generalization level, we also adopted an early stopping strategy that concludes the fine-tuning process once the loss on the *validation part* starts increasing [21]. In this case, the *best_ep* model is always that of the penultimate epoch of the early stopped fine-tuning. Furthermore, we also considered the *prev_ep* and *next_ep* models that correspond to the epochs just before or after
 310 *best_ep*. The motivation was that the “imperfection” of epochs close to the best-performing one may be due to a higher level of generalization.

A major challenge in training *FGSI* models is the imbalance of the labels in two ways: a) The number of positive instances for each label varies among use cases, favoring more popular fine-grained labels, and b) there are many

315 more negative instances than positive ones. In extreme cases, this abundance
of negative instances can prevent the models from learning to recognize positive
ones, as their overall effect on the loss is minimal.

For imbalance among labels, we considered extensions of the well-established
binary cross entropy (BCE) loss [22] that can handle imbalanced training data
320 for multi-label text classification [23]. In particular, the rebalanced focal loss (R-
BCE-FL) is an extension of BCE loss, first introduced to handle label imbalance
in multi-label image classification tasks [24] and recently extended to multi-
label text classification tasks [23]. R-BCE-FL modifies BCE in two ways. First,
the instances are rebalanced, by oversampling based on the overall frequency
325 of corresponding labels. Second, a focusing parameter increases the weight of
more hard-to-classify instances, based on their predicted probability [23].

For the imbalance between positive and negative instances, previous work
has shown that under-sampling can help address this problem [4]. Therefore,
we experimented with an under-sampling procedure removing negative instances
330 to achieve an ideal negative-to-positive ratio (*balance_n*) for each label. This
ideal ratio is not known beforehand, but a reasonable estimate can be provided
through validation. Removing random negative instances to enforce the ideal
negative-to-positive ratio for a label strictly, may deplete the positive or *valid*
negative instances for another. For this reason, the undersampling process of
335 *DBM* does not remove instances that are either positive for any of the labels
or are *valid* for labels with a low negative-to-positive ratio.

Randomness is introduced in several steps of the model development pro-
cess [25], notably in the splitting of training and validation parts, in the under-
sampling of negative instances, and in stochastic optimization [26]. The prelimi-
340 nary experiments showed that starting with different random seeds can produce
different results. For this reason, each experiment was repeated with six differ-
ent seeds, and the predictions were aggregated with a majority voting approach,
which led to a more stable estimate of model performance.

3.4. Competing machine learning approaches

345 Three machine-learning approaches have been included in our experiments
for comparison, based on the most recent related work (Section 2). First, *Beyond MeSH* [4] was considered, as a prior work on the *FGSI* task of refining
MeSH topic annotations at the level of MeSH concepts. In this direction, a
logistic regression (*LR*) model was developed for each fine-grained label consid-
350 ering both lexical and semantic features. For a fair comparison with *DBM*, we
perform the same under-sampling of negative instances, which is repeated for
six seeds combining the results with majority voting. The level of regularization
is selected via grid search on a *validation part* (10%) of the weakly-labelled data,
considering a range of values from 0.01 to one billion¹⁸. A reasonable number of
355 features to be selected was chosen based on the validation experiments presented
in 4.2.

Second, *Coarse2Fine* (*C2F*) [5] was also considered, as a state-of-the-art
work on the broader task of refining existing coarse-grained labels into fine-
grained classes. Although the *C2F* originally relies on the literal occurrence
360 of label names, for a fair comparison against *DBM*, we provided *C2F* with
the same weak supervision. In addition, *C2F* handles multi-class problems,
where each document belongs to exactly one of the fine-grained sub-categories
of a coarse-grained label d . In order to handle the “unlabelled” examples in
FGSI, that belong to d but not to any of its sub-categories, we introduce an
365 additional subcategory (d_{rest}) for each d . The predictions for d_{rest} , however,
are disregarded when estimating the performance of the *C2F* models, which
is based on the actual subcategories, as done for all methods. As *C2F* adds
generated instances to handle imbalance, we did not undersample the data.
In addition, we did not experiment with different random seeds or parameters
370 relying on the original implementation¹⁹.

Finally, in order to investigate the impact of generating instances as per

¹⁸Namely these values were considered: 0.01, 0.1, 1, 10, 100, 1K, 10K, 100K, 1M, 1B.

¹⁹<https://github.com/dheeraj7596/C2F>

C2F, we also developed DBM^{C2F} models using the implementation of *DBM* with the documents generated by *C2F* as weak supervision.

3.5. Evaluation measures

375 As both precision (P) and recall (R) are important for *FGSI*, we adopt the F1-measure ($F1$) as an overall performance measure, considering both macro-averaging ($maF1$, maP , maR) and micro-averaging ($miF1$, miP , miR) across different use cases [27]. In the multi-label *FGSI* datasets articles that are not *valid* for a specific fine-grained label c are not of interest and are disregarded
380 in the evaluation of that label. We refer to this process as *validity filtering* and employ it in the evaluation of all the methods.

4. Experiments and results

RetroBM was applied to fourteen consecutive years (2006-2019) of MeSH evolution for developing respective ground-truth datasets ($test_y$). Then, *DBM*
385 developed respective weakly-labelled training datasets (WS_dev_y) for *FGSI*-model development (Subsection 4.1). Out of these datasets, we used the ones for labels introduced in 2006 (*2006 datasets*) for “validation experiments”, in order to make design choices about the architecture and configurations of the *DBM* method (Subsection 4.2). Finally, the datasets for unseen fine-grained
390 labels introduced in the subsequent years (*2007-2019 datasets*) were used in the “evaluation experiments” (Subsection 4.3), which aimed to confirm the choices made on the *2006 datasets*.

4.1. Retrospective datasets

In total, 88 new fine-grained descriptors were found to match the criteria
395 of *RetroBM*, which was employed to retrospectively process MEDLINE data from 2006 to 2019²⁰. In particular, the 2020 version of the MEDLINE/PubMed

²⁰The full list of descriptors per year is available in the supplementary file S1.

Annual Baseline Repository was used to retrieve documents and MeSH annotations. In addition, SemMedDBv4.0 was used by *DBM* for *CO* and use cases without enough data for training and evaluation were excluded from the experiments²¹.

In particular, *RetroBM* identified 18 descriptors introduced in 2006 developing the *test₂₀₀₆* ground-truth dataset and *DBM* developed the respective weakly-labelled dataset *WS_dev₂₀₀₆*. These two datasets (*2006 datasets*) consisted of more than 400 thousand articles each and were used for validation experiments²². The remaining 70 descriptors, introduced after 2006, were used for the evaluation experiments (*2007-2019 datasets*). The total size of the thirteen weakly-labelled datasets (*WS_dev₂₀₀₇₋₂₀₁₉*) and the respective ground-truth ones (*test₂₀₀₇₋₂₀₁₉*), exceeds 2M and 1M articles respectively²³.

In these datasets, the positive instances are really scarce highlighting the importance of balancing approaches adopted in *DBM*. In *test₂₀₀₆*, in particular, the positive instances for any of the 18 labels are almost 6% of all articles in *test₂₀₀₆*. In addition, about 3% and 5% of the articles in *WS_dev₂₀₀₆* and *test₂₀₀₆* are positive based on the weak labels. These frequencies are similar for the 70 labels of *2007-2019 datasets* as well, considering both the weak and the ground-truth labels²⁴. This suggests that observations on older labels can be useful for labels introduced later, confirming the rationality of a retrospective scenario.

4.2. Validation experiments on 2006 datasets

The first aim of the validation experiments was to confirm the strength of *concept occurrence (CO)* as the weak-supervision source of *DBM*. Table 2 presents LF performance for the *FGSI* task on *test₂₀₀₆*. *CO* was clearly the

²¹We kept use cases with at least 10 articles with a) *c* occurrence in *WS_dev_y*, and b) *d_c* annotations in *test_y*. A too-general use case with more than 1.5M articles for a single *d* was also excluded for practical reasons.

²²*WS_dev₂₀₀₆* 446,377 articles, *test₂₀₀₆* 468,271 articles.

²³*WS_dev₂₀₀₇₋₂₀₁₉* 2,068,965 articles, *test₂₀₀₇₋₂₀₁₉* 1,135,657 articles.

²⁴Detailed statistics for each dataset are available in the supplementary file S1.

Table 2: *FGSI* results of each labeling function on the *test*₂₀₀₆ ground-truth data. *var* stands for the variance of *maF1* across the 18 labels.

| Labeling function (LF) | <i>maP</i> | <i>maR</i> | <i>maF1</i> \pm <i>var</i> | <i>miP</i> | <i>miR</i> | <i>miF1</i> |
|--------------------------------|------------|------------|------------------------------|------------|------------|--------------|
| <i>concept occurrence (CO)</i> | 0.715 | 0.618 | 0.634 \pm 0.06 | 0.737 | 0.531 | 0.618 |
| <i>name exact (NE)</i> | 0.707 | 0.082 | 0.112 \pm 0.04 | 0.901 | 0.027 | 0.052 |
| <i>synonyms exact (SE)</i> | 0.656 | 0.212 | 0.251 \pm 0.12 | 0.889 | 0.046 | 0.088 |
| <i>name lowercase (NL)</i> | 0.649 | 0.441 | 0.446 \pm 0.12 | 0.589 | 0.203 | 0.302 |
| <i>synonyms lowercase (SL)</i> | 0.620 | 0.365 | 0.408 \pm 0.11 | 0.709 | 0.251 | 0.371 |
| <i>name no punct (NNP)</i> | 0.574 | 0.440 | 0.442 \pm 0.13 | 0.583 | 0.196 | 0.293 |
| <i>synonyms no punct (SNP)</i> | 0.597 | 0.349 | 0.385 \pm 0.11 | 0.690 | 0.206 | 0.318 |
| <i>name tokens (NT)</i> | 0.322 | 0.838 | 0.377 \pm 0.07 | 0.157 | 0.807 | 0.263 |
| <i>synonyms tokens (ST)</i> | 0.274 | 0.874 | 0.338 \pm 0.07 | 0.081 | 0.883 | 0.149 |

strongest LF, which confirms previous observations on smaller-scale experiments [4]. In addition, we observe that most LFs achieved better precision than recall, except those based on tokenization (*NT*, *ST*)²⁵. The average correlation between *CO* and each other LF was 0.48 with *SL*, *NNP*, and *SNP* being the most correlated (0.68) and *ST* the least correlated (0.21) with *CO*.
425 Similarly, the average correlation between the dictionary-based LFs themselves was 0.35, with [*SL*, *SNP*] being the most correlated pair (0.95) and [*SL*, *ST*] being the least correlated one (0.06)²⁶.

The second aim of these experiments was to confirm whether *CO* can be
430 enhanced when ensembled with other LFs, choosing a) which LF subset and b) which ensemble approach to adopt in *DBM* for weak-supervision enhancement. In this direction, for each LF subset, the three ensemble approaches were assessed for the *FGSI* task on *test*₂₀₀₆, namely *majority voting (MV)*, *at-least-one (ALO)*, and *Snorkel label model (LM)*. The results confirmed that several

²⁵Some tokens in these LFs can be too-general such as “disease”, “type”, and “a” in Table 1.

²⁶The full correlation matrix for all LFs is available in the supplementary file S2.

435 LF subsets led to performance improvements upon *CO* alone²⁷.

Overall, two LF subsets²⁸ achieved the top *maF1* (0.689) and *miF1* (0.646) scores, under the *ALO* or the *LM* ensemble, outperforming *CO* by 5.5pp and 2.8pp respectively. In addition, six more LF subsets led to these top scores under the *ALO* approach only. Interestingly, *CO*, *NL*, and *SL* are present in
440 all top-performing LF subsets, regardless of the ensemble approach, and the combination of only these three LFs was sufficient to achieve top performance with both *ALO* and *LM*. As a result, the subset [*CO*, *NL*, *SL*] was adopted for weak supervision enhancement in *DBM*.

The fact that *ALO* achieved top results with more LF subsets than *LM*²⁹,
445 suggests that *ALO* was more robust to the choice of LF subset in this setup. A reason for this could be the high precision of most LFs and the fact that none of the top-performing subsets included the low-precision LFs (*NT*, *ST*), as *ALO* is known to be good at combining precise voters by improving the total recall [16]. In particular, *ALO* assigns a label to an article if at least one labeling function
450 does, not affected by potential negative votes of more cautious LFs. As a result, the *ALO* ensemble of [*CO*, *NL*, *SL*] (*ALO*₃) was adopted as the enhanced weak supervision schema in *DBM*.

Another aim of the validation experiments on *2006 datasets* was to make specific choices finalizing the model-development part of *DBM*. In particular,
455 lar, we need to a) estimate a reasonable value for the negative-to-positive ratio (*balance_n*), and b) choose the model of which epoch to use for *DBM* predictions, that is, whether to use the model of the *best_{ep}* that minimized the loss on the *validation part* of *WS_{dev_y}*, or the epoch just before (*prev_{ep}*) or just after *best_{ep}* (*next_{ep}*).

460 In this direction, different *DBM* models were developed on *WS_{dev₂₀₀₆}* and assessed for their *FGSI* performance on *test₂₀₀₆*. In particular, experi-

²⁷The full results for all LFs subsets are available in the supplementary file S3.

²⁸[*CO*, *NL*, *SL*] and [*CO*, *NL*, *SL*, *NNP*].

²⁹8 top-performance subsets with *ALO* against 2 with *LM*.

Table 3: *FGSI* validation results on *test*₂₀₀₆ data. WS stands for the source of weak supervision for trained models. A subscript in *LR* indicates the number of features selected. A subscript in *DBM* indicates the training epoch, with *p*, *b*, and *n* corresponding to *prev_ep*, *best_ep*, and *next_ep* respectively. *var* stands for the variance of *maF1* across the 18 labels.

| LF/model | WS | <i>maP</i> | <i>maR</i> | <i>maF1</i> ± <i>var</i> | <i>miP</i> | <i>miR</i> | <i>miF1</i> |
|--------------------------------|-------------------------|------------|------------|--------------------------|------------|------------|--------------|
| <i>CO</i> | - | 0.715 | 0.618 | 0.634 ±0.06 | 0.737 | 0.531 | 0.618 |
| <i>LR</i> ₅ | <i>CO</i> | 0.427 | 0.822 | 0.474 ±0.10 | 0.295 | 0.655 | 0.406 |
| <i>LR</i> ₁₀ | <i>CO</i> | 0.358 | 0.807 | 0.396 ±0.09 | 0.102 | 0.629 | 0.176 |
| <i>LR</i> ₁₀₀ | <i>CO</i> | 0.275 | 0.809 | 0.318 ±0.06 | 0.085 | 0.620 | 0.149 |
| <i>LR</i> ₁₀₀₀ | <i>CO</i> | 0.279 | 0.798 | 0.310 ±0.07 | 0.134 | 0.644 | 0.221 |
| <i>ALO</i> ₃ | - | 0.717 | 0.698 | 0.689 ±0.03 | 0.713 | 0.591 | 0.646 |
| <i>DBM</i> _{<i>p</i>} | <i>ALO</i> ₃ | 0.681 | 0.744 | 0.700 ±0.03 | 0.689 | 0.624 | 0.655 |
| <i>DBM</i> _{<i>b</i>} | <i>ALO</i> ₃ | 0.688 | 0.738 | 0.702 ±0.03 | 0.690 | 0.626 | 0.656 |
| <i>DBM</i> _{<i>n</i>} | <i>ALO</i> ₃ | 0.689 | 0.741 | 0.702 ±0.03 | 0.695 | 0.623 | 0.657 |

ments with different *balance_n* values ranging from 1 to 100, led to the choice of 10-to-1 as a reasonable negative-to-positive ratio, which was adopted for all remaining experiments. In addition, the *FGSI* results of *DBM* experiments with the three alternatives for the training epoch (Table 3), suggest that the model corresponding to *next_ep* (*DBM*_{*n*}) led to the best results for both metrics. Therefore, the model of *next_ep* was adopted as the final one in *DBM*.

Finally, validation experiments with logistic regression (*LR*) models on the same *FGSI* task (*2006 datasets*), also presented in Table 3, allowed the selection of the number of features to be considered in the evaluation experiments. In particular, *LR* models were trained on *WS_dev*₂₀₀₆ considering different numbers of features for feature selection with the F-ANOVA statistic, namely 5, 10, 100, and 1000. The results reveal that no *LR* model managed to outperform *CO* overall. In particular, the best *LR* model, namely *LR*₅ with five features,

475 managed to outperform *CO* in *F1* for about 40% of the 18 labels on *test*₂₀₀₆³⁰.

4.3. Evaluation on new labels introduced after 2006

In this last set of experiments, we assessed *DBM* and the competing methods on *FGSI* with the 70 unseen labels of *2007-2019 datasets*. The enhancement of *CO* supervision into *ALO*₃ is a key part of *DBM*. However, in order to
480 investigate *ALO*₃ importance we also trained *DBM* models directly on *CO* alone (*coDBM*). For the same reason, we trained the competing models on both the original (*CO*) and the enhanced (*ALO*₃) supervision. The results of these experiments on the *FGSI* task are presented in Table 4³¹.

These results confirm that *CO* is a strong heuristic for *FGSI*, outperforming
485 again the dictionary-based heuristics. In addition, the weak-supervision enhancement schema (*ALO*₃) achieved an overall *maF1* about 5pp higher than *CO* alone, which is statistically significant³². The performance of the *Snorkel Label model* of the same LF subset (*LM*₃) was not significantly lower than that of *ALO*₃, as the two approaches achieved practically the same *F1* score for most
490 of the 70 labels. However, *ALO*₃ achieved *F1* performance equal to or better than that of *LM*₃ for all 70 labels.

Regarding trained models, the results show that logistic regression (*LR*) models did not manage to outperform the supervision provided to them overall, neither *CO* nor *ALO*₃. The overall *maF1* performance of the *C2F* models was
495 also low, achieving top recall but low precision. In particular, *C2F* trained on *CO* did better than *CO* itself for only 13 out of the 70 labels. The low performance of a state-of-the-art method like *C2F* is an indication that the *FGSI* task and datasets have particular challenges, such as the high abundance of “unlabelled” documents that do not belong to any of the labels of interest.

500 On the other hand, the *coDBM*_{*n*} model developed using the weak supervision of *CO* alone led to a limited overall performance gain of about 1.5pp in

³⁰Detailed validation results per descriptor are available in the supplementary file S4.

³¹Detailed evaluation results per descriptor are available in the supplementary file S5.

³²Wilcoxon signed-rank test, $H_1 : F1\ CO < F1\ ALO_3$, p-value = 0.000001.

Table 4: *FGSI* evaluation results on *test*_{2007–2019} data. WS stands for the source of weak supervision for trained models. A subscript in *LR* models indicates the number of features selected. A subscript in *DBM* models indicates the training epoch considered (*n* for *next.ep*). *var* stands for the variance of *maF1* across the 70 labels.

| LF/model | WS | <i>maP</i> | <i>maR</i> | <i>maF1</i> ± <i>var</i> | <i>miP</i> | <i>miR</i> | <i>miF1</i> |
|--|-------------------------|------------|------------|--------------------------|------------|------------|--------------|
| <i>CO</i> | - | 0.766 | 0.590 | 0.626 ±0.06 | 0.848 | 0.554 | 0.670 |
| <i>NL</i> | - | 0.647 | 0.369 | 0.408 ±0.12 | 0.880 | 0.217 | 0.348 |
| <i>SL</i> | - | 0.567 | 0.253 | 0.287 ±0.09 | 0.809 | 0.125 | 0.217 |
| <i>LR</i> ₅ | <i>CO</i> | 0.420 | 0.807 | 0.458 ±0.09 | 0.204 | 0.783 | 0.324 |
| <i>C2F</i> | <i>CO</i> | 0.206 | 0.946 | 0.276 ±0.09 | 0.074 | 0.981 | 0.138 |
| <i>coDBM</i> _{<i>n</i>} | <i>CO</i> | 0.718 | 0.647 | 0.641 ±0.06 | 0.826 | 0.573 | 0.677 |
| <i>coDBM</i> _{<i>n</i>} ^{<i>C2F</i>} | <i>CO</i> | 0.712 | 0.626 | 0.625 ±0.06 | 0.831 | 0.568 | 0.675 |
| <i>MV</i> ₃ | - | 0.782 | 0.481 | 0.542 ±0.08 | 0.873 | 0.281 | 0.462 |
| <i>LM</i> ₃ | - | 0.756 | 0.652 | 0.661 ±0.06 | 0.842 | 0.569 | 0.679 |
| <i>ALO</i> ₃ | - | 0.764 | 0.668 | 0.677 ±0.05 | 0.839 | 0.583 | 0.688 |
| <i>LR</i> ₅ | <i>ALO</i> ₃ | 0.403 | 0.825 | 0.454 ±0.09 | 0.196 | 0.797 | 0.315 |
| <i>C2F</i> | <i>ALO</i> ₃ | 0.203 | 0.959 | 0.272 ±0.09 | 0.076 | 0.983 | 0.141 |
| <i>DBM</i> _{<i>n</i>} | <i>ALO</i> ₃ | 0.725 | 0.695 | 0.672 ±0.04 | 0.839 | 0.621 | 0.714 |
| <i>DBM</i> _{<i>n</i>} ^{<i>C2F</i>} | <i>ALO</i> ₃ | 0.726 | 0.679 | 0.656 ±0.05 | 0.803 | 0.588 | 0.679 |

terms of *maF1* over *CO* itself. The performance of the *DBM*_{*n*} model, using the *enhanced supervision* (*ALO*₃), was the best trained model, achieving about 4.6pp improvement over *CO* in *maF1*, which is statistically significant³³. However, even *DBM*_{*n*} achieves only comparable performance to *ALO*₃ itself. In fact, *ALO*₃ alone has a slightly higher *maF1*, without statistical significance³⁴. Overall, the *DBM* method, starting from *CO* as a strong heuristic, managed to enhance it further and achieved the refinement of several coarse-grained labels to the level of 70 unseen fine-grained concepts, without access to ground-truth

³³Wilcoxon signed-rank test, $H_1 : F1\ CO < F1\ DBM_n$, p-value = 0.0065.

³⁴Wilcoxon signed-rank test, $H_1 : F1\ ALO_3 \neq F1\ DBM_n$, p-value = 0.088.

510 labels. These results suggest that some DL models adopting the design choices of *DBM* can indeed improve upon the weak supervision that they rely on.

5. Discussion and conclusions

In this work, we investigate the task of fine-grained semantic indexing (*FGSI*), without ground-truth training data, and the development of weakly-supervised 515 deep-learning (DL) models for it. Additionally, we propose an automated approach to generating large-scale ground-truth datasets for *FGSI*, based on MEDLINE articles. Hence, the main contributions of the paper are:

- The confirmation of *concept occurrence (CO)* as a strong heuristic for *FGSI* at a large-scale.
- 520 • A new weakly-supervised method, *Deep Beyond MeSH (DBM)*, for the *FGSI* task. This method introduces a novel schema for enhancing the original weak supervision (*ALO₃*) and makes specific design choices that allowed the development of DL models that can bring improvements upon the weak supervision itself.
- 525 • A new method, *Retrospective Beyond MeSH (RetroBM)*, for the automated development of ground-truth datasets for the large-scale evaluation of *FGSI* predictions.

RetroBM, exploiting the extension of MeSH towards more fine-grained labels, developed *FGSI* data sets with ground-truth annotations for 88 distinct fine- 530 grained labels. To the best of our knowledge, this is the first large-scale set of ground-truth datasets for *FGSI*, consisting of more than one million articles in total. These data allowed us to confirm the strength of *CO* as a heuristic for *FGSI* of biomedical literature and proceed with the validation and evaluation experiments of this work.

535 *DBM*, first uses the *CO* heuristic to develop weakly labelled *FGSI* datasets. Then, it significantly enhances these weak labels by combining *CO* with the dictionary-based heuristics *name lowercase (NL)* and *synonyms lowercase (SL)*,

in an *at-least-one* ensemble (ALO_3). The two heuristics examine the literal occurrence of the name and synonyms of the concept in an article. This enhancement is particularly inserting, considering that NL and SL rely on the same concept terms already exploited by CO itself, but use it in different ways leading to the detection of relevant articles missed by CO alone.

Finally, DBM develops DL models for the $FGSI$ task without access to ground-truth labels by using the enhanced ALO_3 supervision and building upon the state-of-the-art approach of fine-tuning pre-trained models. Experiments with existing methods, namely *logistic regression* (LR) and *Coarse2Fine* ($C2F$), concluded that the $FGSI$ task is challenging and they cannot improve upon the results of the weak labels themselves. DBM on the other hand adopts design choices to tackle the challenges of the $FGSI$ task. In particular, for handling the imbalance of the $FGSI$ data, DBM uses the R-BCE-FL loss and under-sampling of negative instances. For avoiding overfitting on the weak supervision DBM relies on early stopping considering the epoch next to the one that minimises loss (*next_ep*). Finally, for limiting the effect of randomness it adopts majority voting among models developed with different random seeds.

The evaluation results reveal that no trained model managed to outperform the overall macro F1 (*maF1*) performance of ALO_3 itself, with the DBM models achieving only a comparable performance. Nevertheless, the DBM models achieved competitive $FGSI$ performance overall, reaching micro F1 (*miF1*) levels above the very competitive enhanced weak labeling (ALO_3) itself on the new unseen labels of the evaluation experiments. These results suggest that DL-based weakly-supervised $FGSI$ with DBM is a promising direction for supporting fine-grained access to the biomedical literature.

As regards *RetroBM*, an interesting extension would be to develop datasets for retrospective investigation of other indexing tasks, such as indexing with descriptors for groups of concepts, by adjusting the use-case selection criteria. As regards the improvement of the promising DBM method for DL-based weakly-supervised $FGSI$, we plan to use additional label information, such as hierarchical relationships, more fine-grained text representations based on con-

trastive learning approaches, and probabilistic weak labels.

570 **6. Acknowledgment**

This research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the HFRI Ph.D. Fellowship grant (Fellowship Number: 697). The data resources considered in this work, including MeSH, MEDLINE/PubMed, and SemMedDB, were accessed courtesy of the U.S. National Library of Medicine. We are grateful to Georgios Katsimpras for the 575 useful discussions on Deep Learning.

References

- [1] A. Nentidis, G. Katsimpras, E. Vandorou, A. Krithara, A. Miranda-Escalada, L. Gasco, M. Krallinger, G. Paliouras, Overview of BioASQ 580 2022: The Tenth BioASQ Challenge on Large-Scale Biomedical Semantic Indexing and Question Answering, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Vol. 13390 LNCS, 2022, pp. 337–361. [arXiv:2210.06852](https://arxiv.org/abs/2210.06852), [doi:10.1007/978-3-031-13643-6_22](https://doi.org/10.1007/978-3-031-13643-6_22).
- [2] Y. Gu, R. Tinn, H. Cheng, M. Lucas, N. Usuyama, X. Liu, T. Naumann, 585 J. Gao, H. Poon, Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing, *ACM Transactions on Computing for Healthcare* 3 (1) (2022) 1–23. [doi:10.1145/3458754](https://doi.org/10.1145/3458754).
- [3] A. R. Aronson, J. G. Mork, C. W. Gay, S. M. Humphrey, W. J. Rogers, 590 The NLM Indexing Initiative’s Medical Text Indexer., *Studies in health technology and informatics* 107 (Pt 1) (2004) 268–72. [doi:10.3233/978-1-60750-949-3-268](https://doi.org/10.3233/978-1-60750-949-3-268).
- [4] A. Nentidis, A. Krithara, G. Tsoumakas, G. Paliouras, Beyond MeSH: Fine-grained semantic indexing of biomedical literature based on weak su-

- 595 pervision, *Information Processing & Management* 57 (5) (2020) 102282.
doi:10.1016/j.ipm.2020.102282.
- [5] D. Mekala, V. Gangal, J. Shang, Coarse2Fine: Fine-grained Text Classification on Coarsely-grained Annotated Data, in: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Stroudsburg, PA, USA, 2021, pp. 600 583–594. arXiv:2109.10856, doi:10.18653/v1/2021.emnlp-main.46.
- [6] I. S. Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, *Language Models are Unsupervised Multitask Learners*, Tech. Rep. May (2020).
605 URL https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf
- [7] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. doi:10.18653/v1/N19-1423.
610
- [8] G. Tsatsaronis, G. Balikas, P. Malakasiotis, I. Partalas, M. Zschunke, M. R. Alvers, D. Weissenborn, A. Krithara, S. Petridis, D. Polychronopoulos, et al., An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition, *BMC bioinformatics* 16 (1) (2015) 615 138. doi:10.1186/s12859-015-0564-6.
- [9] R. You, Y. Liu, H. Mamitsuka, S. Zhu, BERTMeSH: deep contextual representation learning for large-scale high-performance MeSH indexing with full text, *Bioinformatics* 37 (5) (2021) 684–692. doi:10.1093/bioinformatics/btaa837.
620
- [10] A. R. Rae, J. G. Mork, D. Demner-Fushman, A neural text ranking approach for automatic MeSH indexing, in: *CEUR Workshop Proceedings*,

Vol. 2936, 2021, pp. 302–312.

625 URL <http://ceur-ws.org/Vol-2936/paper-22.pdf>

- [11] S. J. Darmoni, L. F. Soualmia, C. Letord, M.-C. Jaulent, N. Griffon, B. Thirion, A. Névéal, Improving information retrieval using Medical Subject Headings Concepts: a test case on rare and chronic diseases, *Journal of the Medical Library Association : JMLA* 100 (3) (2012) 176–183.
630 doi:10.3163/1536-5050.100.3.007.
- [12] A. Nentidis, A. Krithara, G. Tsoumakas, G. Paliouras, Beyond MeSH: Fine-Grained Semantic Indexing of Biomedical Literature Based on Weak Supervision, in: 2019 IEEE 32nd International Symposium on Computer-Based Medical Systems (CBMS), IEEE, 2019, pp. 180–185. doi:10.1109/CBMS.2019.00045.
635 CBMS.2019.00045.
- [13] A. R. Aronson, F.-M. Lang, An overview of MetaMap: historical perspective and recent advances, *Journal of the American Medical Informatics Association* 17 (3) (2010) 229–236. doi:10.1136/jamia.2009.002733.
- [14] A. Nentidis, A. Krithara, G. Tsoumakas, G. Paliouras, What is all this new MeSH about?, *International Journal on Digital Libraries* 22 (4) (2021)
640 319–337. arXiv:2101.08293, doi:10.1007/s00799-021-00304-z.
- [15] A. Ratner, C. De Sa, S. Wu, D. Selsam, C. Ré, Data Programming: Creating Large Training Sets, Quickly (Nips) (2016) 1–9. arXiv:1605.07723.
- [16] N. Kambhatla, Minority Vote: At-Least-N Voting Improves Recall for Extracting Relations, in: COLING/ACL 2006 - 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Main Conference Poster Sessions, no. July, 2006, pp. 460–466.
645 URL <https://aclanthology.org/P06-2060>
- [17] A. Ratner, S. H. Bach, H. Ehrenberg, J. Fries, S. Wu, C. Ré, Snorkel: Rapid Training Data Creation with Weak Supervision, Proceedings of the
650

VLDB Endowment 11 (3) (2017) 269–282. [arXiv:1711.10160v1](#), [doi:10.14778/3157794.3157797](#).

- 655 [18] M. Tänzler, S. Ruder, M. Rei, Memorisation versus Generalisation in Pre-trained Language Models 03. [arXiv:2105.00828](#).
- [19] D. Hendrycks, K. Lee, M. Mazeika, Using Pre-Training Can Improve Model Robustness and Uncertainty, 36th International Conference on Machine Learning, ICML 2019 2019-June (2018) (2019) 4815–4826. [arXiv:1901.09960](#).
- 660 [20] A. Garcia-Silva, J. M. Gomez-Perez, Classifying scientific publications with BERT - is self-attention a feature selection method?, in: D. Hiemstra, M. Moens, J. Mothe, R. Perego, M. Potthast, F. Sebastiani (Eds.), Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 - April 1, 2021, Proceedings, Part I, Vol. 12656 of Lecture Notes in Computer Science, Springer, 2021, pp. 161–175. [doi:10.1007/978-3-030-72113-8_11](#).
- 665 [21] L. Prechelt, Early Stopping — But When?, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 53–67. [doi:10.1007/978-3-642-35289-8_5](#).
- [22] Y. Bengio, A. Courville, P. Vincent, Representation Learning: A Review and New Perspectives, IEEE Transactions on Pattern Analysis and Machine Intelligence 35 (8) (2013) 1798–1828. [doi:10.1109/TPAMI.2013.50](#).
- 670 [23] Y. Huang, B. Giledereli, A. Köksal, A. Özgür, E. Ozkirimli, Balancing Methods for Multi-label Text Classification with Long-Tailed Class Distribution, Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing [arXiv:2109.04712](#).
- 675 [24] T. Wu, Q. Huang, Z. Liu, Y. Wang, D. Lin, Distribution-Balanced Loss for Multi-label Classification in Long-Tailed Datasets, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelli-

- gence and Lecture Notes in Bioinformatics), Vol. 12349 LNCS, 2020, pp.
680 162–178. [arXiv:2007.09654](#), [doi:10.1007/978-3-030-58548-8_10](#).
- [25] J. Dodge, G. Ilharco, R. Schwartz, A. Farhadi, H. Hajishirzi, N. Smith, Fine-Tuning Pretrained Language Models: Weight Initializations, Data Orders, and Early Stopping [arXiv:2002.06305](#).
- [26] D. P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, 3rd
685 International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings (2014) 1–15 [arXiv:1412.6980](#).
- [27] G. Tsoumakas, I. Katakis, I. Vlahavas, Mining Multi-label Data, in: Data Mining and Knowledge Discovery Handbook, Springer US, Boston, MA, 2009, pp. 667–685. [doi:10.1007/978-0-387-09823-4_34](#).