MetaboListem and TABoLiSTM: Two Deep Learning Algorithms for Metabolite Named Entity Recognition

Cheng S. Yeung¹, Tim Beck^{2,3,a}, and Joram M. Posma^{1,3,a}

¹Section of Bioinformatics, Division of Systems Medicine, Department of Metabolism, Digestion and Reproduction, Imperial College London, SW7 2AZ, United Kingdom

²Department of Genetics and Genome Biology, University of Leicester, LE1 7RH, United Kingdom

³Health Data Research (HDR) UK, United Kingdom

^aThese authors contributed equally. ✉

Abstract

Reviewing the metabolomics literature is becoming increasingly difficult because of the rapid expansion of relevant journal literature. Text-mining technologies are therefore needed to facilitate more efficient literature review. Here we contribute a standardised corpus of full-text publications from metabolomics studies and describe the development of two new metabolite named entity recognition (NER) methods. We introduce two deep learning methods for metabolite NER based on Bidirectional Long Short-Term Memory (BiLSTM) networks incorporating different transfer learning techniques. Our first model (MetaboListem) follows prior methodology using GloVe word embeddings. Our second model exploits BERT and BioBERT for embedding and is named TABoLiSTM (Transformer-Affixed BiLSTM). The methods are trained on a novel corpus annotated using rule-based methods, and evaluated on manually annotated metabolomics articles. MetaboListem (F1 score 0.890, precision 0.892, recall 0.888) and TABoLiSTM (BioBERT version: F1 score 0.909, precision 0.926, recall 0.893) have achieved state-of-the-art performance on metabolite NER. A corpus with >1,200 full-text Open Access metabolomics publications and >116,000 annotated metabolites was created. This work demonstrates that deep learning algorithms are capable of identifying metabolite names accurately and efficiently in text. The proposed corpus and NER algorithms can be used for metabolomics text-mining tasks such as information retrieval, document classification and literature-based discovery.

Availability: The corpus and NER algorithms are freely available with detailed instructions from Github at <https://github.com/omicsNLP/MetaboliteNER>.

biomedical literature corpus | deep learning | metabolomics | named entity recognition | natural language processing | text mining

Correspondence: timbeck@leicester.ac.uk and jmp111@ic.ac.uk

Introduction

Since the late 1990s, metabolomics has contributed to better the understanding of the roles that small molecules play in health and disease (1). The on-going developments of high-

throughput data acquisition ensures that the technology is accessible to more researchers (2), evident by the increasing number of studies year-on-year for studies in any organism as well as those only done in humans (Figure 1). With the growth of the number of publications comes the difficulty of literature review, and performing a literature review of all papers relating to a specific phenotype has become an almost impossibly difficult task for researchers (3). To facilitate more efficient review processes, computational literature mining tools are needed.

Natural Language Processing (NLP) is a branch of computer science that aims to understand natural human language computationally; in the biomedical field, NLP often refers to the technologies that allow computational tools to interpret scientific texts written by humans. In recent years, NLP has been successfully applied in many biomedical text mining tasks, including information retrieval, document classification and literature-based discovery (4–7). By identifying biologically important entities in articles and computationally inferring the connection between them, NLP facilitates understanding of biological relationships from textual information. Such efficient and automatic knowledge mining systems can provide a comprehensive viewpoint to researchers by analysing a vast amount of articles.

Some advances in metabolomics have been made for using computer algorithms to read scientific text (8–11), however these suffer from a few drawbacks such as using rule-based annotation methods and dictionary matching. While these methods can indeed correctly capture a vast amount of metabolite names, they are dependent on the library of compound names used - and no library is exhaustive. The latest version of the Human Metabolome Database (HMDB) (12) reports over 18,557 metabolites have been detected and quantified in the human body, and 91,822 more are expected or predicted. Moreover, these numbers are still increasing, from 2,180 (HMDB 1.0 in 2007) to 114,100 (HMDB 4.0 in 2017) in a decade (13) and predicted to be an order of magnitude higher (1). Hence there is a need for generalisable approaches for automated metabolite recognition.

Within NLP, Named Entity Recognition (NER) is a field concerned with automatically detecting specific entities from texts. To facilitate biomedical text-mining, and NER in

Fig. 1. PubMed search results for metabolomics studies (2000-2021), performed on 7 October 2021. PubMed was searched using the following terms in the title and/or abstract: metabolomics, metabonomics, metabolome-wide association, metabolic phenotyping, metabolic fingerprinting, metabolomic, metabonomic, metabolome, metabonome, metabolic profiling, MWAS. The number of projected metabolomics articles in 2021 was calculated as the number on 7 October 2021 (day 278) multiplied by 365/278 to obtain an estimate at the end of the year. (A) All metabolomics studies. (B) Metadata mining by PubMed, such as the studies species (e.g. human), is updated as data becomes available which can be up to 2 years. Here, we estimated the proportion of human studies for 2020 and 2021 based on a robust second-order polynomial regression on the data from 2000-2019 from the total. From 2006-2019 this proportion has been relatively stable between 39-46%.

particular, BioCreative (Critical Assessment of Information Extraction in Biology) (14) have proposed multiple tasks/challenges and provided corpora for chemical entity recognition, such as the CHEMDNER (chemical compound and drug name recognition, BioCreative V4) (15) and CEMP (Chemical Entity Mention in Patents, BioCreative V5) (16) challenges. The use of machine learning techniques such as Conditional Random Fields (CRF) has shown promising results with F1-scores >0.87 of tmChem (17) in the V4 challenge, however the CRF methods were surpassed by a deep learning (DL) system using Bidirectional Long Short-Term Memory (BiLSTM) recurrent neural networks called ChemListem (18) in the V5 challenge (F1-score of 0.90).

Adding a word embedding layer to NER DL models is well known to improve performance (19), and ChemListem used context-free word embeddings/representation using global vectors (GloVe) (20) for this purpose. However, in recent years big improvements have been made in NLP by using contextual embedding with transformers such as BERT (Bidirectional Encoder Representations from Transformers) (21). BERT was trained on a corpus containing 3.3 billion words mainly from the English Wikipedia and BooksCorpus. Unlike the context-free embedding techniques, BERT's contextual embedding makes the vector representation of a word dependent on its context. Therefore, for biomedical text mining the BioBERT model (22) was pretrained on PubMed abstracts and PubMed Central (PMC) articles (23), and has been used to obtain F1-scores on the BioCreative V4 and V5 of 0.93 and 0.94, respectively, and F1-scores of 0.78-0.85 for recognising genes and proteins (24).

There are three main issues with using these methods for metabolite NER. First, as mentioned, with these methods the context matters and therefore training algorithms to recognise chemical entities in abstracts does not necessarily mean they work as well for full-text articles. Second, while a metabolite is a chemical entity, a chemical entity is not necessarily a metabolite, therefore these existing algorithms may result in many false positives when using them in metabolomics.

Last, each metabolite can have multiple (sometimes ten or even a hundred) different synonyms - over 1.2 million synonyms are reported in HMDB - adding to the complexity of using a dictionary of synonyms as new names can be used by different authors. Hence, new algorithms are needed that focus specifically on recognising metabolites as well as not constrain these algorithm to abstracts, but on full-text paragraphs. Here, we describe the development of a standardised, machine-readable metabolomics corpus of full-text Open Access (OA) PMC articles by using the Auto-CORPUS (Automated and Consistent Outputs from Research Publications) package (25) for text standardisation in conjunction with semi-automatic annotation of metabolites in full texts using a combination of dictionary searching (using HMDB to stay in line with prior work (10)), regular expression matching and rule-based approaches. We then use this corpus to train two DL-based algorithms to perform metabolite NER with the aim to obtain a generalisable model that can be used to speed up metabolomics literature review.

Materials and Methods

Dataset. The data used in the study consists of Open Access metabolomics publications from PMC (n=1,218). The metabolomics corpus is constituted of human-related articles in 18 categories; 8 of which are selected by traits and 10 are by journals (see Table 2). The following search terms were used for filtering the metabolomics corpus:

(metabolomics[Abstract] OR metabolome[Abstract] OR metabolomic[Abstract] OR metabolome-wide[Abstract] OR metabonomics[Abstract] OR metabonome[Abstract] OR metabonomic[Abstract] OR metabonome-wide[Abstract] OR lipidomics[Abstract] OR lipidome[Abstract] OR lipidomic[Abstract] OR lipidome-wide[Abstract] OR metabolites[Abstract] OR metabolite[Abstract] OR metabolic profiling[Abstract] OR metabolic phenotyping[Abstract] OR metabolomics[Title] OR metabolome[Title] OR metabolomic[Title] OR metabolome-wide[Title] OR metabonomics[Title] OR metabonome[Title] OR metabonomic[Title] OR metabonome-wide[Title] OR lipidomics[Title] OR lipidome[Title] OR lipidomic[Title] OR lipidome-wide[Title] OR metabolites[Title] OR metabolite[Title] OR metabolic profiling[Title] OR metabolic phenotyping[Title]) AND (urine[Abstract] OR urinary[Abstract] OR blood[Abstract] OR serum[Abstract] OR plasma[Abstract] OR faecal[Abstract] OR faeces[Abstract] OR cerebrospinal fluid[Abstract] OR CSF[Abstract] OR biofluid[Abstract] OR stool[Abstract] OR feces[Abstract] OR fecal[Abstract] OR stool[Title] OR feces[Title] OR fecal[Title] OR urine[Title] OR urinary[Title] OR blood[Title] OR serum[Title] OR plasma[Title] OR faecal[Title] OR faeces[Title] OR cerebrospinal fluid[Title] OR

CSF[Title] OR biofluid[Title]) AND (nmr spectroscopy[Abstract] OR nuclear magnetic resonance[Abstract] OR NMR[Abstract] OR mass spectrometry[Abstract] OR LCMS[Abstract] OR GCMS[Abstract] OR UPLCMS[Abstract] OR LC-MS[Abstract] OR GC-MS[Abstract] OR UPLC-MS[Abstract] OR CEMS[Abstract] OR CE-MS[Abstract] OR nmr spectroscopy[Title] OR nuclear magnetic resonance[Title] OR NMR[Title] OR mass spectrometry[Title] OR LCMS[Title] OR GCMS[Title] OR UPLCMS[Title] OR LC-MS[Title] OR GC-MS[Title] OR UPLC-MS[Title] OR CEMS[Title] OR CE-MS[Title]) AND (human[Abstract] OR patients[Abstract] OR subjects[Abstract] OR participants[Abstract] OR human[Title] OR patients[Title] OR subjects[Title] OR participants[Title]) NOT (mouse[Abstract] OR mice[Abstract] OR rat[Abstract] OR rats[Abstract] OR dog[Abstract] OR dogs[Abstract] OR animal[Abstract] OR cell culture[Abstract] OR dose[Abstract] OR review[Title] OR proteomics[Abstract] OR diet[Abstract] OR proteomic[Abstract] OR proteome[Abstract] OR transcriptomics[Abstract] OR transcriptomic[Abstract] OR transcriptome[Abstract] OR NMR[Title] OR mice[Title] OR rat[Title] OR rats[Title] OR dog[Title] OR dogs[Title] OR animal[Title] OR cell culture[Title] OR dose[Title] OR proteomics[Title] OR diet[Title] OR proteomic[Title] OR proteome[Title] OR transcriptomics[Title] OR transcriptomic[Title] OR transcriptome[Title]) AND ("01/01/2000"[Publication Date] : "3000"[Publication Date])

The publications were extracted on the following dates: cancer (27 Nov, 2020), gastrointestinal (15 Apr, 2020), liver disease (4 Jun, 2021), metabolic syndrome (27 Nov, 2020), neurodegenerative, psychiatric, and brain illnesses (27 Nov, 2020), respiratory (16 Apr, 2020), sepsis (8 Apr, 2020), smoking (4 Jun, 2021); *Analytica Chimica Acta* (14 Apr, 2020), *Analytical and Bioanalytical Chemistry* (15 Apr, 2020), *Analytical Chemistry* (22 Apr, 2020), *Journal of Chromatography A* (14 Apr, 2020), *Journal of Proteome Research* (8 Apr, 2020), *Metabolites* (11 Sep, 2020), *Metabolomics* (15 Apr, 2020), *PLOS One* (27 Nov, 2020), *Proceeding of the National Academy of Sciences of the United States of America* (14 Apr, 2020), *Scientific Reports* (15 Apr, 2020).

The search terms above are formatted in segments, and the segments were to filter PMC articles by: types of metabolomics studies, type of sample (biofluid), technology/assay, human study, avoiding non-human and non-metabolomics studies and publication date respectively. After we filtered for metabolomics articles, a list of search terms was used to search by phenotypes/journals. Journals were filtered by adding a term such as '(Metabolites[Journal])', whereas the phenotypes were filtered by adding the following:

Neurodegenerative, psychiatric, and brain illnesses: (Alzheimer[Abstract] OR Dementia[Abstract] OR Creutzfeldt-Jakob[Abstract] OR Stroke[Abstract] OR Multiple Sclerosis[Abstract] OR Guillain-Barre[Abstract] OR epilepsy[Abstract] OR Parkinson[Abstract] OR migraine[Abstract] OR headache[Abstract] OR Schizophrenia[Abstract] OR Huntington[Abstract] OR autism[Abstract] OR depression[Abstract] OR Alzheimer[Title] OR Dementia[Title] OR Creutzfeldt-Jakob[Title] OR Stroke[Title] OR Multiple Sclerosis[Title] OR Guillain-Barre[Title] OR epilepsy[Title] OR Parkinson[Title] OR migraine[Title] OR headache[Title] OR Schizophrenia[Title] OR Huntington[Title] OR autism[Title] OR depression[Title])

Metabolic syndrome: (obesity[Abstract] OR adiposity[Abstract] OR BMI[Abstract] OR fat mass[Abstract] OR obese[Abstract] OR hypertension[Abstract] OR blood pressure[Abstract] OR diabetes[Abstract] OR CVD[Abstract] OR lipidemia[Abstract] OR waist circumference[Abstract] OR insulin resistance[Abstract] OR fasting glucose[Abstract] OR obesity[Title] OR adiposity[Title] OR BMI[Title] OR fat mass[Title] OR obese[Title] OR hypertension[Title] OR blood pressure[Title] OR diabetes[Title] OR CVD[Title] OR lipidemia[Title] OR waist circumference[Title] OR insulin resistance[Title] OR fasting glucose[Title])

Sepsis: (Sepsis[Abstract] OR Septic shock[Abstract] OR Septic[Abstract] OR Systemic inflammatory response syndrome[Abstract] OR Multiple organ dysfunction syndrome[Abstract] OR Acute pancreatitis[Abstract] OR toxic shock syndrome[Abstract] OR atypical Kawasaki disease[Abstract] OR Sepsis[Title] OR Septic shock[Title] OR Septic[Title] OR Systemic inflammatory response syndrome[Title] OR Multiple organ dysfunction syndrome[Title] OR Acute pancreatitis[Title] OR toxic shock syndrome[Title] OR atypical Kawasaki disease[Title])

Cancer: (cancer[Abstract] OR tumor[Abstract] OR tumour[Abstract] OR cancerous[Abstract] OR carcinogen[Abstract] OR carcinogenic[Abstract] OR carcinoma[Abstract] OR leukaemia[Abstract] OR leukemia[Abstract] OR leukaemic[Abstract] OR leukemic[Abstract] OR lymphoma[Abstract] OR malignancy[Abstract] OR premalignancy[Abstract] OR pre-malignancy[Abstract] OR melanoma[Abstract] OR metastasis[Abstract] OR sarcoma[Abstract] OR adjuvant[Abstract] OR neoadjuvant[Abstract] OR chemotherapy[Abstract] OR

chemo therapy[Abstract] OR chemo-therapy[Abstract] OR malignant[Abstract] OR premalignant[Abstract] OR pre-malignant[Abstract] OR precancerous[Abstract] OR pre-cancerous[Abstract] OR adenocarcinoma[Abstract] OR metastatic[Abstract] OR cancer[Title] OR tumor[Title] OR tumour[Title] OR cancerous[Title] OR carcinogen[Title] OR carcinogenic[Title] OR carcinoma[Title] OR leukaemia[Title] OR leukemia[Title] OR leukaemic[Title] OR leukemic[Title] OR lymphoma[Title] OR malignancy[Title] OR premalignancy[Title] OR pre-malignancy[Title] OR melanoma[Title] OR metastasis[Title] OR sarcoma[Title] OR adjuvant[Title] OR neoadjuvant[Title] OR chemotherapy[Title] OR chemo therapy[Title] OR chemo-therapy[Title] OR malignant[Title] OR premalignant[Title] OR pre-malignant[Title] OR precancerous[Title] OR pre-cancerous[Title] OR adenocarcinoma[Title] OR metastatic[Title])

Gastrointestinal: (IBS[Abstract] OR IBD[Abstract] OR Crohn[Abstract] OR ulcerative colitis[Abstract] OR inflammatory bowel disease[Abstract] OR irritable bowel syndrome[Abstract] OR microscopic colitis[Abstract] OR collagenous colitis[Abstract] OR lymphocytic colitis[Abstract] OR bile acid malabsorption[Abstract] OR celiac[Abstract] OR diversion colitis[Abstract] OR Behçet[Abstract] OR Behcet[Abstract] OR indeterminate colitis[Abstract] OR IBS[Title] OR IBD[Title] OR Crohn[Title] OR ulcerative colitis[Title] OR inflammatory bowel disease[Title] OR irritable bowel syndrome[Title] OR microscopic colitis[Title] OR collagenous colitis[Title] OR lymphocytic colitis[Title] OR bile acid malabsorption[Title] OR celiac[Title] OR diversion colitis[Title] OR Behçet[Title] OR Behcet[Title] OR indeterminate colitis[Title])

Smoking: (smoking[Abstract] OR cigarette[Abstract] OR tobacco[Abstract] OR vaper[Abstract] OR vaping[Abstract] OR smoker[Abstract] OR smokers[Abstract] OR e-cig[Abstract] OR e-cigarette[Abstract] OR e-smoking[Abstract] OR e-smoker[Abstract] OR nicotine patch[Abstract] OR nicotine dependence[Abstract] OR cigar[Abstract] OR pack-year[Abstract] OR pack-years[Abstract] OR pack year[Abstract] OR packyear[Abstract] OR nicotine use[Abstract] OR cannabis use[Abstract] OR smoking[Abstract] OR cigarette[Abstract] OR tobacco[Abstract] OR vaper[Abstract] OR vaping[Abstract] OR smoker[Title] OR smokers[Title] OR e-cig[Title] OR e-cigarette[Title] OR e-smoking[Title] OR e-smoker[Title] OR nicotine patch[Title] OR nicotine dependence[Title] OR cigar[Title] OR pack-year[Title] OR pack-years[Title] OR pack year[Title] OR packyear[Title] OR nicotine use[Title] OR cannabis use[Title])

Respiratory: (pneumonia[Abstract] OR respiratory tract infection[Abstract] OR COVID[Abstract] OR SARS[Abstract] OR MERS[Abstract] OR respiratory syndrome[Abstract] OR coronavirus[Abstract] OR pulmonary fibrosis[Abstract] OR COPD[Abstract] OR Chronic Obstructive Pulmonary Disease[Abstract] OR SARS-CoV-2[Abstract] OR Acute Respiratory Syndrome[Abstract] OR Middle East respiratory syndrome[Abstract] OR severe acute respiratory syndrome[Abstract] OR respiratory virus[Abstract] OR pneumonia[Title] OR respiratory tract infection[Title] OR COVID[Title] OR SARS[Title] OR MERS[Title] OR respiratory syndrome[Title] OR coronavirus[Title] OR pulmonary fibrosis[Title] OR COPD[Title] OR Chronic Obstructive Pulmonary Disease[Title] OR SARS-CoV-2[Title] OR Acute Respiratory Syndrome[Title] OR Middle East respiratory syndrome[Title] OR severe acute respiratory syndrome[Title] OR respiratory virus[Title])

Liver disease: (hepatitis[Abstract] OR haemochromatosis[Abstract] OR hemochromatosis[Abstract] OR cirrhosis[Abstract] OR non-alcoholic fatty liver disease[Abstract] OR NAFLD[Abstract] OR nonalcoholic fatty liver disease[Abstract] OR nonalcoholic steatohepatitis[Abstract] OR non-alcoholic steatohepatitis[Abstract] OR liver fibrosis[Abstract] OR hepatocellular carcinoma[Abstract] OR HCC[Abstract] OR hepatocellular cancer[Abstract] OR acute fatty liver disease[Abstract] OR hepatic encephalopathy [Abstract] OR primary biliary cholangitis[Abstract] OR primary sclerosing cholangitis[Abstract] OR chronic liver disease[Abstract] OR alcoholic liver disease[Abstract] OR ALD[Abstract] OR alcohol-related liver disease[Abstract] OR ARLD[Abstract] OR steatosis[Abstract] OR hepatitis[Title] OR haemochromatosis[Title] OR hemochromatosis[Title] OR cirrhosis[Title] OR non-alcoholic fatty liver disease[Title] OR NAFLD[Title] OR nonalcoholic fatty liver disease[Title] OR non-alcoholic steatohepatitis[Title] OR non-alcoholic steatohepatitis[Title] OR liver fibrosis[Title] OR hepatocellular carcinoma[Title] OR HCC[Title] OR hepatocellular cancer[Title] OR acute fatty liver disease[Title] OR hepatic encephalopathy [Title] OR primary biliary cholangitis[Title] OR primary sclerosing cholangitis[Title] OR chronic liver disease[Title] OR alcoholic liver disease[Title] OR ALD[Title] OR alcohol-related liver disease[Title] OR ARLD[Title] OR steatosis[Title])

The metabolomics articles were stored in HyperText Markup Language (HTML) format, which were then processed by the Auto-CORPus package (25), and standardised into machine-readable JavaScript Object Notation (JSON) documents.

Auto-CORPus outputs three JSON files based on the input of an HTML file of an article: 'maintext', 'table' and 'abbreviation'. In the maintext file, the textual content in the HTML is split into subsections, and each subsections have five attributes: 'section_heading', 'subsection_heading', 'body', 'IAO_term' and 'IAO_ID', where IAO stands for Information Artifact Ontology (26). Here we focused on four section types based on standardised IAO terms, namely the textual abstract ('A'), methods ('M'), results ('R') and discussion ('D') sections. We did not use the introduction sections as

Fig. 2. An example of sentence ID. The ID refers to the second sentence in the second subsection of the method section. Indexes start with 0.

these do not typically contain information pertaining to the study itself or (how) new results (were) acquired).

Rule-based annotation pipeline. The metabolomics corpus was annotated using a semi-automated pipeline to deal with the large volume of texts and prepare it for supervised learning. The pipeline takes the Auto-CORPUS-processed articles as JSON input, and outputs the location (character number) of metabolites in each text (an example can be found in Figure 4 in the Results section). The annotation pipeline is comprised of three steps: pre-processing, identifying metabolites and locating recognised entities.

Pre-processing. (Sub)sections are split into sentences using sentence tokenisation with the BioBERT (22) model (version biobert-base-cased-v1.2) using the Huggingface Transformer package (27). Each sentence is uniquely identifiable by two attributes: the PMC identifier (PMCID) of the article that the sentence belongs to, and a sentence identifier. A sentence identifier is a string that contains a letter and five digits (Figure 2) where the letter indicates the relevant sections types (textual abstract (A), methods (M), results (R) and discussion (D)). These section types are determined by the IAO_term key in the Auto-CORPUS full-text JSON file.

Metabolite identification. The next step is to identify any metabolite mentioned in each tokenised sentence using both dictionary searching and regular expression matching. The dictionary used here was based on the list of metabolites (and their synonyms) that are classified in HMDB (12) as ‘quantified’ or ‘detected’. A list of all metabolites and synonyms was first downloaded on 25 June 2020, and subsequently refreshed on 16 August 2021, after which it was filtered for the ‘quantified’ or ‘detected’ metabolites. Short abbreviations (up to 5 characters) that are not exclusive for metabolites were removed (e.g., ‘PC’ can refer to the metabolite ‘phosphatidylcholine’, but is also a common abbreviation for ‘principal component’ and ‘pancreatic cancer’), and it was then further cleaned of other non-specific words such as ‘result’ (a synonym of ‘omeprazole’).

To complement dictionary searching, the regular expressions are designed to provide partial matches for the metabolite entities (Table 1). It is expected that more metabolites can be captured in this manner compared to complete matching. Entities with a partial match are then expanded to include the entire word and included a full metabolite name or part of a metabolite name; the latter scenario is treated further in post-processing. To craft the regular expressions, the HMDB dictionary was used as a reference. From the dictionary, the entities with one or more spaces and the ones with symbols

like hyphen (‘-’) and colon (‘:’) were selected to form a ‘reference set’. The reference set meant to include metabolite names that are less trivial (and thereby more regular). Trivial terms such as singletons are usually more difficult to be captured with regular expressions without drastically increasing the number of false positives by including common words. The regular expression approach is used to complement the HMDB dictionary, therefore trivial terms are omitted from the reference set as they are assumed to be listed in HMDB already. This was used to create a set of regular expressions such that most, if not all, of the reference set can be (at least partially) matched by at least one regular pattern in the set, while avoiding matching unwanted terms.

The regular expressions were created in a recursive manner: while not all terms in the reference set are matched and there remain some observable patterns in the remaining terms, either a new expression was created or an existing expression augmented in order to match the remaining regular terms. Examples of regular expressions and the terms that they are expected to match are shown in Table 1.

Post-processing. Post-processing of the matches was utilised to integrate the result of dictionary searching and regular expression matching, including those from partial matching. This step is comprised of three components: combining adjacent words, balancing the number of brackets, and merging overlapped entities.

For each metabolite entity that has been recognised its adjacent words are evaluated. The term before the entity is prepended to the annotation only if certain rules are satisfied: the preceding word ends with a hyphen (‘-’), ends with a number, the recognised entity itself starts with a comma and number (e.g. ‘,3’) or the current word starts with a hyphen. If the preceding word is a stop word, or if none of the rules can be met, then the procedure is complete and no term is prepended to the entity. The rules for the next word are the same as described above, except now the entity becomes the ‘preceding’ word in the above scenario and one more rule is evaluated (whether the next word is one of either ‘acid’, ‘isomer’, ‘ester’ or ‘ether’). This process is recursive and is done until no adjacent words satisfy any of the rules.

For all entities the number of brackets is evaluated to ensure all parentheses, square brackets and curly brackets have a starting and ending bracket in the entity. If an opening or closing bracket is missing then the entity is enclosed by the corresponding missing bracket.

All entities that overlap or separated only by a (white)space character are merged into a single entity. The entire post-processing step is executed recursively, i.e. until there is no more change in the list of identified metabolite entities. Once this has finished, two files are created that are used to train the DL models. The first file contains all sentences that have at least one recognised metabolite, along with their sentence identifiers (see Figure 2) and the PMCID of the articles the sentences originate from. The second file is in table with 5 fields: PMCID, sentence identifier, position of the start character of the entity, position of the end character of the entity, and the entity itself. These two files are exemplified in the

Regular expression	Example metabolite
$(\backslash(?!)\w+(oxy yl ol ine ene ose ane ate ester isomer ium ino ide one oxo en estra))+(-\w+)\w+(?!)(ic)?(=[\w+\w+\w+])$	<u>3-Methoxy-4-hydroxy-phenyl</u> -acetate
$([a-z] -[0-9])*([0-9]+([a-z] \alpha \beta \sigma \omega \tau \epsilon \zeta [\alpha-\omega])^*[\w+,\w+\w+\w+])+(([a-z] -[0-9])^+)+$	<u>5alpha-Androstan-3alpha,17beta</u> -diol disulfate
$(([a-z]?[0-9]+(:)[0-9]+)/[a-z]?[0-9]+(:)[0-9]+)^*$	d18:1/16:0
$\backslash d+[a-z]?(\backslash d+[a-z]?)+\backslash)?-$	(1S,10R,11S,15S)-5-hydroxy-4-methoxy-15-methyltetracycloheptadecan-2(7),3,5-trien-14-one
$([1-9])\backslash ((\alpha \beta \sigma \omega \tau \epsilon \zeta [\alpha-\omega]) Z E L R S -\w+)?\w+)-$	(11E)-Octadecenoic acid
$(oline dino etin olin ecine idine noid oate gmine [a-z]dine none igine zole pine nime anide acene prost sone idone hane cine cin tril erol [a-wy-z]amine zime [a-z]toxin(e?) otin(e?) ione oside zine olone xone [a-z]dene lene lasin(e?) atin(e?) ffin(e?) [a-np-z]yline xin(e?) zin(e?) bene stine ythin(e?) astin(e?) xitin(e?))(s?)$$	Ribofuranosylcreatine

Table 1. Examples of regular expressions used in the rule-based annotation pipeline. Example metabolites that each regular expression matches are shown on the right; the matching parts of the entities are underlined.

middle panel of Figure 4.

Training, validation and test set generation and evaluation.

The 1,218 metabolomics publications were split at random into a training, validation and test set a 75:10:15 ratio, corresponding to 913 training publications, 122 publications for validation and 183 publications for the test set. The publications in the training and validation set were processed using the rule-based annotation pipeline described above, creating a semi-automatically annotated corpus. The publications in the test set, this included entities in the full text as well as in tables, were manually annotated to avoid bias and to provide a ground truth to evaluate the algorithms on.

Although metabolite abbreviations and entities in tables were not deliberately annotated in the training set with the annotation pipeline, they were manually annotated in the test set for completeness; these types of metabolite names were considered separately while employing models and evaluating performances (detailed in Post-processing). The performance of algorithms was assessed by evaluating the number of true positives (TP), false positives (FP) and false negatives (FN) and calculating the precision ($\frac{TP}{TP+FP}$), recall ($\frac{TP}{TP+FN}$), F1-score ($\frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$) and F* measure (28) ($\frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall} - \text{precision} \times \text{recall}}$) on the test set.

Metabolite NER using LSTM. A BiLSTM model architecture that was used previously on chemical NER task achieving F1-scores >0.9 (18) was used here for metabolite NER. We first evaluated the original ChemListem model on the corpus, we then developed two new methods (MetaboListem, TABoLiSTM) based on the prior methodology to improve the performance and described below. Our new models were trained in 50 epochs, and the epoch that achieved the highest F1-score on the validation set was used for evaluation of the

test set. All F1-scores reported here are for the test set only.

Pre-processing. Following prior work (18), a tokeniser designed for chemical text-mining (OSCAR4) (29) was employed to tokenise words in the metabolomics corpus for MetaboListem. For the TABoLiSTM model, we used BioBERT as tokeniser (27) because TABoLiSTM uses the BioBERT token embeddings.

For both word tokenisation systems, the BIOES tag scheme was employed to label token sequences. BIOES tag scheme is a standard sequence labeling technique that classifies tokens into five categories: ‘B’ labels the tokens at the beginning of entities, ‘I’ labels internal tokens of entities, ‘E’ labels the ending tokens of entities, ‘S’ labels singletons (i.e. tokens that are a complete entity on their own), and ‘O’ the tokens that do not belong to an entity.

ChemListem used a pre-classifier random forest to predict the probability of a token being (part of) a chemical entity (‘B’, ‘I’, ‘E’ or ‘S’) based on a set of ChEBI-derived (30) chemicals and chemical elements. For MetaboListem and TABoLiSTM we implemented the same pre-classifier approach using the set of metabolites from HMDB (12) rather than chemicals from ChEBI, where tokens were segmented using OSCAR4 and BioBERT for MetaboListem and TABoLiSTM, respectively. The pre-classifier produces name-internal features that were used in the model training.

ChemListem and MetaboListem. For ChemListem (18), name-internal features for each token was computed in the pre-processing stage using the pre-classifier and treated as one of the two inputs to the model (Figure 3). The features go through a convolutional layer with width of 3 and are concatenated with the numeric representation of words (i.e. a unique number assigned to each word) in the GloVe (20) vector space. Each token now has 256+300=556 dimensions,

Fig. 3. The architecture of the ChemListem network. The input tensors 'wordn' and 'ni' correspond to the numerical representations of word (token) sequences and the name-internal features of word (token) sequence determined by the pre-classifier, respectively.

where 256 is the output neuron number of the convolutional layer, and 300 the GloVe vector representation. The concatenated tensor is then fed into the BiLSTM layer, before outputting a 5-dimensional result (number of tags used in BIOES tagging) via a TimeDistributed Dense neuron.

Our MetaboListem algorithm is based on the same neural network architecture (and GloVe embedding) as ChemListem, hence we named the algorithm in the same manner. It differs from ChemListem only in the training data (metabolomics corpus opposed to BioCreative V5 (16)), pre-classifier data (HMDB instead of ChEBI) and TensorFlow implementation (version 2.0 versus 1.3).

TABoLiSTM (Transformer-Affixed Bidirectional LSTM). Our TABoLiSTM model differs from MetaboListem mainly in the use of pre-trained BERT (21) (*bert-base-cased*) and BioBERT (22) (*biobert-base-cased-v1.2*) embedding. The name-internal features and the numeric presentation of words were all re-computed using (Bio)BERT tokens instead of words and the last hidden layer of the (Bio)BERT model replaces the GloVeEmbeddingLayer in Figure 3. The (Bio)BERT layer provides contextualised word embeddings, and the embeddings were concatenated with the convolution output of name-internal features as previously, except for larger tensors of $256+768=1214$ dimensions. A Spatial-Dropout1D layer with dropout rate 0.25 is applied on top of the (Bio)BERT layer to prevent overfitting and improve model performance by embedding the dropout (31).

Post-processing. Post-processing of our DL models involves the same procedure as that of the annotation pipeline (section) to fix pathological terms such as incomplete entities: to check the neighbouring entities, to fix unbalanced parentheses if possible, and to merge overlapped positions. At this stage further rule-based approaches can be employed to further improve models, for MetaboListem and TABoLiSTM these were not used, but rather represent potential further improvements that can be made to the post-processing.

The DL algorithms were trained on sentences with full metabolite names only, and were not trained on abbreviations

or data in tables. In the post-processing step, the algorithms were applied to the Auto-CORPus JSON output of the list of abbreviations (abbreviations and definitions) and any recognised entity (definition) was then linked to its abbreviation and all occurrences in the full text marked as a recognised entity. This was only done when the entire definition was recognised as metabolite entity, i.e. the rules to fix incomplete entities are not applied here. A direct approach to replace all abbreviations in the text with their definitions was investigated but discarded since the current implementation of the abbreviation recognition does not detect all types of abbreviations. The same approach was undertaken on the Auto-CORPus JSON output of all tabular data, where each cell was evaluated by the DL models as a sentence. Any recognised entity was annotated and evaluated against the manually annotated entities in the test set.

Results

Metabolomics corpus. The corpus generated by means of the annotation pipeline (see Materials and Methods) is of similar size to that of the BioCreative V5 dataset. We extracted 54,493 sentences that contain metabolites from the 1,218 metabolomics articles, and in these have identified a total of 116,393 metabolites. For comparison, the BioCreative V5 dataset contains 21,000 abstracts and 99,632 chemical entities. Table 2 details the metabolomics corpus and number of annotated unique entities illustrating that most metabolomics articles report between 20 and 30 unique metabolites. While metabolomics articles on cancer are the largest single group within the corpus, the average number of metabolites reported in cancer articles is similar to that of other disease areas.

Model performance. We first evaluated the application of the rule-based annotation pipeline on the 183 manually annotated test set articles. Our pipeline yields an F1-score of 0.8893, with precision and recall of 0.8850 and 0.8937 respectively (see Table 3). Next, we evaluated the performance of ChemListem (18), trained on the BioCreative V5 dataset, on the metabolomics corpus where it achieves an F1-score of 0.7669. Our first DL model uses the same architecture as ChemListem but was trained on the metabolomics corpus, therefore this model was named MetaboListem, and it achieved an F1-score of 0.8900 (precision of 0.8923 and recall of 0.8877). Our second DL model uses a different embedding layer (BERT or BioBERT) to MetaboListem, but was trained on the same dataset. The TABoLiSTM (Transformed-Affixed Bidirectional LSTM) model with BERT embeddings achieved an F1-score of 0.9004, with precision and recall rates of 0.9187 and 0.8829 respectively; the same model with BioBERT embeddings results in a 0.85% improvement of the F1-score.

The workflow including the structure of the metabolomics corpus, the annotation and result of the TABoLiSTM application on the abstract of a publication is shown in Figure 4. We exemplify the differences in the output from the TABoLiSTM model and our annotation pipeline using two snippets of text

Fig. 4. Example of model employment on the abstract of PMC6316856. The Auto-CORPus output is illustrated at the top, followed by the manual annotations (formatted in the same way as rule-based annotations) in the middle and the TABoLiSTM prediction result (bottom). The prediction results were visualised as an annotated HTML file. In this example all entities are true positives and are highlighted in green.

Category	# of metabolites	# of articles	Average # of (unique) metabolites per article
Journal			
<i>Analytical and Bioanalytical Chemistry</i>	716	34	21.1
<i>Analytical Chemistry</i>	863	48	18.0
<i>Journal of Chromatography A</i>	220	9	24.4
<i>Journal of Proteome Research</i>	694	24	28.9
<i>Metabolites</i>	1279	50	25.6
<i>Metabolomics</i>	1267	43	29.5
<i>PLOS One</i>	5400	175	30.9
<i>PNAS USA</i>	429	14	30.6
<i>Scientific Reports</i>	4278	126	34.0
Trait			
cancer	12662	492	25.7
gastrointestinal	876	37	23.7
liver disease	1959	121	16.2
metabolic syndrome	7081	286	24.8
neurodegenerative, psychiatric, and brain illnesses	2546	113	22.5
respiratory diseases	398	37	10.8
sepsis	552	22	25.1
smoking	1360	124	11.0

Table 2. Summary of number of unique metabolites mentioned per category. Each metabolite is only counted once per category, even when it is reported multiple times in different articles of the category.

Model	Training data	Embedding	F1-score	Precision	Recall	F*-score	Model size
Annotation pipeline			0.8893	0.8850	0.8936	0.8006	
ChemListem	CEMP BioCreative V.5	GloVe	0.7669	0.7301	0.8075	0.6219	26MB
MetaboListem	metabolomics corpus	GloVe	0.8900	0.8923	0.8877	0.8018	26MB
TABoLiSTM	metabolomics corpus	BERT	0.9004	0.9187	0.8829	0.8189	827MB
TABoLiSTM	metabolomics corpus	BioBERT	0.9089	0.9255	0.8928	0.8329	827MB

Table 3. Summary of model performances. Models are evaluated in terms of F1-score, precision and recall rate, and F*-score using the metabolomics corpus test set (183 manually annotated full-text articles). The best results for each metric is highlighted in bold.

in Figure 5 that illustrate the typical examples of how the DL model predictions differ from the rule-based annotations. These examples include spelling errors, metabolites not included in the dictionary and metabolite abbreviations.

Beyond prediction accuracy, another aspect to evaluate is the model sizes. Despite the performance advantages, TABoLiSTM is much larger than both ChemListem and MetaboListem; the sizes of the latter two models are around 26MB, whereas the size of the TABoLiSTM weight file exceeds 412MB. Furthermore, pretrained BERT/BioBERT (cased versions) are required to employ the TABoLiSTM model, which adds an extra 415MB.

The performance of each of the 5 models was evaluated for separate sections (abstract, materials/methods, results, results&discussion, discussion) in the test set (Figure 6). Across all models, the best performance was achieved on separate results and discussion sections, with lowest performance for material/methods and results&discussion sections.

Automated literature review of the metabolomics corpus. We applied TABoLiSTM on the abstract, results (including table data) and discussion sections of each of the 8 traits (see Table 2) in the metabolomics corpus to summarise the application of the algorithm in literature review. Figure 7 visualises the results for the 492 cancer articles in the corpus,

(a)
Results: Established analytical method based on urine metabolomics GC-MS of Type II diabetes, determine the urine succinic acid, L-leucine, L-isoleucine, tyrosine, slanine, acetoace acid, mannose, L-isoleucine, L-threonine, Phenylalanine, fructose, D-glucose, palmi acid, oleic acid and arachidonic acid were significantly were significantly changed.

(b)
In line with our expectations, the strongest positive associations with diabetes are observed for the numerous sugar metabolites that were observed on the three platforms. Concentrations of glucose, mannose, Desoxyhexose (primarily Deoxyglucose), uronic acid (primarily glucuronic acid), dihexose (primarily maltose), and several products from the biosynthesis or the degradation of glycosylated proteins or glycolipids (H3-HNAC2-NANA, HNAC, HNAC-H2-dH) are all found increased by up to 90% in the diabetes group ($p = 1.3 \times 10^{-4}$ to 2.5×10^{-9}). Furthermore, there is a significant decrease of average 5-anhydroglucitol concentrations by 37.8% ($p = 5.1 \times 10^{-6}$) in participants with diabetes when compared to the control group (Figure 2).

Fig. 5. Examples of TABoLiSTM and rule-based annotation results. The green entities are recognised by both methods, the yellow entities are only recognised by TABoLiSTM, and the red only by rule-based annotation. (a) is excerpted from PMC4969426 and exemplifies typos; (b) is excerpted from PMC2978704.

Fig. 6. Model performance per publication section. The model performance (precision, recall and F1-score) are given for 5 sections (A, abstract; M, materials/methods; R, results; RD, results and discussion (as single section); D, discussion) for the 5 models.

where only metabolites found in at least 3 different articles are visualised. The size of the text of each metabolites is proportional to the number of articles that reported the metabolite and the position is based on how frequent metabolites are found together. In total, 664 metabolites are recognised in at least three of the 492 articles on cancer (compared to 12,662 unique metabolite names across all 492 articles). The five most frequent metabolites amongst these are glucose, glutamine, lactate, alanine and glutamate, which are mentioned in 22-29% of articles (see Table 4).

In the cancer summary graph, no apparent patterns or clusters are observable based on the node colours due to the three-article threshold applied to the graph. This is in contrast to other networks such as one created for the smoking trait where the restriction of 3 articles is not applied and all 1,360 metabolites reported in 124 articles are visualised (see Figure 8). This reflects the low number of common metabolites for this trait in the corpus (Table 4). Glucose is the only metabolite that appears in the top 10 for all traits, with it being reported in 62% of articles on metabolic syndrome.

Discussion

We have described three new DL models, using different word embedding layers, that have both achieved state-of-the-art performance (F1-score > 0.89) for metabolite NER on a manually annotated dataset with OA metabolomics articles. Our DL models (F1-scores of 0.91, 0.90 and 0.89 for TABoLiSTM (BioBERT), TABoLiSTM (BERT) and MetaboListem, respectively) surpass by a considerable margin previous methods (using CRFs) for metabolite NER which achieved maximum F1-scores of 0.78 (8) and 0.87 (9). Our methods, compared to similar methods for chemical entity recognition (17, 18, 24) and metabolite NER (8, 9), have the benefit of having been trained not only on abstracts and titles, but also on full-text paragraphs that are relevant for information retrieval from studies (results and discussion sections). The corpora these algorithms were trained and evaluated on are made available for future re-use and algorithm development, and to the best of our knowledge are the first of its kind for metabolomics. In addition, we developed a semi-automated annotation pipeline combining regular expression and rule-based approaches with dictionary searching, both

for processing the training data for and as an alternative to DL methods, which too achieved outstanding performance (F1-score > 0.88) which also surpassed prior DL algorithms (8, 9).

Annotation pipeline. Recent work that used NLP methods via dictionary searching for text mining in metabolomics highlighted the considerable user effort (1-4hrs as per (10)) required to run these type of analyses. Our annotation pipeline consists of a dictionary of metabolite names and synonyms (from HMDB) which can easily be added to by including data from other databases such as ChemSpider (34) and ChEBI (30). However, any entity added to the dictionary must undergo some form of quality control to ensure it is a metabolite. Here, we have limited ourselves to only including metabolites listed in HMDB that were classified as either ‘quantified’ or ‘detected’ and have not included those that are ‘expected’ or ‘predicted’.

The annotation pipeline also includes regular expression and rule-based methods used to recognise metabolites in text and these add additional advantages: they are time-efficient, easily customisable and explainable. While it requires minutes or even hours for an skilled expert to annotate paragraphs and entire full-text articles, the time to process one full article using these methods is less than 15 seconds on average. Once experts establish additional rules focusing on recognising naming conventions, then these can easily be added into the pipeline. Moreover, for each identified entity it is immediately obvious which rule(s) resulted in it being predicted as a metabolite entity, and therefore more interpretable than a DL model. For this reason, any false positives or false negatives can be analysed and new sets of rules established to correctly annotate the metabolites in text. This also allows for developing a user-feedback pipeline in which new rules can be shared to obtain better results. We have used this semi-automated annotation pipeline to prepare the training corpus for the DL methods, therefore any improvement in the rules for annotation may also improve DL models. Another application of the pipeline is to run it in parallel with the DL-based methods and combine the results as a post-processing step. The main limitation of the annotation system is that the system assumes that the dictionary of metabolite names contains all known metabolites with irregular names. We used

Fig. 7. Graphical representation of metabolites mentioned in the abstract, results and discussion sections of cancer-related metabolomics articles. Each node represents a metabolite and is displayed as the (lower-cased) metabolite name. Two nodes are connected by an edge if the corresponding metabolites appear in the same article at least once. The displayed names are sized proportionally to the number of articles that report the metabolite, and are coloured according to the article that mentions the metabolite most (if such article is not unique, the one with smaller PMCID number is selected). The graph layout is generated using ForceAtlas2 (32) implemented in Gephi (33). Because of the large number of metabolites (see Table 2), only the metabolites that appear in three or more articles are used for the graph construction. Here, edges are hidden for sake of visibility.

the HMDB as dictionary because it is the most comprehensive database of human metabolites thus far, however our annotation approach may still result in missing metabolites. One example is endoxifen, which is an active metabolite of the breast cancer drug tamoxifen. Although its synonym 4-hydroxy-*N*-desmethyltamoxifen can be recognised using regular expressions (for example because of the prefix), the word endoxifen itself is not recognised by the pipeline because of its irregularity and not being included in our dictionary (it exists in HMDB as ‘expected’). Augmenting the dictionary with the metabolites listed as ‘expected’/‘predicted’ may appear to be a feature and obvious solution, however adding all of these terms would generate a dictionary 4 times larger than the current one and therefore drastically increase the runtime for the rule-based annotation. Also, >90% of the ‘expected’ metabolites are lipids (12), and these names are usually recognisable using regular expressions.

Overall, the annotation pipeline is dependent on its data and rules, it can therefore not detect metabolites that have been misspelled, nor will it flag certain terms as false positives if they appear in the dictionary as synonyms for a metabolite. For example, common words or words with multiple meanings (depending on the context) would be recognised by the system such as ‘result’ (synonym of ‘omeprazole’) and ‘retinal’ (both a vitamin-derivative and a medical term for ‘eye’). HMDB also includes chemical entities such as ‘ammonium acetate’ and ‘silica’, which are commonly mentioned in methods sections of metabolomics articles. As the HMDB includes more than 18,000 ‘detected’ and ‘quantified’ metabolites and their synonyms, it was not feasible to examine all entities as pre-processing in this study and only those with 5 or less characters were manually checked; therefore, further cleaning of the dictionary would likely improve the precision of the annotation pipeline. Our results

cancer		gastrointestinal		liver disease		metabolic syndrome	
glucose	142 (29%)	glucose	12 (32%)	glucose	29 (24%)	glucose	178 (62%)
glutamine	121 (25%)	tyrosine	12 (32%)	cholesterol	25 (21%)	cholesterol	108 (38%)
lactate	117 (24%)	lactate	11 (30%)	creatinine	21 (17%)	valine	83 (29%)
alanine	111 (23%)	acetate	10 (27%)	tyrosine	20 (17%)	leucine	73 (26%)
glutamate	109 (22%)	phenylalanine	9 (24%)	glycine	18 (15%)	tyrosine	71 (25%)
tyrosine	105 (21%)	leucine	9 (24%)	phenylalanine	17 (14%)	triglycerides	70 (24%)
glycine	105 (21%)	tryptophan	9 (24%)	lactate	16 (13%)	alanine	67 (23%)
valine	104 (21%)	alanine	8 (22%)	bilirubin	16 (13%)	isoleucine	65 (23%)
choline	103 (21%)	arginine	8 (22%)	pyruvate	15 (12%)	phenylalanine	62 (22%)
creatinine	95 (19%)	citrate	8 (22%)	valine	15 (12%)	glycine	59 (21%)
neurodegenerative, psychiatric, and brain illnesses		respiratory		sepsis		smoking	
tryptophan	31 (27%)	lactate	8 (22%)	phenylalanine	11 (50%)	creatinine	24 (19%)
glucose	26 (23%)	glycine	6 (16%)	glucose	9 (41%)	cotinine	15 (12%)
glycine	25 (22%)	leucine	6 (16%)	lactate	9 (41%)	nicotine	14 (11%)
glutamate	22 (19%)	pyruvate	6 (16%)	arginine	8 (36%)	glucose	14 (11%)
kynurenine	22 (19%)	glucose	6 (16%)	urea	6 (27%)	leucine	11 (9%)
phenylalanine	21 (19%)	lysine	5 (14%)	nitric oxide	6 (27%)	isoleucine	11 (9%)
serotonin	18 (16%)	creatine	5 (14%)	methionine	6 (27%)	valine	11 (9%)
tyrosine	18 (16%)	isoleucine	5 (14%)	pyruvate	6 (27%)	cholesterol	11 (9%)
creatinine	18 (16%)	glutamine	5 (14%)	creatinine	6 (27%)	lactate	10 (8%)
lactate	17 (15%)	acetate	4 (11%)	ATP	5 (23%)	glutamate	9 (7%)

Table 4. Top ten metabolites reported for each trait in the metabolomics corpus. The number of articles that report the metabolite are given including the percentage of the total. The total number of articles in each category can be found in Table 2.

Fig. 8. Graphical representation of metabolites mentioned in the smoking-related metabolomics articles. Each node represents a metabolite and is displayed as the (lower-cased) metabolite name. Two nodes are connected by an edge if the corresponding metabolites appear in the same article at least once. The displayed names are sized proportionally to the number of articles that report the metabolite, and are coloured based on the article that reports the metabolite most (if such article is not unique, the one with lower PMCID number is selected). The graph is generated using ForceAtlas2 (32) implemented in Gephi (33). Only the metabolites in the abstract, result (including those in tables) and discussion sections are plotted here. Unlike Figure 7, all metabolites (including the ones appear in less than three articles) are used in the graph construction. Edges are hidden for sake of visibility.

indicate that the combination of regular expressions, dictionary searching and rules can indeed correctly capture a vast

amount of metabolite names as also reported by Majumder *et al* (10).

Metabolomics corpus. To the best of our knowledge, our metabolomics corpus, which consists of the abstract, method, result and discussion sections of 183 manually annotated and 835 semi-automatically annotated PMC articles, is the first available corpus that is designed for human metabolomics research. While some databases include some references to articles where metabolites are mentioned, the databases do not contain enough information (e.g. annotated texts) for NLP model development. An advantage of our metabolomics corpus, compared to common chemical NER corpora (15, 16), is that it contains not only abstracts but also sentences from the method, result and discussion sections which is relevant for context-dependent NER. While abstracts indeed provide a rich set of context, the structure of sentences are often different from the main text and therefore training and employing contextual NLP models only on abstracts could potentially result in overlooking valuable information while mining the full text of articles.

The limitation of the current metabolomics corpus is that it was designed only to label metabolites that appear in full in sentences; that is, if the recognition of a metabolite requires semantic interpretation then it would not be correctly detected. For example, if the string ‘1- and 3-methylhistidine’ is encountered by the pipeline, then only ‘3-methylhistidine’ is annotated, whereas the metabolite ‘1-methylhistidine’ is also in the string. To overcome this limitation of the corpus, it needs to be processed further using semantic analysis techniques such as parse trees (35, 36).

We are currently working on addressing a limitation of the Auto-CORPus package (25) that we used to process the full-text. In investigating our results, we found that any abbreviations that cannot be mapped to full names are not annotated, this includes for example cases where Greek letters are abbreviated by letters of Latin alphabet (α as 'a'). We also encountered instances where metabolites are abbreviated without the abbreviation (PC, PE, SM, and others) being defined, because the abbreviations are commonly used in the (lipidomics) community. These abbreviations were also not annotated in the (semi-automatically annotated) metabolomics corpus, but could easily be added as a set of rules to the algorithm. Lastly, superscripts and subscripts in the corpus are encoded differently from normal text by Auto-CORPus (25) and this causes negative impact to the annotation algorithms. We therefore anticipate periodically updating the metabolomics corpus based on user feedback and further research.

Deep learning models. Four LSTM-based models were evaluated on the same manually annotated test set as the annotation pipeline. ChemListem (18), which achieved an F1-score >0.90 on the BioCreative V5 (16) challenge data, achieved a lower F1-score of 0.77 when evaluated on the metabolomics corpus test set. This is in contrast with our MetaboListem model that used the same neural network architecture as ChemListem, but was trained on the metabolomics corpus, achieving an F1-score of ~0.89. The lower precision and recall rates of ChemListem when applied to our corpus are the results of higher false positives and false negatives, respectively. The false positives are expected due to the fact ChemListem was trained to recognise chemical entities and not metabolites specifically. The higher false negative rate is indicative of limited power for detecting metabolites - despite these also being chemical entities. The recall rates of ChemListem evaluated on the CEMP and metabolomics corpora are 0.89 (18) and 0.81 (here), respectively, which is indicative of the fact that metabolites appear less frequently in the BioCreative V5 dataset. Our MetaboListem model outperforms ChemListem on all three metrics for the metabolomics corpus, thus indicating that the BiLSTM network structure is generalisable to finer NER tasks and that training on relevant data further improves performance. This allows the algorithms to deal with changes in nomenclature and overcome the issue of having to define complex domain knowledge with limited sets of rules which was observed in NLP tasks for recognising genes and proteins previously (37). Moreover, we have shown how each algorithm achieves different levels of performance on individual sections of a publication, therefore algorithms developed only on abstracts (as these are easily available) may not perform as well on full-text sections. Our algorithms have been trained on a variety of sections (including abstracts) and have shown good performance across all sections.

Our second model, TABoLiSTM, replaced the GloVe embedding (20) that was used in MetaboListem with BERT (21) and BioBERT (22). The TABoLiSTM model with BioBERT achieved the highest F1-score (0.91) of all models evaluated

and improves on all metrics compared to MetaboListem. This improvement is directly attributable to using contextualised token embeddings - enhancing the contextual sensitivity - and using smaller token sequences that are in between character- and word-level embeddings. Smaller token sequences are known to be more suitable for NER tasks where segmenting words is difficult (38), for example with metabolites.

While the annotation pipeline has the highest recall rate, we observed that for many instances with spelling errors in metabolite names that these were recognised by the MetaboListem and TABoLiSTM models but were missed by the rule-based annotation pipeline. This includes for example misspelled entities such as 'slanine' (typo of 'alanine'), 'ace-toace acid' (typo of 'acetoacetic acid') and 'palmi acid' (typo of 'palmitic acid') indicating that the DL models are capable to generalise and identify unseen metabolite names, including the misspelled ones, by recognising the token structures and the contexts. In light of earlier discussion on the annotation pipeline not recognising endoxifen, the DL models do correctly recognise endoxifen as metabolite which is likely because the token structure of endoxifen is similar to that of tamoxifen in addition to the similar context embedding of metabolites.

MetaboListem outperforms the precision of the annotation pipeline by 0.7% which may be explained by the ambiguous synonyms falsely identified by the annotation pipeline. For the TABoLiSTM model (BioBERT embedding achieving higher precision by 4% compared to the annotation pipeline) this may be to learning contexts rather than learning the rules and regular structures designated to the annotation pipeline. The algorithms were (trained and) evaluated on the full text output from Auto-CORPus (25), however Auto-CORPus also provides separate JSON output files for table data and abbreviations and these files mostly contain single terms without context. Empirically, we found that although the DL models are context sensitive by construction (BiLSTM network and BioBERT embedding) they detect entities in tables and abbreviation lists with high accuracy comparable to the full text results.

Overall, the semi-automatic annotation pipeline is ideally used in parallel with the DL methods as combining these can to provide wider coverage of features that are not yet learned by the DL algorithms due to limited training examples. Likewise, this is also relevant for addressing some limitations of the DL methods identified from their false positives. We found that while most bacteria are not recognised, a very small number of bacterial species are picked up as 'metabolites' by the DL methods. These terms all include tokens that can also be found in metabolites such as 'lact' (*Lactobacillus*, *Lactococcus*), 'butyr' (*Butyricimonas*, *Pseudobutyribrio*, *Butyribrio*), 'chloro' (*Chloroflexi*), 'succini' (*Succinivibrio*), 'sulfo' (*Desulfovibrio*), and 'acid' and 'amino' (*Acidaminobacter*). As part of a wider effort to provide NLP tools to the omics community, we are developing an NER algorithm for microbiome studies which can also be used to filter out false positives from our DL models when complete. Other metabolites are identified (by both the DL and rule-

based methods) that are in fact part of an enzyme or pathway name. In our DL models we have not included negative training examples of terms in these contexts, however more elaborate rules can also be developed to filter these false positives out in a post-processing step. For example, we observed in the training set that the words ‘process’ and ‘success’ are falsely recognised as metabolites in certain context. Based on the dictionary list, no known metabolite name/synonym ends with ‘ss’, so therefore all terms that match the regular pattern ‘ss\$’ could be excluded. This type of rule-based post-processing was not used for our models because of the risk of introducing bias in model evaluation, where an n-gram of ‘success’ is found in metabolite names such as succinic acid/succinate.

In terms of usability of the DL algorithms, the model size is an important factor. The TABoLiSTM models require 827MB of storage, whereas ChemListem and MetaboListem only need 26MB. Therefore, TABoLiSTM may be too heavy to be leveraged on web services and would ideally be used on local systems. Therefore, employing the lighter MetaboListem model in an online environment is more feasible, regardless of the cost of 1.8% in the F1-score. On average, a full-text article can be annotated by MetaboListem in under 10s, whereas for TABoLiSTM this is close to 40 seconds.

Biological interpretation. We have exemplified how the algorithms can be used to speed up literature review by providing a summary of all identified metabolites in the full text. The most common metabolites reported appear to be those that can be measured using virtually all major assays. Glucose is reported most frequently for 4 of the traits, and as one of the top 5 metabolites for the other traits. For cancer studies, the high co-occurrence of glucose and lactate is indicative of the Warburg effect (39) where glucose is converted to lactate for cancer cell energy metabolism. However, the remaining top reported metabolites are mostly amino acids, which is also observed for other traits. At the same time, some trait-specific metabolites are found such as bilirubin and cholesterol for liver diseases, cholesterol and triglycerides for metabolic syndrome, cotinine and nicotine for smoking, and tryptophan, kynurenine and serotonin for neurological/psychiatric/brain diseases. The latter being reflective of the serotonin hypothesis (40) that has resulted in antidepressant drugs such as kynurenine uptake inhibitors. The proximities between nodes in the network graphs are known to express community structures (41), hence it is not surprising to find so many amino acids and other common metabolites near the centre of the graphs.

Conclusion

In summary, this work contributed the first annotated metabolomics corpus to facilitate future NER development for metabolomics research. Our four novel NER algorithms, including a rule-based annotation algorithm and three DL NER systems, are implemented in Python. All data generated, including the annotated metabolomics corpus (manually annotated test set and semi-automatically anno-

tated training and validation sets), code to generate models, the MetaboListem model and code for post-processing, is available via <https://github.com/omicsNLP/MetaboliteNER>. These algorithms have all shown state-of-the-art performance (F1 score >0.88) for metabolite NER tasks (with the best model achieving an F1-score of 0.91), thereby facilitating future text-mining of metabolomics literature.

AUTHOR CONTRIBUTIONS

Conceptualization, J.M.P. and T.B.; data curation, C.S.Y. and J.M.P.; formal analysis, C.S.Y.; funding acquisition, J.M.P. and T.B.; investigation, C.S.Y.; methodology, C.S.Y. and J.M.P.; project administration, J.M.P. and T.B.; resources, J.M.P.; software, C.S.Y.; supervision, J.M.P. and T.B.; validation, J.M.P.; visualization, C.S.Y.; writing—original draft preparation, C.S.Y.; writing—review and editing, J.M.P. and T.B. All authors have read and agreed to the published version of the manuscript.

This work has been supported by Health Data Research (HDR) UK and the Medical Research Council via an UKRI Innovation Fellowship to TB (MR/S003703/1) and a Rutherford Fund Fellowship to JMP (MR/S004033/1).

The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

¹orcidgreenORCID: 0000-0002-4971-9003 (JMP).

Bibliography

- David S. Wishart. Metabolomics for investigating physiological and pathophysiological processes. *Physiological Reviews*, 99(4):1819–1875, 2019. doi: 10.1152/physrev.00035.2018.
- Paul Miggiels, Bert Wouters, Gerard J.P. van Westen, Anne-Charlotte Dubbelman, and Thomas Hankemeier. Novel technologies for metabolomics: More for less. *TrAC Trends in Analytical Chemistry*, 120:115323, 2019. ISSN 0165-9936. doi: 10.1016/j.trac.2018.11.021.
- Lutz Bornmann and Rüdiger Mutz. Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references: Growth Rates of Modern Science: A Bibliometric Analysis Based on the Number of Publications and Cited References. *Journal of the Association for Information Science and Technology*, 66(11):2215–2222, November 2015. ISSN 23301635. doi: 10.1002/asi.23329.
- Richard G Jackson, Rashmi Patel, Nishamali Jayatilake, Anna Kolliakou, Michael Ball, Genevieve Gorrell, Angus Roberts, Richard J Dobson, and Robert Stewart. Natural language processing to extract symptoms of severe mental illness from clinical text: the clinical record interactive search comprehensive data extraction (CRIS-CODE) project. *BMJ Open*, 7(1):e012012, January 2017. doi: 10.1136/bmjopen-2016-012012.
- Sayedmostafa Sheikhalishahi, Riccardo Miotto, Joel T Dudley, Alberto Lavelli, Fabio Rinaldi, and Venet Osmani. Natural language processing of clinical notes on chronic diseases: Systematic review. *JMIR Med Inform*, 7(2):e12239, 4 2019. ISSN 2291-9694. doi: 10.2196/12239.
- Nikola Milosevic, Cassie Gregson, Robert Hernandez, and Goran Nenadic. A framework for information extraction from tables in biomedical literature. *International Journal on Document Analysis and Recognition (IJДАР)*, 22(1):55–78, 2 2019. doi: 10.1007/s10032-019-00317-0.
- Volodymyr Kuleshov, Jialin Ding, Christopher Vo, Braden Hancock, Alexander Ratner, Yang Li, Christopher Ré, Serafim Batzoglou, and Michael Snyder. A machine-compiled database of genome-wide association studies. *Nature Communications*, 10(1), July 2019. doi: 10.1038/s41467-019-11026-x.
- Chikashi Nobata, Paul D. Dobson, Syed A. Iqbal, Pedro Mendes, Jun’ichi Tsujii, Douglas B. Kell, and Sophia Ananiadou. Mining metabolites: extracting the yeast metabolome from the literature. 7(1):94–101, October 2010. doi: 10.1007/s11306-010-0251-6.
- Wutthipong Kongburan, Praisian Padungweang, Worarat Krathu, and Jonathan H. Chan. Metabolite named entity recognition: A hybrid approach. In Akira Hirose, Seiichi Ozawa, Kenji Doya, Kazushi Ikeda, Minh Lee, and Derong Liu, editors, *International Conference on Neural Information Processing: Neural Information Processing*, pages 451–460. Springer International Publishing, 2016. ISBN 978-3-319-46687-3.
- Erica L.-W. Majumder, Elizabeth M. Billings, H. Paul Benton, Richard L. Martin, Amelia Palermo, Carlos Guijas, Markus M. Rinschen, Xavier Domingo-Almenara, J. Rafael Montenegro-Burke, Bradley A. Tagtow, Robert S. Plumb, and Gary Siuzdak. Cognitive analysis of metabolomics data for systems biology. *Nature Protocols*, 16(3):1376–1418, March 2021. ISSN 1754-2189, 1750-2799. doi: 10.1038/s41596-020-00455-4.
- Dritan Kodra, Petros Pousinis, Panagiotis A. Vorkas, Katerina Kademoglou, Theodoros Liapikos, Alexandros Pechlivanis, Christina Virgilou, Ian D. Wilson, Helen Gika, and Georgios Theodoridis. Is current practice adhering to guidelines proposed for metabolite identification in lc-ms untargeted metabolomics? a meta-analysis of the literature. *Journal of Proteome Research*, 0(0):online first, 0. doi: 10.1021/acs.jproteome.1c00841.
- David S Wishart, Yannick Djoumbou Feunang, Ana Marcu, An Chi Guo, Kevin Liang, Rosa Vázquez-Fresno, Tanvir Sajed, Daniel Johnson, Carin Li, Naama Karu, Zinat Sayeeda, Elvis

- Lo, Nazanin Assempour, Mark Berjanskii, Sandeep Singhal, David Arndt, Yonjie Liang, Hasan Badran, Jason Grant, Arnau Serra-Cayuela, Yifeng Liu, Rupa Mandal, Vanessa Neveu, Allison Pon, Craig Knox, Michael Wilson, Claudine Manach, and Augustin Scalbert. HMDb 4.0: the human metabolome database for 2018. *Nucleic Acids Research*, 46 (D1):D608–D617, January 2018. ISSN 0305-1048, 1362-4962. doi: 10.1093/nar/gkx1089.
13. D. S. Wishart, C. Knox, A. C. Guo, R. Eisner, N. Young, B. Gautam, D. D. Hau, N. Psychogios, E. Dong, S. Bouatra, R. Mandal, I. Sinelnikov, J. Xia, L. Jia, J. A. Cruz, E. Lim, C. A. Sobsey, S. Shrivastava, P. Huang, P. Liu, L. Fang, J. Peng, R. Fradette, D. Cheng, D. Tzur, M. Clements, A. Lewis, A. De Souza, A. Zuniga, M. Dawe, Y. Xiong, D. Clive, R. Greiner, A. Nazirova, R. Shaykhtudinov, L. Li, H. J. Vogel, and I. Forsythe. HMDB: a knowledgebase for the human metabolome. *Nucleic Acids Research*, 37(Database):D603–D610, January 2009. ISSN 0305-1048, 1362-4962. doi: 10.1093/nar/gkn810.
14. Lynette Hirschman, Alexander Yeh, Christian Blaschke, and Alfonso Valencia. Overview of BioCreAtivE: critical assessment of information extraction for biology. *BMC Bioinformatics*, 6(Suppl 1):S1, 2005. ISSN 14712105. doi: 10.1186/1471-2105-6-S1-S1.
15. Martin Krallinger, Florian Leitner, Odbulja Rabal, Miguel Vazquez, Julien Oyarzabal, and Alfonso Valencia. CHEMDNER: The drugs and chemical names extraction challenge. *Journal of Cheminformatics*, 7(S1):S1, December 2015. ISSN 1758-2946. doi: 10.1186/1758-2946-7-S1-S1.
16. Qinghua Wang, Shabbir S. Abdul, Lara Almeida, Sophia Ananiadou, Yalbi I. Balderas-Martinez, Riza Batista-Navarro, David Campos, Lucy Chilton, Hui-Jou Chou, Gabriela Contreras, Laurel Cooper, Hong-Jie Dai, Barbra Ferrell, Julianne Fluck, Socorro Gama-Castro, Nancy George, Georgios Gkoutos, Afroza K. Irin, Lars J. Jensen, Silvia Jimenez, Toni R. Jue, Ingrid Keseler, Sumit Madan, Sérgio Matos, Peter McQuilton, Marija Milacic, Matthew Mort, Jeyakumar Natarajan, Evangelos Pafilis, Emiliano Pereira, Shruti Rao, Fabio Rinaldi, Karen Rothfels, David Salgado, Raquel M. Silva, Onkar Singh, Raymund Stefancsik, Chu-Hsien Su, Suresh Subramani, Hamsa D. Tadepally, Loukia Tsaprouni, Nicole Vasilevsky, Xiaodong Wang, Andrew Chatr-Aryamontri, Stanley J. F. Lauderdalekind, Sheri Matis-Mitchell, Johanna McEntyre, Sandra Orchard, Sangya Pundir, Raul Rodriguez-Esteban, Kimberly Van Aukun, Zhiyong Lu, Mary Schaeffer, Cathy H. Wu, Lynette Hirschman, and Cecilia N. Arighi. Overview of the interactive task in BioCreative V. *Database*, 2016:baw119, 2016. ISSN 1758-0463. doi: 10.1093/database/baw119.
17. Robert Leaman, Chih-Hsuan Wei, and Zhiyong Lu. tmChem: a high performance approach for chemical named entity recognition and normalization. *Journal of Cheminformatics*, 7 (S1):S3, December 2015. ISSN 1758-2946. doi: 10.1186/1758-2946-7-S1-S3.
18. Peter Corbett and John Boyle. Chemlistem: chemical named entity recognition using recurrent neural networks. *Journal of Cheminformatics*, 10(1):59, December 2018. ISSN 1758-2946. doi: 10.1186/s13321-018-0313-8.
19. Maryam Habibi, Leon Weber, Mariana Neves, David Luis Wiegandt, and Ulf Leser. Deep learning with word embeddings improves biomedical named entity recognition. *Bioinformatics*, 33(14):i37–i48, July 2017. ISSN 1367-4803, 1460-2059. doi: 10.1093/bioinformatics/btx228.
20. Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1162.
21. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North*, pages 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.
22. Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, pages 1234–1240, September 2019. ISSN 1367-4803, 1460-2059. doi: 10.1093/bioinformatics/btz682.
23. Jeff Beck. Report from the Field: PubMed Central, an XML-based Archive of Life Sciences Journal Articles. Montréal, Canada. doi: 10.4242/BalisageVol6.Beck01.
24. Cong Sun, Zhihao Yang, Lei Wang, Yin Zhang, Hongfei Lin, and Jian Wang. Biomedical named entity recognition using BERT in the machine reading comprehension framework. *Journal of Biomedical Informatics*, 118:103799, June 2021. ISSN 15320464. doi: 10.1016/j.jbi.2021.103799.
25. Tim Beck, Tom Shorter, Yan Hu, Zhuoyu Li, Shujian Sun, Casiana M. Popovici, Nicholas A. R. McQuibban, Filip Makraduli, Cheng S. Yeung, Thomas Rowlands, and Joram M. Posma. Auto-CORPus: A Natural Language Processing Tool for Standardising and Reusing Biomedical Literature. *Frontiers in Digital Health*, 4(788124):788124, January 2022. ISSN 2673-253X. doi: 10.3389/fdgh.2022.788124.
26. Werner Ceusters. An information artifact ontology perspective on data collections and associated representational artifacts. *Studies in health technology and informatics*, 180:68–72, 2012. ISSN 0926-9630.
27. Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.6.
28. David J. Hand, Peter Christen, and Nishadi Kirielle. F*: an interpretable transformation of the f-measure. *Machine Learning*, 110(3):451–456, March 2021. doi: 10.1007/s10994-021-05964-1.
29. David M Jessop, Sam E Adams, Egon L Willighagen, Lezan Hawizy, and Peter Murray-Rust. OSCAR4: a flexible architecture for chemical text-mining. *Journal of Cheminformatics*, 3(1): 41, December 2011. ISSN 1758-2946. doi: 10.1186/1758-2946-3-41.
30. Paula de Matos, Rafael Alcántara, Adriano Dekker, Marcus Ennis, Janna Hastings, Kenneth Haug, Immaculada Spiteri, Steve Turner, and Christoph Steinbeck. Chemical Entities of Biological Interest: an update. *Nucleic Acids Research*, 38(suppl_1):D249–D254, January 2010. ISSN 0305-1048, 1362-4962. doi: 10.1093/nar/gkp886.
31. Yarin Gal and Zoubin Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1019–1027, 2016.
32. Mathieu Jacomy, Tommaso Venturini, Sebastien Heymann, and Mathieu Bastian. ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLoS ONE*, 9(6):e98679, June 2014. ISSN 1932-6203. doi: 10.1371/journal.pone.0098679.
33. Mathieu Bastian, Sebastien Heymann, and Mathieu Jacomy. Gephi: An Open Source Software for Exploring and Manipulating Networks. 2009. doi: 10.13140/2.1.1341.1520.
34. Antony J. Williams. ChempSpider: A Platform for Crowdsourced Collaboration to Curate Data Derived From Public Compound Databases. In Sean Ekins, Maggie A. Z. Hupcey, and Antony J. Williams, editors, *Collaborative Computational Technologies for Biomedical Research*, pages 363–386. John Wiley & Sons, Inc., Hoboken, NJ, USA, May 2011. ISBN 9781118026038 9780470638033. doi: 10.1002/9781118026038.ch22.
35. Katrin Fundel, Robert Küffner, and Ralf Zimmer. RelEx—Relation extraction using dependency parse trees. *Bioinformatics*, 23(3):365–371, 12 2006. ISSN 1367-4803. doi: 10.1093/bioinformatics/btl1616.
36. Guodong Zhou, M. Zhang, D. Ji, and Qiaoming Zhu. Tree kernel-based relation extraction with context-sensitive structured parse tree information. In *EMNLP*, 2007.
37. Ramón A-A. Erhardt, Reinhard Schneider, and Christian Blaschke. Status of text-mining techniques applied to biomedical text. *Drug Discovery Today*, 11(7):315–325, 2006. ISSN 1359-6446. doi: 10.1016/j.drudis.2006.02.011.
38. Mourad Gridach. Character-level neural network for biomedical named entity recognition. *Journal of Biomedical Informatics*, 70:85–91, June 2017. ISSN 15320464. doi: 10.1016/j.jbi.2017.05.002.
39. M. G. Vander Heiden, L. C. Cantley, and C. B. Thompson. Understanding the Warburg Effect: The Metabolic Requirements of Cell Proliferation. *Science*, 324(5930):1029–1033, May 2009. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1160809.
40. I.P. Lapin and G.F. Oxenkrug. Intensification of the central serotonergic processes as a possible determinant of the thymoleptic effect. *The Lancet*, 293(7586):132–136, 1969. ISSN 0140-6736. doi: 10.1016/S0140-6736(69)91140-4.
41. Andreas Noack. Modularity clustering is force-directed layout. *Physical Review E*, 79(2): 026102, February 2009. ISSN 1539-3755, 1550-2376. doi: 10.1103/PhysRevE.79.026102.