

Klarigi

Slater, Luke T; Williams, John A; Schofield, Paul N; Russell, Sophie; Pendleton, Samantha C; Karwath, Andreas; Fanning, Hilary; Ball, Simon; Hoehndorf, Robert; Gkoutos, Georgios V

DOI:

[10.1016/j.combiomed.2022.106425](https://doi.org/10.1016/j.combiomed.2022.106425)

License:

Creative Commons: Attribution (CC BY)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Slater, LT, Williams, JA, Schofield, PN, Russell, S, Pendleton, SC, Karwath, A, Fanning, H, Ball, S, Hoehndorf, R & Gkoutos, GV 2023, 'Klarigi: characteristic explanations for semantic biomedical data', *Computers in Biology and Medicine*, vol. 153, 106425. <https://doi.org/10.1016/j.combiomed.2022.106425>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.



Contents lists available at ScienceDirect

Computers in Biology and Medicine

journal homepage: www.elsevier.com/locate/complbiomed

Klarigi: Characteristic explanations for semantic biomedical data[☆]

Luke T. Slater^{a,b,f,i,*}, John A. Williams^{a,b,i}, Paul N. Schofield^g, Sophie Russell^a,
Samantha C. Pendleton^a, Andreas Karwath^{a,b,f,i}, Hilary Fanning^{b,i}, Simon Ball^{b,i},
Robert Hoehndorf^h, Georgios V. Gkoutos^{a,b,c,d,e,f,i}

^a College of Medical and Dental Sciences, Institute of Cancer and Genomic Sciences, University of Birmingham, UK

^b Institute of Translational Medicine, University Hospitals Birmingham, NHS Foundation Trust, UK

^c NIHR Experimental Cancer Medicine Centre, UK

^d NIHR Surgical Reconstruction and Microbiology Research Centre, UK

^e NIHR Biomedical Research Centre, UK

^f MRC Health Data Research UK (HDR UK), Midlands, UK

^g Department of Physiology, Development, and Neuroscience, University of Cambridge, UK

^h Computer, Electrical and Mathematical Sciences & Engineering Division, Computational Bioscience Research Center, King Abdullah University of Science and Technology, UK

ⁱ University Hospitals Birmingham NHS Foundation Trust, Edgbaston, Birmingham, UK

ARTICLE INFO

Dataset link: <https://github.com/reality/klarigi>
, <https://github.com/reality/klarigi>

Keywords:

Semantic explanation
Ontology
Semantic analysis
Enrichment analysis
Phenotypes
Explicability
Phenotype profiles

ABSTRACT

Annotation of biomedical entities with ontology classes provides for formal semantic analysis and mobilisation of background knowledge in determining their relationships. To date, enrichment analysis has been routinely employed to identify classes that are over-represented in annotations across sets of groups, such as biosample gene expression profiles or patient phenotypes, and is useful for a range of tasks including differential diagnosis and causative variant prioritisation. These approaches, however, usually consider only univariate relationships, make limited use of the semantic features of ontologies, and provide limited information and evaluation of the explanatory power of both singular and grouped candidate classes. Moreover, they are not designed to solve the problem of deriving cohesive, characteristic, and discriminatory sets of classes for entity groups. We have developed a new tool, called Klarigi, which introduces multiple scoring heuristics for identification of classes that are both compositional and discriminatory for groups of entities annotated with ontology classes. The tool includes a novel algorithm for derivation of multivariable semantic explanations for entity groups, makes use of semantic inference through live use of an ontology reasoner, and includes a classification method for identifying the discriminatory power of candidate sets, in addition to significance testing apposite to traditional enrichment approaches. We describe the design and implementation of Klarigi, including its scoring and explanation determination methods, and evaluate its use in application to two test cases with clinical significance, comparing and contrasting methods and results with literature-based and enrichment analysis methods. We demonstrate that Klarigi produces characteristic and discriminatory explanations for groups of biomedical entities in two settings. We also show that these explanations recapitulate and extend the knowledge held in existing biomedical databases and literature for several diseases. We conclude that Klarigi provides a distinct and valuable perspective on biomedical datasets when compared with traditional enrichment methods, and therefore constitutes a new method by which biomedical datasets can be explored, contributing to improved insight into semantic data.

[☆] GVG and LTS acknowledge support from the NIHR Birmingham ECMC, UK, NIHR Birmingham SRMRC, UK, Nanocommons H2020-EU (731032), MAESTRIA, UK (Grant agreement ID 965286), HYPERMARKER, UK (Grant agreement ID 101095480), and the MRC Heath Data Research UK (HDRUK/CFC/01), an initiative funded by UK Research and Innovation, Department of Health and Social Care (England) and the devolved administrations, and leading medical research charities, UK. The views expressed in this publication are those of the authors and not necessarily those of the NHS, the National Institute for Health Research, the Medical Research Council or the Department of Health. RH, PNS and GVG were supported by funding from King Abdullah University of Science and Technology (KAUST), Saudi Arabia Office of Sponsored Research (OSR) under Award No. URF/1/3790-01-01. AK was additionally supported by the Medical Research Council, UK (MR/S003991/1). PNS and GVG acknowledge the support of The Alan Turing Institute, UK.

* Corresponding author at: College of Medical and Dental Sciences, Institute of Cancer and Genomic Sciences, University of Birmingham, UK.

E-mail address: l.slater.1@bham.ac.uk (L.T. Slater).

<https://doi.org/10.1016/j.complbiomed.2022.106425>

Received 30 August 2022; Received in revised form 4 December 2022; Accepted 13 December 2022

Available online 22 December 2022

0010-4825/© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Background

The power of ontologies for representation of biomedical data has led to ontological annotation becoming a preferred method for the characterisation of large biomedical datasets [1,2]. Increasingly, domains such as medicine are adopting ontologies for annotation, aggregation, and analysis of patient data, opening up new methods for analysis hitherto unavailable using standard terminologies and nosologies [3]. The ability to link data to these symbolic representations reduces inaccuracy and ambiguity, and permits integration with background knowledge through contextual aggregation. As such, biomedical ontologies and their instance data constitute a massive source of formalised knowledge.

This repository of formalised knowledge greatly facilitates secondary use, taking advantage of the semantic features of ontologies to integrate and explain the large and multi-modal datasets annotated with them. This is typically used for the purposes of ontology-based data access, integration, and preparation, enhancing standard Extract Transform Load (ETL) processes [4,5]. In addition, semantic analysis methods are now applied to many knowledge synthesis and classification tasks, including prediction of protein–protein interactions and identification of rare disease genetic variants [6]. In the clinical space, semantic methods have been applied to tasks including diagnosis of rare and common diseases [7,8] and identification of subtypes of diseases such as autism [9]. In addition, the synthesis of classical ontology-based methods and machine learning (ML) is seen as increasingly powerful [10].

Of great importance to the trust in, and consequently willingness to adopt, new analytical methods on the part of both the medical profession and the public, is the problem of explainability [11]. While application of a machine learning (ML) algorithm can generate a classifier based on large amounts of data, for example in patient population stratification, current methods make understanding the clinical basis for classification difficult. Despite the increasing use of semantics in biomedical analysis, the subsequent derivation of semantic explanations for classifications, outcomes, labels, or groups generated by those analyses, remains a challenging task, and a major practical, ethical, and technical issue [12,13]. This is a problem for any new algorithmic approach, though current discussion has centred on the new ML approaches emerging in the health sciences.

We define the task of semantic explanation as the identification of characteristic sets of classes for groups of entities described by ontology classes. What constitutes a characteristic set of classes depends upon the parameters of a particular investigation, but can most often be defined by classes that are some combination of explanatory, compositional, and discriminatory.

One traditional approach to semantic explanation is enrichment analysis. For example, and most prominently, gene set enrichment analysis (GSEA) identifies classes that are significantly over-expressed in a set of genes [14]. Enrichment analysis techniques have also been used to explore over-representation of terms in groups of entities associated with other datasets and ontologies [15,16]. There are, however, multiple drawbacks to the use of enrichment analysis methods for semantic explanation:

- Traditional enrichment methods using a univariate statistical test provide limited information about the relationship of each explanatory term with the group of interest.
- They usually do not consider the multivariate contribution of sets of terms to a non-redundant characterisation of the group of interest.
- Most make little or no use of the semantic features of ontologies, or information-theoretic methods of evaluating candidate terms.
- They provide limited tools and metrics for evaluation of the overall explanatory power of groups of terms with respect to the group of interest.

- They can be difficult or impossible to use with non-preconfigured ontologies, and often have a dedicated focus on the Gene Ontology (GO) [17].

Most methods include built-in calls to a single ontology: fgsea [18], GSEABase [19], and GREAP [20] leverage GO; DOSE [21] implements the Disease Ontology; ReactomePA makes use of the Reactome pathway database [22], with several recent works centring pathway enrichment [23–25]; MamPhea [26] uses the Mammalian Phenotype Ontology. Other methods implement multiple ontologies as packages or online tools, including DAVID [27], GProfiler [28], and XGR [29].

Of these, only XGR allows users to provide their own custom ontology, and this must be a graph object mutated from an OWL file. Additionally the XGR package is the only one of these which does not explicitly rely on pre-populated, species-specific gene-ontology annotation databases. This limits analyses to both species-specific questions and to genes or gene products, and these tools have little direct relevance to non-genetic translational medicine. In traditional GSEA, ranking genes in order of relevance of biological assays and defining a function as the number of genes above that sorted list, can be formulated as a random walk along sorted features. This approach implements multiple sorted lists from multiomics experiments and implements joint random walks among assays. Another recent development uses Bayesian regression to model the relationship between phenotypes and the proportion of patients in a population with rare alleles, and interpreting the probability of associations being equal or depending on the phenotypic similarity of the patients with rare alleles [30]. This approach requires informed assumptions about existing gene/phenotype relationships in literature to weight priors. All of these methods are specific to the genetic domain, and cannot therefore be easily applied to other kinds of biomedical data.

While methods do exist to help correct for the dependence of statistical tests on related ontology classes, these post-hoc corrections do not take full advantage of the ontology to make inferential associations. Some methods such as XGR [29] and Func [31] do take ontology structure into account, with both requiring preprocessing beforehand if using ontologies other than GO (Func) or using ontologies outside a predefined set or without pre-computed annotation databases (XGR). In particular, OntoFunc provides inferential preprocessing for Func [32]. Recent Bayesian approaches to ontology enrichment likewise make use of ontology structure, but require additional numerical data to construct informative priors [33], or are restricted to GO [34].

Furthermore, most enrichment approaches consider only univariate relationships for individual terms; some investigations have approached multivariable enrichment, with one using a random walk approach on omics datasets [35], and another relying on pre-processing into principal components [36].

Moreover, enrichment methods are not intended to identify full and cohesive sets of characteristic terms for groups of annotated entities, and these methods therefore provide a limited solution to the task of semantic explanation. We propose to solve these problems by developing a new method for semantic explanation.

The method introduces metrics for quantification of multiple aspects of the relationship between classes, sets of terms, and groups of annotated entities. These include a configurable measure of information content, facilitating the involvement of information theoretic methods of evaluating candidates. This is combined with a new algorithm for derivation of multivariable explanatory sets of terms, optimised for discrimination and/or composition. The method makes use of ontological inference via an ontology reasoner, and can be used directly with any consistent Ontology Web Language (OWL) ontology without preprocessing. We also include native support for the Phenopacket format [37] to facilitate efficient use for large-scale import any analysis of clinical phenotype data. To provide additional ability to determine the discriminatory power of explanatory sets, it also includes a classification method that can be re-applied to the given dataset, or applied

to unseen data. In aggregate, these features address the limitations of traditional enrichment approaches, while also providing similar information through the implementation of permutation-derived p -values for univariate scoring metrics.

In this work, we formally describe the method and its implementation. We discuss existing approaches and their potential limitations, explaining how our proposed approach solves these problems. We then describe development of the tool, called Klarigi, that implements the method, resulting in an analysis framework that can be used to explore semantic datasets. The name of the implementation, ‘Klarigi,’ is adopted from the Esperanto word meaning ‘to explain.’ The work is related to a previous investigation, which describes an early version of the approach described in the context of a cluster explanation algorithm developed for a particular application [38]. In this work, the approach is heavily modified and improved upon, including many changes to the algorithm and scoring system, additional features for characteristic and discriminatory analysis, and a formal description of the design, in addition to its implementation and evaluation in the form of a generally applicable software framework. These modifications are described in more detail throughout the manuscript.

To evaluate our approach, we develop two use-cases for characterisation of patient populations. We compare and contrast the results of the analysis with those of an enrichment analysis, and compare with existing literature and manually curated biomedical databases as gold standard. In doing this, we show that Klarigi can be used to generate insights into biomedical datasets via semantic explanation, providing a perspective that is both distinct and complementary to the use of enrichment analysis.

Design and implementation

The fundamental challenge of semantic explanation is to determine, given a set of groups, and sets of entities described by ontology classes associated with those groups, a set of characteristic classes for each group. Since we developed this approach to be applicable to any biomedical ontology, we will describe it, including a data model and algorithms, in abstract terms first. We will then describe their implementation into the Klarigi tool.

Data modelling

Klarigi aims to determine sets of classes that characterise a group of entities annotated by ontology classes. Doing this requires, at minimum, a corpus with three features:

1. A set of entities, each associated with a set of ontology classes.
2. A set of groups, which are each associated with a set of entities.
3. An OWL ontology describing all classes in the corpus.

We can define these as sets, which will allow us to define metrics and heuristics on their contents. G is a set of groups, where G_i is the set of entities ascribed to the group i . E is a set of entities, where E_j is the set of classes associated with entity j . O is a set of classes in the ontology, where O_x is a class in the ontology.

To illustrate concepts throughout this section, we will define an example corpus. For argument, we have 1000 patient phenotype profiles, all described by classes from the human phenotype ontology, HPO [39]. All of these patients have been diagnosed with either pulmonary embolism or pneumonia, and we are interested in characterising and distinguishing between those diseases based on information from the phenotype profiles. To represent this in our data model above, we consider patients to be our entities, and E will be populated with the set of phenotype profiles associated with each patient. O is populated with all classes in HPO, since this is the ontology used to describe the patients. If we are interested in characterising these two diagnoses, we could create a data model with two groups in G , pulmonary embolism and pneumonia, identifying the patients. In this case, entities are patients or

patient profiles, and so E would be populated with the set of phenotype profiles associated with each patient. The elements of G define sets that identify patients in E with each disease. Note that these groups are not mutually exclusive, however for simplicity in this example we will consider that each patient belongs to *either* the pulmonary embolism or pneumonia group.

Metrics

To identify ‘useful’ sets of explanatory classes for groups, we need to balance multiple qualities of the relationship between candidate classes and the group. Specifically, we want to identify classes that are, in comparison with other candidates:

Compositional Associated with greater proportion of entities in the considered group.

Discriminatory More associated with entities in the considered group than in other groups.

Informative More specific.

Linking this with our example, the first item concerns identifying HPO classes that appear in a greater proportion of the phenotype profiles for either disease. The second concerns classes that appear in phenotype profiles for one disease more so than another. The third determines how informative or specific a class is; for example, general cardiovascular symptoms are less specific than particular symptoms like atrial fibrillation.

To quantify these qualities and thus make them comparable, we define three univariate scores to measure them, encoding numeric measures of the explanatory power of each candidate class for the given group. We first define a simple function that enables us to determine whether or not a particular entity in the corpus is annotated with a class (or any of its subclasses):

$$member(C, N) = \begin{cases} 1 & |C \cap N| > 0, \\ 0 & otherwise \end{cases} \quad (1)$$

Given two sets C and N , any two sets of ontology classes, this function returns 1, if the size of the intersection of those sets is at least 1, and 0 otherwise. This allows us to identify whether any classes associated with a particular entity include our candidate class. Using this function, we can then define the first score, inclusion:

$$inclusion(O_x, G_j) = \frac{\sum_{\substack{z=1 \\ z \in G_j}}^{|E|} member(subclass(O_x), E_z)}{|G_j|} \quad (2)$$

This score measures the proportion of entities in a given group that are annotated with either the candidate term or any of its subclasses. The *subclass* function is a call to the ontology reasoner, and returns a set of all transitive subclasses of the class passed as argument, including the original class. This enables the account of the deductive inferences made by the ontology reasoner, entailed by the axioms asserted in the considered ontology. For example, if 250 of 500 phenotype profiles in the pneumonia group were associated with supraventricular arrhythmia (HP:0005115) or any of its subclasses, the *inclusion* score for supraventricular arrhythmia (HP:0005115) with respect to pneumonia would be 0.5. As such, this measure is a quantification of membership proportion enhanced by the notion of subsumptive hierarchy in the ontology, either explicitly asserted or inferred by the reasoner.

We next measure how uniquely a candidate class characterises the considered group, or its over-representation in the group with respect

to others. Our previous work [38] defined a measure of exclusion ($exclusion_{old}$) as:

$$exclusion_{old}(O_i, G_j) = 1 - \frac{\sum_{x \neq j}^{|G|} \sum_{y=1}^{|E|} member(subclass(O_i), E_y)}{\sum_{x \neq j}^{|G|} |G_x|} \quad (3)$$

This provides an inverse measure of the proportion of entities annotated with the class that are attributed to *other* groups: the fewer entities external to the considered group that are annotated with the class, the greater the exclusion score. For example, in the case of a dataset with 700 pneumonia patients and 300 pulmonary embolism patients and an inclusion of 0.5 for a particular phenotype in both groups, the exclusion score would be 0.5 in both cases: equal to one minus the *inclusion* score of the phenotype for the other group, which captures the overall absolute equality, though is unintuitive in its provision of a non-zero value.

In GSEA, measurement of comparative expression is usually achieved by an *enrichment score*, which is usually defined by z-scores and/or odds and relative risk ratios, that measure the strength of association between the class and group membership; the ratio of B given the presence of A, and its symmetric inverse. However, their interpretation is challenging due to the behaviour of ratios and confusion between odds, relative, and absolute risk [40]. In the above example, the relative risk of both pneumonia and pulmonary embolism would be 1. However, this can become misleading in the presence of larger class imbalances [41]. For example, if we had 950 cases of pneumonia, and 50 cases of pulmonary embolism. If a phenotype appears thrice in this dataset, and two of those are associated with pulmonary embolism, we would get the extremely large values of 37.96 and 39.54 for relative risk and odds ratio respectively. This example also illustrates the limitation of the $exclusion_{old}$ score, which would be 0.96 for pneumonia and 0.998 for pulmonary embolism, capturing in both cases the low prevalence of the class in the other group. While these are valid associations, they do not take into account the absolute risk of the event occurring with respect to the disparity in the group sizes, and so large values are given for uncommon and unlikely phenotypes, which can be easily misinterpreted. To address this issue, we define exclusion as:

$$exclusion(O_i, G_j) = \frac{\sum_{z=1}^{|E|} member(subclass(O_x), E_z)}{\sum_{z=1}^{|E|} member(subclass(O_x), E_z)} - \frac{|G_j|}{|E|} \quad (4)$$

In this measure, the proportion of entities associated with the class that belong to the considered group is determined. From this figure, the proportion of the total entities in the dataset that are associated with the considered group is subtracted. This measures the exclusive representation of the class in the group versus the representation of the group in the overall population, and can be explained in terms intuitively similar to the calculation of an odds ratio: the odds of the event occurring in the exposed group minus the odds of exposure.

As such, it reflects the representation of a class in a considered group, balanced with the absolute likelihood of the group appearing in the whole dataset. In the example above, the exclusion score for the phenotype with 0.5 inclusion with respect to both pneumonia and pulmonary embolism would be 0. This is because neither are more exclusively associated with the phenotype when their relative frequencies are taken into account. In the latter example, our proposed phenotype association profile would lead to pulmonary embolism receiving an exclusion score of 0.617 and pneumonia an exclusion score of -0.617 . These scores express the exclusivity of a phenotype's association with a group in the context of the overall appearance of the group in the dataset, providing a more intuitive metric for whether a class exclusively characterises entities in a group.

We further address the problem of balancing the internal and external characteristic power of candidate classes by introducing an additional heuristic to provide a balanced measure of our previously

defined *inclusion* and *exclusion* scores. We note that the inclusion score is equivalent to the 'absolute risk' of the class being associated with an entity in the given group. On the basis of those two scores, we can further define:

$$r\text{-score}(O_x, G_j) = 2 \cdot \frac{inclusion(O_x, G_j) \cdot exclusion(O_x, G_j)}{inclusion(O_x, G_j) + exclusion(O_x, G_j)} \quad (5)$$

In the case of our latter example with a large class imbalance, our inclusion score for pulmonary embolism is 0.04, while it is 0.001 for pneumonia. Our *r-scores* would then be 0.075 and 0.002, respectively. This provides a balanced metric between the internal representation of the class in the group, and its external exclusivity in the context of the larger dataset. Since both figures are taken into account, the *r-score* metric avoids extreme and potentially misleading values of z-score and odds ratio, which are also unhelpful for identification of characteristic classes.

In the case where the program is being used with a corpus that only describes a single group, the *r-score* measure is alternatively defined as equivalent to *inclusion* (note that the program can be forced into this mode, see 'exclusive group loading' later). In this case, the algorithm does not consider exclusion at all, and only provides solutions in the context of *inclusion* – describing the internal composition of the group.

To test our third quality, we introduce a measurement of how informative or specific a candidate class is. Since this score does not rely on any relationship between groups and entities, it is calculated only once per ontology class:

$$specificity(O_x) = IC(O_x) \quad (6)$$

Specificity is defined as the result of a particular information content measure. These provide an information-theoretic numerical measure for how informative a particular class in an ontology is, and are provided through the Semantic Measures Library [42], which implements many different information content measures. Currently, a choice between the Zhou [43] and Resnik [44] measures is available, which are defined as:

$$IC_{Resnik}(x) = -\log\left(\frac{|I(x)|}{|I(T)|}\right) IC_{Zhou}(x) = \frac{\log(deep(x))}{\log(deep_{max})} \quad (7)$$

The Resnik measure is defined as the reverse log probability of the class appearing in the corpus, and will therefore be greater for classes that are annotated to entities more infrequently. For example, using the Resnik measure, the *specificity* would be greater for arrhythmia (HP:0011675) would be greater the less frequently it appears in phenotype profiles associated with any group throughout a corpus. Meanwhile, the Zhou method is a measure of how deep the given class is when representing the ontology as a directed acyclic graph, calculated as a proportion of its maximum depth. Therefore, classes that are deeper in the ontology will have more intermediate subsuming classes between them and the root class, and as a result will receive greater values. For example, using this measure, supraventricular arrhythmia (HP:0005115) would receive a greater *specificity* score than arrhythmia (HP:0011675), since it is deeper in the ontology than the former. In all cases, the score for a single class would be equivalent, independent of which group is currently being considered. We chose these two frequently-used information content measures to provide the choice between a statistical method of measuring specificity (Resnik) and one dependent on the ontology structure (Zhou), with the best choice potentially differing depending on the dataset, ontology, or intended results [45]. These are also known as extrinsic and intrinsic measures, respectively.

Scoring

Now that we have defined a number of measures for the explanatory power of candidate discriminating classes, the next steps depend upon identifying and scoring candidate classes in the relevant biomedical ontology. At this stage, we load and classify the ontology, verifying

internal consistency and reflecting any structural inferences. Klarigi uses the ELK reasoner to perform this task [46], which supports the \mathcal{EL} subset of OWL. Since we sought to support a maximal number of ontologies, datasets, and settings, we chose ELK since it ensures maximum polynomial time classification with respect to the number of axioms in the ontology. Since Klarigi is implemented using OWLAPI [47], however, use with more expressive reasoners can be easily implemented, or ontologies can be pre-processed before passing to Klarigi.

While scores can be derived, theoretically, for all classes in the biomedical ontology, in practice this is not necessary. This is because only classes that subsume classes appearing in our corpus can receive inclusion and exclusion scores above zero. We can therefore define a new set of candidate terms C , the set of candidate terms in the ontology that we will score, as a subset of O :

$$C = \{x \in O \mid \exists z \in E \text{ member}(\text{subclass}(O_x), E_z)\} \quad (8)$$

In practice, this is implemented more efficiently by iterating directly all classes appearing explicitly in the annotation corpus and identifying their transitive superclasses. Once C is determined, the scores are calculated as above for each combination of group and candidate class, excepting for specificity where the score is calculated only once per class.

Candidate restriction and univariate analysis

We have now created the set of C candidate explanatory classes, which consists of the set of all classes that either directly annotate, or subsume classes that directly annotate, all entities in the corpus. The next step is to identify characteristic candidate variables for each group of interest. We do this by creating a new set U_j for each group:

$$U_j = \{x \in C \mid r\text{-score}(O_x, G_j) \geq \text{min-}r\text{-score} \wedge \text{inclusion}(O_x, G_j) \geq \text{min-inclusion} \\ \wedge \text{exclusion}(O_x, G_j) \geq \text{min-exclusion} \wedge \text{specificity}(O_x) \geq \text{min-ic}\} \quad (9)$$

This new subset U_j can be determined for each group, and contains a subset of the C candidate classes that meet the set of minimum restrictions for each score in relation to the considered group: *r-score*, inclusion, exclusion, and specificity. The process up to this stage consists the univariate mode of operation for Klarigi. At this point, the set U_j can be output for a group of interest, and subsequently interpreted. The minimum cut-offs for scores are controlled by configurable parameters, described in Table 1.

Multivariable explanatory sets

In addition to the univariate mode of operation, we also determine an additional method to identify sets of candidate terms that together characterise the group of interest. To do this, we can consider a solution space S , where S_m is any potential subset of U_j , the set of candidate classes that meet the minimum cut-offs for the univariate scores. Our goal is to identify such an S_m that is a ‘good’ characterisation of entities in the group. To identify what constitutes a good characterisation of the group, we turn to our previously defined scores. For the purposes of our solution, we only consider *r-score* and *specificity* to evaluate the explanatory power of individual terms, relying on the ability of *r-score* to balance between *inclusion* and *exclusion* scores. To evaluate the overall fitness of a set of terms S_m , we define an additional heuristic:

$$\text{overall_inclusion}(S_c, G_j) = \frac{\sum_{z=1}^{|E|} \text{member}(\text{subclass}(S_c), E_z)}{|G_j|} \quad (10)$$

This is a measure of overall inclusion, the proportion of entities in the group of interest that are annotated by at least one of the classes in the proposed solution, or their subclasses. This measures how well a set of candidate classes covers the full set of entities in the group.

Table 1

Names, descriptions, and default parameters for candidate class restrictions. These parameters define minimum values for the scoring heuristics, which restrict the set of candidate variables that will appear in the univariate analysis output, and be considered in the multivariable analysis stage. These default values were set based upon our observations using the algorithm, and are configurable by the user.

Parameter	Default	Description
min-inclusion	0	Candidate terms with inclusion values below this level will not appear in, or contribute to, explanatory sets.
min-exclusion	0.05	Candidate terms with exclusion values below this level will not appear in, or contribute to, explanatory sets.
min-r-score	0.1	Candidate terms with <i>r-score</i> values below this level will not appear in, or contribute to, explanatory sets.
min-ic	0.4	Candidate terms with information content values below this level will not appear in, or contribute to, explanatory sets.

The challenge of obtaining a good solution lies with the optimisation of several variables, and can therefore be considered as a multiple objective optimisation problem, considering the scoring heuristics as objective functions. The objective functions can be defined as the cut-offs for our scores. In the case of scores that measure the explanatory power of individual terms, this is the cut-off for candidate terms to appear in S , while the *overall_inclusion* function can be defined as the cut-off for S to be considered acceptable.

This can be considered in terms of the ϵ -constraints solution to multiple objective optimisation problems. In these solutions, one objective function is retained, while the rest are transformed into a set of constraints between which the remaining objective function can be optimised [48]. Our solution is inspired by this approach, selecting *overall_inclusion* as the objective function. However, instead of using static constraints for the other cutoffs, we develop an algorithm that steps down through acceptable values of these constraints in a priority order to dynamically identify high values for constraints in the context of objective function optimisation.

To do this, we define an order of priority for parameters:

1. *overall_inclusion*
2. *r-score*
3. *specificity*

The objective function can be considered as the highest priority constraint (and is also one that has no lower boundary). To identify a solution that maximises *overall_inclusion*, while optimising the other values within their configured boundaries, the algorithm steps down through acceptable values of each constraint in order of priority, from lowest to highest. *overall_inclusion* is chosen as the objective function because it is the only measure that encodes the overall characteristic quality of a set of candidate terms, while *r-score* is placed second due to its encoding of the more specific compositional and discriminatory information concerning the relationships between classes and groups, while *specificity* is retained at the lowest priority, as it provides additional background information concerning the terms, which in most cases will be less directly relevant to characterisation of the groups than the other measures.

A satisfactory solution is defined as a set of ontology terms that meets a current minimum value for *overall_inclusion*, and in which every member meets the current minimum constraint for *r-score* and *specificity*. Upon each step of the algorithm, the subset of classes that meet the current individual criteria is identified, and if that set meets the current cutoff for *overall_inclusion*, this is returned as the solution. If not, the current constraint settings are stepped down in order of priority, with a new check for a satisfactory solution at each step. First, the cut-off for *r-score* is reduced stepwise to its lower limit, at

which point the cut-off for *specificity* will be reduced by one step, and the cut-off for *r-score* will be reset. Once both *r-score* and *specificity* reach their lower limits, they are reset to their original values, and the *overall_inclusion* cut-off will be reduced by one step. This process is repeated until a satisfactory solution is found, checked every time a constraint is changed, or until the *overall_inclusion* cutoff reaches zero (a null response). The result of this process is a set of terms S_{good} , a subset of C , that both maximises the value of the objective function, and maximises values of the constraints according to the order of priority. The algorithm is given in Algorithm 1. The time complexity of the stepdown algorithm is linear with respect to the number of candidate classes in U_j ; every class in U_j will be examined once each time the constraints are reduced. In the worst case, where no solution is found, this will be $((top - total - inclusion \cdot step) \cdot (((top - r - score - bot - r - score) \cdot step) \cdot ((top - ic - bot - ic) \cdot step)))$ times.

Data: Refer to Table 2.

Result: A set of ontology terms that explanatory of the cluster.

specificityCutoff = **top-ic**

r-scoreCutoff = **top-r-score**

totalInclusionCutoff = **top-total-inclusion**

while True do

```

     $S_p = \{x \in U_j \mid r\text{-score}(O_x, G_j) \geq r\text{ScoreCutoff} \wedge \text{specificity}(O_x) \geq \text{specificityCutoff}\}$ 
     $S_{pm} = \{x \in S_p \mid r\text{-score}(O_x, G_j) \geq \text{max-r-score} \wedge \text{specificity}(O_x) \geq \text{max-specificity} \wedge \text{inclusion}(O_x, G_j) \geq \text{max-inclusion} \wedge \text{exclusion}(O_x, G_j) \geq \text{max-exclusion}\}$ 
    if  $\text{overall\_inclusion}(S_{pm}, G_j) < \text{totalInclusionCutoff}$  then
        if  $r\text{ScoreCutoff} \leq \text{bot} - r\text{Score}$  then
             $r\text{ScoreCutoff} = \text{top} - r\text{Score}$ 
            if  $\text{specificityCutoff} > \text{bot} - ic$  then
                 $\text{specificityCutoff} = \text{specificityCutoff} - \text{step}$ 
            else
                 $\text{specificityCutoff} = \text{top} - ic$ 
                 $\text{totalInclusionCutoff} = \text{totalInclusionCutoff} - \text{step}$ 
             $r\text{ScoreCutoff} = r\text{ScoreCutoff} - \text{step}$ 
        return  $S_p$ 

```

Algorithm 1: Algorithm for identifying characterising ontology terms for a cluster, by stepping down through *r-score* and *specificity* thresholds.

The algorithm also supports maximal constraints, that are used to define the maximum score at which a candidate term can contribute to the *overall_inclusion* score. This is implemented in Algorithm 1 as the determination of the S_{pm} subset of the proposed solution S_p . By default, they are set to values that will not be triggered (shown in Table 2), but they can be configured to prevent the result set being dominated by a small number of highly explanatory classes, where a larger and more varied set is desired. For example, the appearance of the term hypertension (HP:0008071) at 0.95 inclusion for a group defined by a diagnosis of hypertension may not be very informative. Thus setting an upper limit to score values that can contribute to *overall_inclusion* will force the algorithm to continue to step down to find additional explanatory terms. Since these terms still characterise the group, however, they are still included in the explanatory set, and are therefore included in the output (and included in the final *overall_inclusion* value associated with the output). These upper parameters can be used to encourage the algorithm to seek larger, less monolithic explanatory sets.

While the upper and lower boundaries for cut-offs are configurable, reasonable defaults have been set based on our observations using the algorithm. These can be defined for any of the term-wise scores: *r-score*, *inclusion*, *exclusion*, and *specificity*.

Discrimination and significance

On top of the univariate and multivariate modes of producing explanatory sets of classes for groups, we have also implemented several

Table 2

Descriptions and default values for multivariable parameters for the stepdown algorithm. These default values were set based upon our observations using the algorithm, and are configurable by the user.

Parameter	Default	Description
max-exclusion	1	Candidate terms with <i>r-score</i> values above this level will not contribute to calculation of overall inclusion, though they will still appear in explanatory sets.
max-inclusion	1	Candidate terms with <i>r-score</i> values above this level will not contribute to calculation of overall inclusion, though they will still appear in explanatory sets.
max-r-score	1	Candidate terms with <i>r-score</i> values above this level will not contribute to calculation of overall inclusion, though they will still appear in explanatory sets.
top-ic	0.7	The maximum stepdown value for IC.
bot-ic	0.4	The minimum stepdown value for IC.
top-r-score	0.8	The maximum stepdown value for <i>r-score</i> .
bot-r-score	0.1	The minimum stepdown value for <i>r-score</i> .
top-total-inclusion	1	Maximum stepdown value for overall inclusion.
step	0.05	Cutoffs will be reduced by this step on each iteration of the stepdown algorithm.

additional methods by which the fitness of solutions can be evaluated. To identify the discriminatory power of a set of explanatory classes, we can define a function that will, for each combination of entity and group in the corpus, produce a score:

$$\text{score}(E_i, G_j, S_l) = \prod_{x=1}^{|S_l|} (1 + (\text{exclusion}(O_x, G_j) \cdot \text{member}(\text{subclass}(O_x), E_i))) \quad (11)$$

This score can either be calculated on the results of the multivariable stepdown approach, the full set of univariate results, or a manually specified list of term associations. In the case that we are using the univariate mode of execution only, S_l will be equivalent to U_j . Given a proposed solution for a group, consisting of a set of classes, a score is calculated for each entity in the corpus, defined as the product of one plus the exclusion score of every class in the solution set that the entity is annotated with.

This definition of a predictive score allows us to construct a model to classify group membership for entities, from which an Area Under the receiver operating Characteristic (AUC) score can be calculated. This provides a simple measure of how well the set of explanatory variables, univariate or multivariate, was able to reclassify the set it was learned from. This can help to inform modifications to parameters or judge quality of the dataset in general.

The difference between these scores between the univariate and multivariate results can also provide insight on the quality of the modularisation performed by the stepdown algorithm. In the case that this value is small, there is a small amount of loss of total discrimination. Where values are large, there is a large decrease in performance when restricting scoring to the derived set of characteristic values, and this may be caused by a model that is poorly fit to the training data and does not reflect the full dataset. The program also contains the possibility to apply this classifier to unseen data, calculating an AUC in the same way. This provides the ability to perform a test-set validation, or even external validation, to more properly judge the quality of the solution. Where this approach is applied to a new dataset, the *exclusion* score used is the one learned from the training corpus.

The second method of results evaluation is significance testing. In many instances, it is useful to access the type 1 error (false positive) rate to enable comparison with existing tools which provide *p*-values, and also to provide a conservative aid to interpreting inclusion

and exclusion results. We implemented standard Monte Carlo methods to approximate the empirical p -value of the inclusion and exclusion statistics [49,50]. Briefly, we sample with replacement a set of class associations for each entity of the same size as its original association set.

A test statistic (inclusion/exclusion) is generated for that set of patient profiles. This procedure is repeated 1,000 times to create a vector of test statistics. The statistics are ranked, and the empirical p -value is obtained as:

$$\hat{p} = \frac{r+1}{n+1} \quad (12)$$

where r is the number of test statistics ranked greater or equal to the observed value, and n is the total number of permutations. The full algorithm is shown in Algorithm 2.

Data: nPerm = Number of monte carlo permutations

AE = All unique classes in O that appear in E

Result: testDist = Sorted list of calculated test statistics, length = nPerm + 1

pValue = Calculated empirical p -value

for n in nPerm **do**

ET _{j} = |E|;

for p in E **do**

ET _{jp} = random_sample(AE, length(E _{p}));

testDist _{n} = score(ET);

testDist = rank(testDist);

r = length(testStat \geq testDist);

pValue = $\frac{r+1}{nPerm+1}$;

return pValue;

Algorithm 2: Algorithm for generating p -values for inclusion and exclusion scores through permutation.

The number of permutations can be changed by the user, and we suggest at least 1,000 to ensure a (log-)normal distribution and sufficiently precise p -values possible (minimum of 0.001 in this case). Throughout the evaluations in this article, we use 1,000 permutations. Since we are creating a set of p -values for non-independent phenotypes, we suggest a Bonferroni threshold be applied to determine an adjusted significance threshold:

$$\theta \leq \frac{\alpha_{fwe}}{n} \quad (13)$$

where α_{fwe} is the desired error rate (0.05), n is the total number of phenotypes Klarigi generates p -values for per score. Thus θ is the new threshold for a significant inclusion/exclusion score.

Output, interpretation, and configuration

Klarigi reports its results in either the univariate or multivariate mode, in the form of a table of class associations with each group of interest, reported with the scores assigned to them. These can either be received in a \LaTeX table, a plain text table, or in Tab-Separated Values format. Each class is associated with four scores describing the relationship between the class and the group, which are given short descriptions in Table 3.

The program is highly configurable, and we have previously discussed the parameters that are available to modify the operation and results of the program both for candidate restriction (see Table 1) and identification of multivariable modules (see Table 2). In addition, there are additional parameters that control different modes of operation and calculation of scores. For example, there is a choice between which information content measure to use for the calculation of the *specificity* score. Descriptions of parameters are available in the documentation for the program. Klarigi can also be run in exclusive group loading (EGL) mode, which only loads the group of interest, and does not consider any others. In this case, the stepdown algorithm operates upon *inclusion*, and neither *exclusion* nor *r-score* are calculated. This can be useful

Table 3

Short description of Klarigi class scores and their interpretation.

Score	Description
Inclusion	Measures the extent to which this term composes the group of interest. Greater values indicate a greater proportion of entities in the group being annotated with this class or its subclasses.
Exclusion	Measurement of how discriminative this term is for the group of interest. Greater values indicate that this class is more exclusively associated with the group.
<i>r-score</i>	A balanced measure of how compositional and discriminative the class is for the group.
IC	The information content of the class. The greater the value, the more informative or specific the class is, according to the chosen information content measure.

for investigating group constituency without taking into account the exclusivity of terms, and in the manuscript we subsequently refer to this as a compositional analysis, in contrast to the discriminatory analysis that results from loading all groups.

Because of this configurability, and the dependence of results on that configurability, we anticipate a general workflow of using the application to consist of interactive modification of those parameters according to intermediate results, and desired outcomes. For example, if an empty set is returned upon analysis, the **min-ic** parameter could be reduced in an attempt to identify explanatory terms with a lower *specificity* than previously considered.

Results

We describe here the development and implementation of Klarigi, and make the tool freely available at <https://github.com/reality/klarigi>. This repository includes the source code, pre-compiled binaries, documentation, and a tutorial in notebook format that walks through the functionality of the software, explaining the significance of different parameters and outputs.

To demonstrate and evaluate the use of Klarigi for exploration of biomedical datasets, we developed two use-cases. Both describe clinical entities annotated using the Human Phenotype Ontology (HPO) [39]: in the first case text-derived phenotype profiles for a set of Medical Information Mart for Intensive Care III (MIMIC-III) admissions [51], and the second using a set of phenopackets [52], describing patients with rare diseases reported in literature [53]. We use Klarigi to explore both of these datasets, comparing and contrasting those results with the medical literature and enrichment analysis.

To further evaluate the results Klarigi produced for these use-cases, we also used an enrichment method to identify over-represented terms. To do this, we used XGR [29] to identify over-represented phenotypes using both one-tailed binomial and Fisher tests. P -values were adjusted with Bonferroni correction, propagating patient annotations to each term according to the True Path Rule. P -values were calculated with two background distributions: once using traits annotated to each patient and their frequencies as background, and additionally using only those annotated to the parent class of each entity tested; the maximum of each p -value was selected to help correct for the interdependence of testing. In the case of Klarigi, we also tested the normality of the permutation distribution for inclusion and exclusion scores by visually examining the distributions using histograms and QQ-plots. An R script containing these visualisations is provided in experiment repository.

The code used to obtain our results, as well as to perform the quantitative evaluations and create the tables for the paper, are available online at https://github.com/reality/klarigi_paper. All experiments were run on a regular computer running a Linux operating system with an Intel i7 processor and 16 GB of memory.

Use case: Pulmonary embolism

Pulmonary embolism (PE) is a condition in which a thrombus (usually a blood clot) forms in, or migrates from a distal site, to the pulmonary arteries and occludes the pulmonary circulation, blocking the entry of blood into or through the lungs. This may be associated with predisposing conditions such as cancer, cardiovascular disease or lung disease, but is often sporadic, for example following a deep vein thrombosis in the leg. This condition is associated with a considerable mortality rate, and often presents acutely in the emergency room in ways that render it difficult to diagnose when associated with other comorbidities, such as Chronic Obstructive Pulmonary Disease (COPD), and typically shares symptoms with other more common conditions, such as pneumonia and acute bronchitis [54]. The time-critical dependence of treatment on diagnosis makes it important to identify combinations of discriminating symptoms as rapidly as possible [55].

To demonstrate Klarigi's functionality, and to gain insight into the phenotypic presentations associated with pulmonary embolism and pneumonia, we created and evaluated text-derived phenotype profiles describing critical care admissions with the two diagnoses. We identified relevant admissions from MIMIC-III, and extracted phenotype profiles from their associated discharge letters. MIMIC-III is a large, freely-accessible dataset concerning nearly 60,000 visits to the emergency department at the Beth Israel Deaconess hospital [51].

We used the ICD-9 codes 486 and 415.19 to identify 5,437 MIMIC-III admissions with either pneumonia or pulmonary embolism as an associated diagnosis (provided in the coded diagnosis list produced by a coding expert): 4,853 with pneumonia and 912 with pulmonary embolism. We further restricted these to admissions with pneumonia or pulmonary embolism given as the primary diagnosis, and those which did not have the other diagnosis in any position. This led to a final set of 991 admissions, made up of 699 primarily coded with pneumonia, and 292 with pulmonary embolism. We employed the 2021-08-02 release of the Human Phenotype Ontology, available at <http://purl.obolibrary.org/obo/hp/releases/2021-08-02/hp.owl>.

To create phenotype profiles for these admissions, we used the Komenti semantic text-mining framework [56], which implements Stanford CoreNLP [57]. For every considered admission, we collected the text from the corresponding discharge note. We created a lemmatised vocabulary of all labels and synonyms in the Human Phenotype Ontology (HPO) [39], containing 50,265 labels for 16,019 unique classes. We used these to identify ontology terms mentioned in the discharge notes. We then removed all associations that were identified by Komenti as negated or uncertain. We also removed classes equivalent to, or subclasses of, pneumonia (HP:0002090), pulmonary embolism (HP:0002204), and abnormal thrombosis (HP:0001977). To facilitate a holdout validation, we then reserved a randomly sampled 20% of the annotated admissions. In the training set there are a total of 799 records, 568 pneumonia and 232 pulmonary embolism, while in the test set there were 190 records, 60 pulmonary embolism and 131 pneumonia.

In its univariate mode, Klarigi calculates all relevant metrics for every class that contains an instance in the dataset, restricted by the options given in Table 1. The results of this analysis are shown in Supplementary Tables 1 and 2. The subsequent results form different modular subsets of these univariate results, by using the stepdown algorithm with reference to different objective functions, as well as potentially restricting the initial set of univariate results by modifying minimum scoring thresholds.

We then employed Klarigi to perform a compositional analysis of both groups separately, using "exclusive group loading" (EGL) mode to consider each group without reference to its relationship to any other groups in the dataset. As such, only the inclusion and IC scores are used, and the step-down algorithm uses inclusion rather than *r-score* as its optimisation parameter. The results of this analysis are shown in

Table 4

Compositional analysis performed for pneumonia and pulmonary embolism using Klarigi. This is intended to identify a set of variables that characterise each group separately. This is achieved by using Klarigi in Exclusive Group Loading (EGL) mode. As such, the program produces a set of compositional classes for each disease, without considering how or whether those classes can be used to distinguish between the two groups. These results are a subset of the univariate results (visible in Supplementary Tables 1 and 2) sufficient to meet the threshold for *overall_inclusion*. Resnik IC was used. *max-inclusion* was set to 0.7 for pneumonia, to avoid returning only dyspnea (HP:0002094) and abnormal breath sound (HP:0030829), which alone meet a sufficient *overall_inclusion*. *min-inclusion* was set to 0.3 for pulmonary embolism to reduce the size of the result set.

Pneumonia	Inclusion	IC
Dyspnea (HP:0002094)	0.8	0.83
Abnormal systemic blood pressure (HP:0030972)	0.76	0.76
Abnormal breath sound (HP:0030829)	0.69	0.8
Cough (HP:0012735)	0.67	0.87
Increased blood pressure (HP:0032263)	0.64	0.83
Hypertension (HP:0000822)	0.64	0.89
Abnormality of temperature regulation (HP:0004370)	0.48	0.77
Fever (HP:0001945)	0.48	0.83
<i>Overall</i>	99.3	–
Pulmonary embolism	Inclusion	IC
Pain (HP:0012531)	0.81	0.63
Abnormal pattern of respiration (HP:0002793)	0.77	0.7
Abnormal systemic blood pressure (HP:0030972)	0.76	0.76
Dyspnea (HP:0002094)	0.75	0.83
Increased blood pressure (HP:0032263)	0.66	0.83
Hypertension (HP:0000822)	0.66	0.89
Abnormality of fluid regulation (HP:0011032)	0.63	0.67
Edema (HP:0000969)	0.62	0.67
Arrhythmia (HP:0011675)	0.61	0.64
<i>Overall</i>	100.0	–

Table 4. For both conditions, large values of the overall inclusion score show that the given explanatory sets describe almost all of the group. Large values for *inclusion* on individual classes also indicate terms that cover a large proportion of the relevant group.

To gain insight on the differences between the phenotypic composition of the two groups, we also employed Klarigi to perform a discriminative analysis. In this case, it uses the opposing group as context, providing the additional *exclusion* and *r-score* measures, using the latter as the optimisation parameter for application of the stepdown algorithm. These results are shown in Table 5. We tested different hyper-parameter options based on the reclassification AUC on the training set, but finally used the default values, obtaining final values of 0.903 and 0.996 for pulmonary embolism and pneumonia respectively. Permutation testing was performed after hyper-parameter optimisation, and so was only run once. Meanwhile, the results of the enrichment analysis using both the Fisher and binomial tests are shown in Supplementary Table 3.

Scores for *inclusion* measure the proportion of entities in the group – in this case PE or pneumonia – that are annotated by the class. Scores for *exclusion* indicate the extent to which an individual annotated to that class is likely to belong to one group rather than another, e.g. more likely to have pneumonia or PE, or vice versa. This is therefore a set of discriminating features that can be used to allocate a patient to one group or the other.

To evaluate how characteristic and discriminatory the explanatory sets were of their respective diseases, we used the 20% holdout set to define a classification task. The results of this analysis are listed in Table 6. In this analysis, we used the sets of characteristic terms identified by Klarigi in univariate mode, Klarigi in multivariate mode, and enrichment. Since the binomial enrichment results were a subset of the Fisher results, we combine these into one 'enrichment' set of explanatory terms. We then constructed classifiers for pulmonary embolism and pneumonia in the test set with scores for each entity identified using Klarigi. To provide additional measures of evaluation,

Table 5

Results for the discriminatory analysis performed using Klarigi. Here, we are attempting to identify an exclusively compositional set of explanatory classes for each group, containing classes that can be used to discriminate between the two groups, introducing the *exclusion* and *r-score* measures to achieve this. These results are contrasted with the purely compositional analysis seen in Table 4, since it forms an alternative subset of the univariate results shown in Supplemental Tables 1 and 2.

Pneumonia	<i>r-score</i>	Inclusion	Exclusion	IC
Cough (HP:0012735)	0.25	0.67	0.16 (p=0.051)	0.87
Congestive heart failure (HP:0001635)	0.22	0.3	0.17 (p=0.105)	0.95
Respiratory distress (HP:0002098)	0.21	0.34	0.15 (p=0.151)	0.92
Respiratory insufficiency (HP:0002093)	0.2	0.25	0.16 (p=0.042)	0.81
Airway obstruction (HP:0006536)	0.2	0.31	0.14 (p=0.068)	0.87
Productive cough (HP:0031245)	0.19	0.16	0.25 (p=0.013)	1.0
Chronic pulmonary obstruction (HP:0006510)	0.19	0.31	0.14 (p=0.157)	1.0
Renal insufficiency (HP:0000083)	0.19	0.37	0.13 (p=0.052)	0.84
Respiratory failure (HP:0002878)	0.19	0.22	0.16 (p=0.117)	1.0
Chills (HP:0025143)	0.19	0.27	0.14 (p=0.173)	1.0
Wheezing (HP:0030828)	0.19	0.38	0.12 (p=0.191)	1.0
Pulmonary edema (HP:0100598)	0.18	0.22	0.16 (p=0.141)	1.0
Aspiration (HP:0002835)	0.18	0.19	0.17 (p=0.102)	1.0
Fever (HP:0001945)	0.18	0.48	0.11 (p=0.267)	0.83
Atrial fibrillation (HP:0005110)	0.17	0.27	0.13 (p=0.216)	0.95
Abnormality of temperature regulation (HP:0004370)	0.17	0.48	0.1 (p=0.289)	0.77
Confusion (HP:0001289)	0.16	0.17	0.16 (p=0.114)	1.0
Rhonchi (HP:0030831)	0.16	0.2	0.14 (p=0.171)	1.0
Atrial arrhythmia (HP:0001692)	0.16	0.28	0.11 (p=0.166)	0.83
Chronic kidney disease (HP:0012622)	0.16	0.15	0.17 (p=0.027)	0.87
Stage 5 chronic kidney disease (HP:0003774)	0.15	0.11	0.22 (p=0.047)	1.0
Reduced consciousness/confusion (HP:0004372)	0.15	0.24	0.11 (p=0.254)	0.8
<i>Overall</i>	–	99.3	–	–
<hr/>				
Pulmonary embolism	<i>r-score</i>	Inclusion	Exclusion	IC
Sinus tachycardia (HP:0011703)	0.17	0.25	0.13 (p=0.057)	1.0
Increased body weight (HP:0004324)	0.17	0.16	0.18 (p=0.023)	0.84
Obesity (HP:0001513)	0.16	0.14	0.17 (p=0.028)	0.87
Lower limb pain (HP:0012514)	0.15	0.1	0.28 (p=0.005)	0.89
Abnormal electrophysiology of sinoatrial node origin (HP:0011702)	0.15	0.26	0.11 (p=0.097)	0.89
Limb pain (HP:0009763)	0.15	0.1	0.27 (p=0.005)	0.84
Pleuritic chest pain (HP:0033771)	0.15	0.12	0.19 (p=0.025)	1.0
Abnormality of body weight (HP:0004323)	0.14	0.21	0.11 (p=0.098)	0.73
Palpitations (HP:0001962)	0.14	0.1	0.22 (p=0.015)	1.0
Vascular dilatation (HP:0002617)	0.13	0.1	0.19 (p=0.035)	0.74
Syncope (HP:0001279)	0.12	0.07	0.28 (p=0.005)	0.89
Tachycardia (HP:0001649)	0.11	0.29	0.06 (p=0.166)	0.74
Arthritis (HP:0001369)	0.11	0.22	0.07 (p=0.155)	0.74
Hypercoagulability (HP:0100724)	0.1	0.06	0.71 (p=0.001)	1.0
Abnormal exteroceptive sensation (HP:0033747)	0.1	0.06	0.29 (p=0.001)	0.8
Somatic sensory dysfunction (HP:0003474)	0.1	0.06	0.29 (p=0.001)	0.75
<i>Overall</i>	–	80.6	–	–

Table 6

Test set classification results for pneumonia and pulmonary embolism, using Klarigi's classifier (see Eq. (11)), and semantic similarity using the Resnik IC and Resnik pairwise method, with results given for both best match average (SS-BMA) and average groupwise (SS-AVG) strategies. OI is the overall inclusion score, defined in Eq. (10), of the solution with respect to the test corpus. The greatest score for each category is emboldened. .95 confidence intervals are given for AUC values in parentheses.

Group	Method	AUC (Klarigi)	AUC (SS-BMA)	AUC (SS-AVG)	OI
Pneumonia	Enrichment	0.989 (0.976–1)	0.748 (0.677–0.82)	0.714 (0.636–0.792)	0.977
	Klarigi (u)	0.996 (0.989–1)	0.769 (0.699–0.834)	0.73 (0.654–0.807)	0.992
	Klarigi (mv)	0.992 (0.982–1)	0.771 (0.699–0.844)	0.752 (0.677–0.828)	0.985
Pulmonary embolism	Enrichment	0.575 (0.529–0.621)	0.543 (0.452–0.634)	0.561 (0.472–0.65)	0.15
	Klarigi (u)	0.858 (0.801–0.916)	0.631 (0.543–0.719)	0.675 (0.59–0.76)	0.716
	Klarigi (mv)	0.825 (0.764–0.886)	0.594 (0.505–0.684)	0.661 (0.576–0.747)	0.65

as well as to evaluate the utility of the exclusion score in measuring term discrimination, we also compared the Klarigi classification method with semantic similarity approaches. To do this, we also built classifiers scored using the groupwise similarity between each phenotype profile in the test set, and the explanatory set for each method. We performed this task for both the best match average and average groupwise strategies, with both cases using the Resnik pairwise method and Resnik information content measure (using the annotations in the training set for the probability distribution). For all classifiers, we calculated an AUC using the scores for each admission in the test set, and their actual diagnosis group for the label.

Use case: Rare disease annotation with phenopackets

Phenopackets are a standardised format for the representation of phenotypic descriptions of patients that use biomedical ontologies to annotate phenotypes. Phenopackets are increasingly being used as a standard format for exchanging, aggregating and analysing human disease information [52]. We developed Klarigi to natively support the phenopackets format, converting it internally into the required data model, which we believe is a forward-looking facility for the anticipated future use of phenopackets, for example as an export format for EHRs.

Table 7

The publications that report the 19 patients with IHPRF3 that are phenotyped in the phenopackets dataset.

Short name	Full title	Patients
Bhoj 2016 [59]	Mutations in TBCK, Encoding TBC1-Domain-Containing Kinase, Lead to a Recognisable Syndrome of Intellectual Disability and Hypotonia	10
Chong 2016 [60]	Recessive Inactivating Mutations in TBCK, Encoding a Rab GTPase-Activating Protein, Cause Severe Infantile Syndromic Encephalopathy	5
Guerreiro 2016 [61]	Mutation of TBCK causes a rare recessive developmental disorder	3
Zapata-Aldana 2019 [62]	Further delineation of TBCK - Infantile hypotonia with psychomotor retardation and characteristic facies type 3	2
Total	-	19

To further evaluate Klarigi in the context of this data representation method, we examined a dataset of 384 rare disease patients whose phenotypes were originally reported in the literature, and have subsequently had their phenotypes transcribed into the phenopacket format [53]. The dataset describes a range of rare diseases, and in our use case, we examine the disease with the most patient descriptions, Hypotonia, infantile, with psychomotor retardation and characteristic facies 3 (OMIM:616900) (IHPRF3), with 19 patients. IHPRF3 is an ultra-rare genetic disorder broadly characterised by global developmental delay and regression, brain atrophy, extreme hypotonia and a wide range of characteristic facial dysmorphologies [58]. It is caused by recessive mutations in the TBCK gene [59]. In the dataset we use, the patient descriptions are derived from four separate publications, listed in Table 7 were taken from these citations. The 2018-03-08 release of HPO was used to annotate these patients, and this was the version we used for the analysis.

Supplementary Table 5 shows the full set of 61 associated classes that constitute the univariate results. Meanwhile, Table 8 shows the multivariable results, whose overall inclusion describes complete coverage of the dataset with at least one of these classes, expressed in a much smaller explanatory set of 33 classes. Enrichment analysis was also performed using the Fisher and binomial tests, and the significant results of this analysis are shown in Supplementary Table 6.

To provide a quantitative evaluation for this analysis involving a small set of patients, we compared them with the disease annotations defined by the HPO database. HPO database annotations were developed through a combination of expert curation and text mining from literature, clinical descriptions, and experimental evidence [63]. There are 46 such annotations, listed in Supplementary Table 4. To measure how characteristic the result sets are, we calculated semantic similarity between each one and the set of HPO database annotations. We used the Resnik method of pairwise similarity, with the Zhou method of information content, with the best match average groupwise strategy. The results of the Fisher and Binomial tests are here treated separately, since neither forms a subset of the other.

Table 9 lists the semantic similarity value for each comparison.

We can also use Klarigi to evaluate alternative ways of grouping the data, or subsets of the data. In this use case, we explored 19 IHPRF3 patient data across different publications, shown in Table 7. In Supplementary Table 7 we show multivariable Klarigi results for IHPRF3 patients grouped by the publication in which they are reported.

Discussion

We have described the design and implementation of Klarigi, and its application to two use-cases, one to the comparison of two phenotypically similar diseases from text-derived phenotypes, and another to

rare disease patients from expert annotations created from literature. In this section, we will interpret and discuss the results, and provide more discussion of Klarigi in context, including limitations and future work.

The work described here is related to our previous work, which describes a semantic characterisation algorithm in the context of a cluster explanation method [38], realised by means of inclusion and exclusion scoring and a stepdown algorithm. This work builds upon it in several ways, redefining the *exclusion* score and introducing the *r-score*, improvement to the stepdown algorithm, a full formal description of the algorithm's design and implementation, introduction of permutation testing and the facility for producing models for discriminatory classification, configurability including modification of parameters and a new univariate mode of operation, and presentation of two practical use cases that compare it with enrichment methods. Moreover, the approach is consolidated into a software implementation that can be used on any compatible semantic dataset.

Interpretation of results

Pulmonary embolism

The large values of *overall_inclusion* calculated for the compositional results in Table 4 indicate a high overall coverage of the groups by the explanatory modules. There are many shared explanatory classes between the two diseases here, expressing their known similarity in presentation, as evidenced by previous compositional analysis of known presenting phenotypes for both conditions [54]. Half of the explanatory terms listed for pneumonia are shared by pulmonary embolism, often with very similar scores for *inclusion*, such as a difference of 0.05 for dyspnea (HP:0002094). Differences in terms for this set may provide candidates for discriminatory terms, such as cough (HP:0012735) and arrhythmia (HP:0011675). Indeed, the enrichment results include cough (HP:0012735) as a significantly over-represented class, being the only significant result produced by the binomial method.

The discriminatory analysis aims to produce explanatory sets that can be used to distinguish between the other groups in the dataset, shown in Table 5. We can see that none of the four redundant terms in the compositional results appear in either set of results, though non-redundant and highly compositional classes such as cough (HP:0012735) are preserved.

We see in the discriminatory analysis inclusion of classes describing known predisposing conditions for each disease. For example, pneumonia is a common complication of chronic renal disease [64] and PE is often a sequel to deep vein thrombosis [65] which is associated with lower limb pain, and obesity is similarly a known associated factor [66]. Tachycardia, hypercoagulability, and syncope are all established characteristics of PE, and it is notable that for pneumonia the most discriminating features are associated with respiratory difficulty; cough, wheezing, fever, and atrial arrhythmias. While fever may be found in some cases of PE it is clear that in our cohort this was not sufficiently common to contribute to the compositional analysis or to prevent it appearing as a discriminating class for pneumonia. The phenotype profiles derived from discharge letters do not contain information on D dimer levels [67], hence this frequently used measure did not appear in the results.

The discriminatory power of the sets is further explicated by the classification analysis, the results of which are given in Table 6. Reduction in performance for prediction of pneumonia with the Klarigi classifier was minimal, indicating that the set of explanatory terms generalised well to the phenotypic distribution of those patients in the test set. This reduction was greater for pulmonary embolism, which tracks the lower overall quality of this explanatory set evident on the training set, however the classifier still achieved a high performance.

For all classification methods and for both groups, models derived from the sets of explanatory terms generated by Klarigi outperformed

Table 8
Multivariable results for IHPRF3 in Klarigi. Reclassification AUC was 1.

OMIM:616900 (19 members)	<i>r</i> -score	Inclusion	Exclusion	IC
Severe muscular hypotonia (HP:0006829)	0.8	0.68 (p=0.001)	0.95 (p=0.009)	1.0
Developmental regression (HP:0002376)	0.63	0.58 (p=0.001)	0.68 (p=0.007)	1.0
Severe global developmental delay (HP:0011344)	0.5	0.42 (p=0.001)	0.62 (p=0.012)	1.0
Abnormality of upper lip vermilion (HP:0011339)	0.46	0.47 (p=0.001)	0.45 (p=0.007)	0.83
Respiratory insufficiency (HP:0002093)	0.42	0.42 (p=0.001)	0.42 (p=0.042)	0.81
Reduced tendon reflexes (HP:0001315)	0.42	0.53 (p=0.001)	0.35 (p=0.018)	0.82
Small basal ganglia (HP:0012697)	0.41	0.26 (p=0.001)	0.95 (p=0.007)	1.0
Exaggerated cupid's bow (HP:0002263)	0.41	0.26 (p=0.001)	0.95 (p=0.009)	1.0
Macroglossia (HP:0000158)	0.39	0.26 (p=0.001)	0.78 (p=0.01)	0.95
Areflexia (HP:0001284)	0.39	0.32 (p=0.001)	0.5 (p=0.012)	0.89
Profound global developmental delay (HP:0012736)	0.36	0.26 (p=0.001)	0.58 (p=0.012)	1.0
Prominent nasal bridge (HP:0000426)	0.36	0.26 (p=0.001)	0.58 (p=0.014)	1.0
Skeletal muscle hypertrophy (HP:0003712)	0.35	0.26 (p=0.001)	0.51 (p=0.013)	0.78
Sloping forehead (HP:0000340)	0.34	0.21 (p=0.001)	0.95 (p=0.006)	1.0
Aplasia/Hypoplasia of the cerebellar vermis (HP:0006817)	0.33	0.26 (p=0.001)	0.45 (p=0.044)	0.87
Highly arched eyebrow (HP:0002553)	0.33	0.26 (p=0.001)	0.45 (p=0.039)	1.0
Cerebellar vermis hypoplasia (HP:0001320)	0.33	0.26 (p=0.001)	0.45 (p=0.044)	0.95
Tented upper lip vermilion (HP:0010804)	0.31	0.21 (p=0.001)	0.62 (p=0.015)	1.0
Coarse facial features (HP:0000280)	0.29	0.26 (p=0.001)	0.34 (p=0.036)	1.0
Small forehead (HP:0000350)	0.29	0.26 (p=0.001)	0.34 (p=0.041)	1.0
Narrow forehead (HP:0000341)	0.29	0.26 (p=0.001)	0.34 (p=0.041)	1.0
Aplasia/Hypoplasia of the corpus callosum (HP:0007370)	0.29	0.37 (p=0.001)	0.24 (p=0.094)	0.95
Aplasia/Hypoplasia of the cerebral white matter (HP:0012429)	0.29	0.37 (p=0.001)	0.24 (p=0.094)	0.95
Hypoplasia of the corpus callosum (HP:0002079)	0.29	0.32 (p=0.001)	0.27 (p=0.094)	1.0
Cerebral white matter hypoplasia (HP:0012430)	0.29	0.32 (p=0.001)	0.27 (p=0.094)	1.0
Brachycephaly (HP:0000248)	0.29	0.21 (p=0.001)	0.45 (p=0.036)	0.95
Absent speech (HP:0001344)	0.28	0.37 (p=0.001)	0.23 (p=0.085)	1.0
Abnormality of the cerebellar vermis (HP:0002334)	0.27	0.26 (p=0.001)	0.28 (p=0.042)	0.8
Diffuse cerebellar atrophy (HP:0100275)	0.27	0.16 (p=0.001)	0.95 (p=0.006)	1.0
Global brain atrophy (HP:0002283)	0.27	0.16 (p=0.001)	0.95 (p=0.006)	1.0
Global developmental delay (HP:0001263)	0.26	0.95 (p=0.001)	0.15 (p=0.055)	0.89
Open mouth (HP:0000194)	0.26	0.16 (p=0.001)	0.7 (p=0.011)	1.0
Diffuse cerebral atrophy (HP:0002506)	0.26	0.16 (p=0.008)	0.7 (p=0.002)	1.0
Visual impairment (HP:0000505)	0.25	0.21 (p=0.001)	0.31 (p=0.052)	0.87
<i>Overall</i>	–	100.0	–	–

Table 9
Resnik+BMA+Zhou semantic similarity between the explanatory set for each method and the definitive set of IHPRF3 HPO annotations defined by the HPO annotations database.

Method	Similarity
Fisher	0.560
Binomial	0.607
Klarigi (u)	0.703
Klarigi (mv)	0.72

those from enrichment by AUC. This is very striking in the case of pulmonary embolism, where performance of the enrichment on the test set approached random by all classification methods, with poor generalisation also evidenced by the very low overall inclusion value on the holdout set. While differences in scores fell within .95 confidence intervals for all values within the same classification methods in the pneumonia setting, they were very distinct in the case of pulmonary embolism, indicating a significant result.

Another insight from the classification analysis on the test set is that the Klarigi classifier consistently performed substantially better than both semantic similarity approaches. AUC scores for the Klarigi classification method fell outside of confidence intervals in all cases except the universally poorly performing set of models formulated from the PE enrichment results. This indicates that the Klarigi method formed a significantly better performing set of classifiers in these settings. This helps to provide evidence that the exclusion score, which makes up the Klarigi classification model, is a good measure of class discrimination, that generalises to use for quantifying discriminatory sets.

Rare disease annotation with phenopackets

The phenopackets dataset is much smaller than the MIMIC dataset, describing far fewer patients. The annotations of patients were also

derived by an expert from literature reports, instead of being mined from text. The task also differs, in that we are evaluating IHPRF3 patients against a large background distribution of patients with other rare genetic diseases, instead of comparing their phenotype profiles head-to-head with those of a single chosen, and phenotypically similar disease. This difference is evidenced by the increased clarity of the results, with large individual values for inclusion and exclusion, overall inclusion, and reclassification AUC. Klarigi was also able to identify a much smaller subset of the univariate sets during the multivariate stage, and this provides an example of the identification of sparse representations of entity groups, in the case that the representation of terms across the group is sufficiently homogenous.

The small number of disease cases also informed the decision not to perform a train-test validation for this use case. Instead, explanatory sets were compared with semantic similarity to the HPO database annotations for the disease. In this evaluation, Klarigi's result sets were shown to be more similar to the definitional annotations than were both enrichment approaches.

We can also use this dataset to identify relationships that are not represented in the structured description of the disease. Several classes identified by the multivariate results, described in Table 8, are not found in the HPO database annotations. For instance, brachycephaly (HP:0000248) is not included, though it affects 21% of described patients with IHPRF3. As such, the findings produced by Klarigi could constitute new disease associations for the disease, not currently expressed in existing scientific databases describing it. In addition, of the HPO database annotations that are shared by Klarigi's representation, we provide additional information about patient representation, information content, and how exclusive these phenotypes are for patients across a wide set of rare diseases, all information that can be used to enrich secondary use of these phenotype associations.

Supplementary Table 8 shows the multivariable Klarigi results for IHPRF3 when stratified by the publication they appear in. This analysis allows us to identify features that may be more specific to particular publications and the patient groups they report. For example, intellectual disability (HP:0001249) is over-represented in the ten patients described by Bhoj et al., and both listed phenotypes for the three siblings in Guerreiro et al. are highly exclusive and inclusive. The terms appearing in these sets are highly discriminatory, and reflect the genetic and phenotypic heterogeneity of the patient groups described. Interestingly, we can see from the metadata associated with the HPO database annotations, that some associations were directly evidenced by the OMIM clinical description, which in turn makes use of publications not represented in the phenopackets dataset, in particular Alazami et al..

Comparison to enrichment analysis

The results of Klarigi cannot easily be directly compared to enrichment analysis. Fundamentally, they solve different problems, since enrichment analysis methods only attempt to identify significantly associated classes, and usually only for univariate relationships. Meanwhile, Klarigi produces its results through class scoring and discrimination rather than a significance cut-off. Therefore, while our quantitative evaluations show that Klarigi out-performs enrichment in all cases, this primarily indicates the different problem that Klarigi solves: semantic explanation, or identification of characteristic and discriminatory sets of terms for an annotated entity group.

Our results can also be used to illustrate where enrichment does and does not perform well for semantic explanation. In the case of pneumonia, there are many individual classes with a strong and significant univariate association with the group, and this leads to a sizeable set of associated classes, which do function together as a discriminatory module, evident by the high performance on the test set, listed in Table 6. In the case of pulmonary embolism, however, the phenotypic distribution is more complex, and only a single term is identified. By contrast, Klarigi identifies a set of terms that, in combination, characterise the group. Here the individual significance of terms becomes less relevant, since they can nevertheless appear in a discriminative module, whose generalisability be checked through means other than a statistical test, such as through classification or measurement of *overall_inclusion*.

Since Klarigi does not identify classes using significance testing, users are able to reactively modify hyper-parameters to identify better explanatory modules for the group, retaining the facility for final confirmatory significance testing without risking multiple testing hazards. For example, if initial results are too general, or if the result set is dominated by a single class, the *max_inclusion* parameter can be modified. Furthermore, while Klarigi provides significance testing for individual classes, even where exclusion scores for particular classes may not be significant, it is highly likely that the combination of explanatory classes will be. The 'fit' of a result set to the dataset is provided by multiple measures: individual *p*-values, individual scores, overall scoring, and discriminatory AUC.

Klarigi's univariate mode is more comparable to enrichment, in that it identifies a set of pairwise relationships between terms and the group of interest. However, Klarigi provides more information about those relationships, and can therefore be used to more clearly explicate and interpret the dataset. For example, enrichment identified *congestive heart failure* (HP:0002098) as a highly over-represented class for pneumonia, with large values for odds ratio and z-score. This association is not incorrect, and in fact it is a strong discriminator, which also appears in the Klarigi solution. However, while its discriminatory power is high, it is not a very highly compositional term, with only 0.3 for *inclusion*. In the case of another highly associated class for pneumonia, *cough* (HP:0012735) appears to strongly characterise pneumonia, however this phenotype also appears with high prevalence in the pulmonary embolism group too, evidenced by the Klarigi-derived

inclusion score of 0.25, compared with 0.67 for pneumonia. In cases like this, large values for zscore and odds ratio may be confusing or difficult to interpret, and may not provide enough information to those attempting to gain a holistic understanding of each group and their relationship with each other. Meanwhile, Klarigi captures this combination of facts in the scores it presents for *congestive heart failure* (HP:0002098) and *cough* (HP:0012735).

Another distinguishing feature is Klarigi's use of an ontology reasoner. Most enrichment methods do not make use of a reasoner, working instead on purely asserted axioms, or in some cases though preferred knowledge graphs produced by preprocessing tools such as OntoFunc. Use of an ontology reasoner provides access to inferential reasoning and relationships between classes when identifying explanatory terms. For example, an additional inferred subclass relationship would affect *inclusion*, *exclusion*, and *specificity* scores as instances are determined across the ontology structure. Furthermore, use of an ontology reasoner confirms internal consistency of the ontology being used to perform the investigation.

Context, limitations, and future work

The design of Klarigi is a development upon principles first explored in our previous work, in which we developed an algorithm for deriving multivariable explanations for semantic clusters identified from Human Phenotype Ontology (HPO) phenotype profiles [38]. In this work, the approach is heavily modified and improved upon, including changes to the algorithm and scoring system, and is generalised to be applicable to any dataset and ontology. In particular, in the Design and Implementation section, we discussed limitations to the previous definition of the *exclusion* score, and created a new definition that is less sensitive to class imbalance.

While one feature of Klarigi is the ability to interactively modify hyperparameters to identify desirable result sets, these parameters also introduce complexity to the process, and understanding the interplay of the parameters and their effect on the results may require manual work, since it will differ depending on the dataset. Moreover, it may not be easily possible for a human operator to find the optimal parameter configuration.

The current defaults for algorithm parameters were determined based on our observations using the algorithm. While these are configurable, and worked well in our settings, they may not do so for all datasets, and their configuration can be time consuming. Grid search could also be employed to automatically identify good settings for parameters, using either the *overall_inclusion* measure or discriminatory AUC. Such a search through possible overall parameters could also add significantly to the execution time. However, overcoming this problem would also help to avoid the problem of providing parameter defaults that perform well across different ontologies, datasets, and algorithmic parameters, and reduce the amount of human input required to obtain quality results. Parameter optimisation could also potentially introduce problems with multiple testing. We controlled for this, however, by choosing not to employ permutation testing until the final set of desired results is received — Klarigi currently requires an additional parameter to be passed for significance testing to be performed. Simulation studies could also be undertaken to identify good default values for different indications, e.g. for different ontologies.

It is also possible that different scores could be introduced, or different definitions of existing scores could be explored. For example, the option to re-weight the *r-score*, in a similar manner to differential weight applications to precision and recall in an *f-score*. Likewise, different priorities of measures could be explored, such as considering information content above *r-score* in the stepdown algorithm, or replacing *overall_inclusion* with AUC. The program also currently only supports two information content measures, but could easily be

extended to support more, or to support the definition of customised measures of ontology-based *specificity*.

The information Klarigi produces about unique characterisation of rare diseases may also benefit development of disease-phenotype associations, which are used in a wide range of analysis approaches, such as rare disease variant prioritisation or diagnosis. This contribution could consist simply in the identification of additional phenotypes through analysis of either larger aggregated datasets similar to the phenopackets dataset, or text-mined data from clinical encounters or literature. Moreover, we showed in Supplementary Table 8 that we were able to identify disease phenotypes that were more exclusively associated with single publications. Taking note of this information could support development of high quality disease-phenotype relationships. For example, particular phenotypes that are very exclusive to one or few publications may be less preferred as candidates for associations, or associations could be provided with a measure of literature agreement. Such a dataset could be produced using a larger set of phenopackets data, or through mining literature and clinical narrative for phenotype profiles. Klarigi parameters could be configured to identify the necessary level of association desired to constitute a disease-phenotype link.

The measures described in our experiment could also be combined with enrichment analysis approaches. Enrichment analysis use a measure of *enrichment score*, which could comprise any of the *r-score*, inclusion, or exclusion scores given herein, then using the usual method of testing for significant effects. Indeed, the use of significance testing upon the final set of multivariable results can be considered as a kind of enrichment analysis. Other potential expansions of functionality could consist in extending the tool to facilitate a wider range of information content measures, involving multi-faceted semantic similarity, or identification of sub-groups through clustering.

Admissions sampling for the MIMIC dataset did not consider whether multiple admissions were attributed to the same patient, and this could potentially be a small source of bias. However, we do not believe it affects the conclusions of this study, since it would not favour a particular algorithm. In the same way, though unlikely, different publications described by the phenopackets dataset could potentially describe the same individuals. Future work could involve fully evaluating the use of Klarigi in a clinical classification setting, either standalone or as a feature selection methodology. We did not compare the approach with another multivariable enrichment method. This is because existing implementations for multivariable enrichment are either limited to GO, or require a complicated setup process. The possibility for Klarigi to be easily applied to any ontology and ontology-annotated dataset is a major benefit of the approach. We addressed this by making these methodological differences clear in our problem description, comparison of the results, and discussion. We anticipate that Klarigi can be used to perform improved semantic descriptive analysis of biomedical entities. This is an area of research that is of recent interest, especially with advancements in structured representation of phenotype profiles through phenopackets, and the automated creation of phenotypic profiles through text mining. For example, one study performed a statistical analysis of text-mined phenotype profiles for 4,095 individuals with Down Syndrome [69], descriptively reporting on phenotypes and their frequency of appearance. This analysis, however, only used a measure of patient frequency to stratify phenotypes. The use of Klarigi in this case would have introduced measures of representation for considered phenotypes, for example overall coverage of unique individuals, and the utility to easily switch groupings to explore the different aspects of the data.

Conclusions

Klarigi provides a new way to solve the task of semantic explanation, addressing limitations in traditional enrichment approaches. As such, it provides an approach that can be used for improved exploration

and interpretation of biomedical datasets linked to ontologies. We have demonstrated that it can be used to gain insight into relationships between biomedical entities with clinical relevance.

Ethical approval

This work makes use of the MIMIC-III dataset, which was approved for construction, de-identification, and sharing by the BIDMC and MIT institutional review boards (IRBs). Further details on MIMIC-III ethics are available from its original publication (DOI:10.1038/sdata.2016.35). Work was undertaken in accordance with the MIMIC-III guidelines. As such, while the code used to process MIMIC data is provided in the GitHub repository, the data itself is not made available.

Abbreviations

HPO: Human Phenotype Ontology, IHPRF3: Hypotonia, infantile, with psychomotor retardation and characteristic facies 3, OMIM: Online Mendelian Inheritance In Man, OWL: Web Ontology Language, GSEA: Gene Set Enrichment Analysis, EGL: Exclusive Group Loading, GO: Gene Ontology, AUC: Area Under receiver operating Characteristic, IC: Information Content, MIMIC-III: Medical Information Mart for Intensive Care III, ML: Machine Learning; ETL: Extract, Transform, Load.

CRedit authorship contribution statement

Luke T. Slater: Conceived of the study and experimental design, Developed novel methods, Performed the experiments, Implemented the software, Created initial manuscript draft. **John A. Williams:** Assisted experimental design, Manuscript development. **Paul N. Schofield:** Performed biological and clinical interpretation of results, Contributed to experimental design, Contributed to manuscript. **Sophie Russell:** Contributed to evaluation of the approach. **Samantha C. Pendleton:** Contributed to evaluation of the approach. **Andreas Karwath:** Contributed to experimental design, methods development, and results analysis. **Hilary Fanning:** Contributed to experimental design and project supervision. **Simon Ball:** Contributed to experimental design and project supervision. **Robert Hoehndorf:** Contributed to experimental design and manuscript. **Georgios V. Gkoutos:** Supervised project, Contributed to experimental design, Contributed to manuscript.

Availability of data and material

Klarigi is available under an open source licence from <https://github.com/reality/klarigi>. The code used to perform experiments is available from <https://github.com/reality/klarigi>.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.combiomed.2022.106425>. Includes additional univariate results and data.

References

- [1] R. Hoehndorf, P.N. Schofield, G.V. Gkoutos, The role of ontologies in biological and biomedical research: a functional perspective, *Brief. Bioinform.* 16 (6) (2015) 1069–1080, <http://dx.doi.org/10.1093/bib/bbv011>.
- [2] O. Bodenreider, Bio-ontologies: current trends and future directions, *Brief. Bioinform.* 7 (3) (2006) 256–274, <http://dx.doi.org/10.1093/bib/bbl027>.
- [3] M.A. Haendel, J.A. McMurry, R. Relevo, C.J. Mungall, P.N. Robinson, C.G. Chute, A census of disease ontologies, *Annu. Rev. Biomed. Data Sci.* 1 (1) (2018) 305–331, <http://dx.doi.org/10.1146/annurev-biodatasci-080917-013459>.
- [4] M. del Carmen Legaz-García, J.A. Miñarro-Giménez, M. Menárguez-Tortosa, J.T. Fernández-Breis, Generation of open biomedical datasets through ontology-driven transformation and integration processes, *J. Biomed. Semant.* 7 (1) (2016) 32, <http://dx.doi.org/10.1186/s13326-016-0075-z>.

- [5] A. Mahmoud, M. Y. Shams, O. M. Elzeki, N. Awadallah Awad, Using semantic web technologies to improve the extract transform load model, *Comput. Mater. Continua* 68 (2) (2021) 2711–2726, <http://dx.doi.org/10.32604/cmc.2021.015293>.
- [6] R. Hoehndorf, P.N. Schofield, G.V. Gkoutos, PhenomeNET: A whole-phenome approach to disease gene discovery, *Nucleic Acids Res.* 39 (18) (2011) e119, <http://dx.doi.org/10.1093/nar/gkr538>.
- [7] S. Köhler, M.H. Schulz, P. Krawitz, S. Bauer, S. Dölken, C.E. Ott, C. Mundlos, D. Horn, S. Mundlos, P.N. Robinson, Clinical diagnostics in human genetics with semantic similarity searches in ontologies, *Am. J. Hum. Genet.* 85 (4) (2009) 457–464, <http://dx.doi.org/10.1016/j.ajhg.2009.09.003>.
- [8] L.T. Slater, A. Karwath, J.A. Williams, S. Russell, S. Makepeace, A. Carberry, R. Hoehndorf, G.V. Gkoutos, Towards similarity-based differential diagnostics for common diseases, *Comput. Biol. Med.* 133 (2021) 104360, <http://dx.doi.org/10.1016/j.cbiomed.2021.104360>.
- [9] O.J. Veatch, J. Veenstra-Vanderweele, M. Potter, M.A. Pericak-Vance, J.L. Haines, Genetically meaningful phenotypic subgroups in autism spectrum disorders, *Genes Brain Behav.* 13 (3) (2014) 276–285, <http://dx.doi.org/10.1111/gbb.12117>.
- [10] M. Kulmanov, F.Z. Smaili, X. Gao, R. Hoehndorf, Machine learning with biomedical ontologies, 2020, <http://dx.doi.org/10.1101/2020.05.07.082164>, bioRxiv, 2020.05.07.082164.
- [11] E. Jermutus, D. Kneale, J. Thomas, S. Michie, Influences on User Trust in Healthcare Artificial Intelligence: A Systematic Review [version 1; peer review: 1 approved with reservations], *Wellcome Open Res.* 7 (65) (2022) <http://dx.doi.org/10.12688/wellcomeopenres.17550.1>.
- [12] B. Babic, S. Gerke, T. Evgeniou, I.G. Cohen, Beware explanations from AI in health care, *Science* 373 (6552) (2021) 284–286, <http://dx.doi.org/10.1126/science.abg1834>, URL <https://www.science.org/doi/abs/10.1126/science.abg1834>. arXiv:<https://www.science.org/doi/pdf/10.1126/science.abg1834>.
- [13] C.M. Cutillo, K.R. Sharma, L. Foschini, S. Kundu, M. Mackintosh, K.D. Mandl, T. Beck, E. Collier, C. Colvis, K. Gersing, V. Gordon, R. Jensen, B. Shabestari, N. Southall, M.I.i.H.W.W. Group, Machine intelligence in healthcare—perspectives on trustworthiness, explainability, usability, and transparency, *Npj. Digital Med.* 3 (1) (2020) 47, <http://dx.doi.org/10.1038/s41746-020-0254-2>.
- [14] R. Gentleman, M. Morgan, W. Huber, *Gene set enrichment analysis*, in: *Bioconductor Case Studies*, Springer, 2008, pp. 193–205.
- [15] Y. Deng, L. Gao, B. Wang, X. Guo, HPOSim: An R package for phenotypic similarity measure and enrichment analysis based on the human phenotype ontology, *PLOS ONE* 10 (2) (2015) e0115692, <http://dx.doi.org/10.1371/journal.pone.0115692>.
- [16] T. Wittkop, E. TerAvest, U.S. Evani, K.M. Fleisch, A.E. Berman, C. Powell, N.H. Shah, S.D. Mooney, STOP using just GO: A multi-ontology hypothesis generation tool for high throughput experimentation, *BMC Bioinformatics* 14 (2013) 53, <http://dx.doi.org/10.1186/1471-2105-14-53>.
- [17] K. Verspoor, *The gene ontology as a source of lexical semantic knowledge for a biological natural language processing application*, 2003.
- [18] G. Korotkevich, V. Sukhov, N. Budin, B. Shpak, M.N. Artyomov, A. Sergushichev, Fast gene set enrichment analysis, 2021, 060012, <http://dx.doi.org/10.1101/060012>.
- [19] M. Morgan, S. Falcon, R. Gentleman, GSEABase: Gene set enrichment data structures and methods, 2021, <http://dx.doi.org/10.18129/B9.bioc.GSEABase>, Bioconductor version: Release (3.14).
- [20] Y. Yang, F. Qian, X. Li, Y. Li, L. Zhou, Q. Wang, X. Zhou, J. Zhang, C. Song, Z. Yu, GREAP: A comprehensive enrichment analysis software for human genomic regions, *Brief. Bioinform.* 23 (5) (2022) bbac329.
- [21] G. Yu, L.-G. Wang, G.-R. Yan, Q.-Y. He, DOSE: An R/Bioconductor package for disease ontology semantic and enrichment analysis, *Bioinformatics* 31 (4) (2015) 608–609, <http://dx.doi.org/10.1093/bioinformatics/btu684>.
- [22] G. Yu, Q.-Y. He, ReactomePA: An R/Bioconductor package for reactome pathway analysis and visualization, *Mol. Biosyst.* 12 (2) (2016) 477–479, <http://dx.doi.org/10.1039/C5MB00663E>.
- [23] H. Liu, M. Yuan, R. Mitra, X. Zhou, M. Long, W. Lei, S. Zhou, Y.-e. Huang, F. Hou, C.M. Eischen, CTpathway: A CrossTalk-based pathway enrichment analysis method for cancer research, *Genome Med.* 14 (1) (2022) 1–20.
- [24] D. Chicco, G. Agapito, Nine quick tips for pathway enrichment analysis, *PLoS Comput. Biol.* 18 (8) (2022) e1010348.
- [25] P. Gollapalli, S. B. S. H. Rimac, P. Patil, S.K. Nalilu, S. Kandagalla, P. Shetty, Pathway enrichment analysis of virus-host interactome and prioritization of novel compounds targeting the spike glycoprotein receptor binding domain–human angiotensin-converting enzyme 2 interface to combat SARS-CoV-2, *J. Biomol. Struct. Dyn.* 40 (6) (2022) 2701–2714.
- [26] M.-P. Weng, B.-Y. Liao, MamPhEA: A web tool for mammalian phenotype enrichment analysis, *Genome Biol.* 11 (1) (2010) P27, <http://dx.doi.org/10.1186/gb-2010-11-s1-p27>.
- [27] B.T. Sherman, M. Hao, J. Qiu, X. Jiao, M.W. Baseler, H.C. Lane, T. Imamichi, W. Chang, DAVID: A web server for functional enrichment analysis and functional annotation of gene lists (2021 update), *Nucl. Acids Res.* 10 (2022).
- [28] U. Raudvere, L. Kolberg, I. Kuzmin, T. Arak, P. Adler, H. Peterson, J. Vilo, G:Profiler: A web server for functional enrichment analysis and conversions of gene lists (2019 update), *Nucleic Acids Res.* 47 (W1) (2019) W191–W198, <http://dx.doi.org/10.1093/nar/gkz369>.
- [29] H. Fang, B. Knezevic, K.L. Burnham, J.C. Knight, XGR software for enhanced interpretation of genomic summary data, illustrated by application to immunological traits, *Genome Med.* 8 (1) (2016) 129, <http://dx.doi.org/10.1186/s13073-016-0384-y>.
- [30] D. Greene, S. Richardson, E. Turro, Phenotype similarity regression for identifying the genetic determinants of rare diseases, *Am. J. Hum. Genet.* 98 (3) (2016) 490–499, <http://dx.doi.org/10.1016/j.ajhg.2016.01.008>.
- [31] K. Prüfer, B. Muetzel, H.-H. Do, G. Weiss, P. Khaitovich, E. Rahm, S. Pääbo, M. Lachmann, W. Enard, FUNC: A package for detecting significant associations between gene sets and ontological annotations, *BMC Bioinformatics* 8 (2007) 41, <http://dx.doi.org/10.1186/1471-2105-8-41>.
- [32] R. Hoehndorf, J.M. Hancock, N.W. Hardy, A.-M. Mallon, P.N. Schofield, G.V. Gkoutos, Analyzing gene expression data in mice with the Neuro Behavior Ontology, *Mamm. Genome* 25 (1) (2014) 32–40.
- [33] A. Hukku, C. Quick, F. Luca, R. Pique-Regi, X. Wen, BAGSE: A Bayesian hierarchical model approach for gene set enrichment analysis, *Bioinformatics* 36 (6) (2020) 1689–1695, <http://dx.doi.org/10.1093/bioinformatics/btz831>.
- [34] S. Zhang, J. Cao, Y.M. Kong, R.H. Scheuermann, GO-Bayes: Gene Ontology-based overrepresentation analysis using a Bayesian approach, *Bioinformatics* 26 (7) (2010) 905–911, <http://dx.doi.org/10.1093/bioinformatics/btq059>.
- [35] K.-L. Tiong, C.-H. Yeang, MGSEA – a multivariate Gene set enrichment analysis, *BMC Bioinformatics* 20 (1) (2019) 145, <http://dx.doi.org/10.1186/s12859-019-2716-6>.
- [36] J. Weiner 3rd, T. Domaszewska, Tmod: An R package for general and multivariate enrichment analysis, Preprint, PeerJ Preprints, 2016, <http://dx.doi.org/10.7287/peerj.preprints.2420v1>.
- [37] J.O.B. Jacobsen, M. Baudis, G.S. Baynam, J.S. Beckmann, S. Beltran, T.J. Callahan, C.G. Chute, M. Courtot, D. Danis, O. Elemento, R.R. Freimuth, M.A. Gargano, T. Groza, A. Hamosh, N.L. Harris, R. Kaliyaperumal, A. Khalifa, P.M. Krawitz, S. Köhler, B.J. Laraway, H. Levhväslaiho, K.C. Lloyd, L. Matalonga, J.A. McMurphy, A. Metke-Jimenez, C.J. Mungall, M.C. Munoz-Torres, S. Ogishima, A. Papakonstantinou, D. Piscia, N. Pontikos, N. Queralt-Rosinach, M. Roos, P.N. Schofield, A. Siapas, D. Smedley, L.D. Smith, R. Steinhaus, J.C. Sundaramurthi, E.M. Swietlik, S. Thun, N.A. Vasilevsky, A.H. Wagner, J.L. Warner, C. Weiland, M.A. Haendel, P.N. Robinson, The GA4GH Phenopacket schema: A computable representation of clinical data for precision medicine, 2021, <http://dx.doi.org/10.1101/2021.11.27.21266944>, medRxiv, 2021.11.27.21266944.
- [38] L.T. Slater, J.A. Williams, A. Karwath, H. Fanning, S. Ball, P. Schofield, R. Hoehndorf, G.V. Gkoutos, Multi-faceted Semantic Clustering With Text-derived Phenotypes, 2021, <http://dx.doi.org/10.1101/2021.05.26.21257830>, MedRxiv, 2021.05.26.21257830.
- [39] S. Köhler, M. Gargano, N. Matentzoglou, L.C. Carmody, D. Lewis-Smith, N.A. Vasilevsky, D. Danis, G. Balagura, G. Baynam, A.M. Brower, T.J. Callahan, C.G. Chute, J.L. Est, P.D. Galer, S. Ganesan, M. Griese, M. Haimel, J. Pazmandi, M. Hanauer, N.L. Harris, M.J. Hartnett, M. Hastreiter, F. Hauck, Y. He, T. Jeske, H. Kearney, G. Kindle, C. Klein, K. Knoflach, R. Krause, D. Lagorce, J.A. McMurphy, J.A. Miller, M.C. Munoz-Torres, R.L. Peters, C.K. Rapp, A.M. Rath, S.A. Rind, A.Z. Rosenberg, M.M. Segal, M.G. Seidel, D. Smedley, T. Talmy, Y. Thomas, S.A. Wiafe, J. Xian, Z. Yüksel, I. Helbig, C.J. Mungall, M.A. Haendel, P.N. Robinson, The human phenotype ontology in 2021, *Nucleic Acids Res.* 49 (D1) (2020) D1207–D1217, <http://dx.doi.org/10.1093/nar/gkaa1043>.
- [40] H.T.O. Davies, I.K. Crombie, M. Tavakoli, When can odds ratios mislead? *Br. Med. J.* 316 (7136) (1998) 989–991.
- [41] S. Tenny, M.R. Hoffman, *Relative Risk*, in: *StatPearls*, StatPearls Publishing, Treasure Island (FL), 2022.
- [42] S. Harispe, S. Ranwez, S. Janaqi, J. Montmain, The semantic measures library and toolkit: Fast computation of semantic similarity and relatedness using biomedical ontologies, *Bioinformatics* 30 (5) (2014) 740–742, <http://dx.doi.org/10.1093/bioinformatics/btt581>.
- [43] Z. Zhou, Y. Wang, J. Gu, A new model of information content for semantic similarity in WordNet, in: 2008 Second International Conference on Future Generation Communication and Networking Symposia, 3, IEEE, 2008, pp. 85–89.
- [44] P. Resnik, Using information content to evaluate semantic similarity in a taxonomy, 1995, arXiv Preprint [cmp-lg/9511007](https://arxiv.org/abs/9511007), arXiv:[cmp-lg/9511007](https://arxiv.org/abs/9511007).
- [45] C. Pesquita, D. Faria, H. Bastos, A. Falcao, F. Couto, Evaluating GO-based semantic similarity measures, in: *Proc. 10th Annual Bio-Ontologies Meeting*, 37, Citeseer, 2007, p. 38.
- [46] Y. Kazakov, M. Kröttsch, F. Simančík, The incredible ELK, *J. Automat. Reason.* 53 (1) (2014) 1–61, <http://dx.doi.org/10.1007/s10817-013-9296-3>.
- [47] M. Horridge, S. Bechhofer, The OWL API: A Java API for OWL ontologies, *Semantic Web* 2 (1) (2011) 11–21, <http://dx.doi.org/10.3233/SW-2011-0025>.
- [48] Y. Haimes, On a bicriterion formulation of the problems of integrated system identification and system optimization, *IEEE Trans. Syst. Man Cybern.* 1 (3) (1971) 296–297.
- [49] B. Phipson, G.K. Smyth, Permutation *P*-values should never be zero: Calculating exact *P*-values when permutations are randomly drawn, *Stat. Appl. Genet. Mol. Biol.* 9 (2010) 39, <http://dx.doi.org/10.2202/1544-6115.1585>.

- [50] B.V. North, D. Curtis, P.C. Sham, A note on the calculation of empirical P values from Monte Carlo procedures, *Am. J. Hum. Genet.* 71 (2) (2002) 439–441.
- [51] A.E.W. Johnson, T.J. Pollard, L. Shen, L.-w.H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L.A. Celi, R.G. Mark, MIMIC-III, a freely accessible critical care database, *Sci. Data* 3 (1) (2016) 1–9, <http://dx.doi.org/10.1038/sdata.2016.35>.
- [52] J.O.B. Jacobsen, M. Baudis, G.S. Baynam, J.S. Beckmann, S. Beltran, O.J. Buske, T.J. Callahan, C.G. Chute, M. Courtot, D. Danis, O. Elemento, A. Essenwanger, R.R. Freimuth, M.A. Gargano, T. Groza, A. Hamosh, N.L. Harris, R. Kaliyaperumal, K.C.K. Lloyd, A. Khalifa, P.M. Krawitz, S. Köhler, B.J. Laraway, H. Lehväslaiho, L. Matalonga, J.A. McMurry, A. Metke-Jimenez, C.J. Mungall, M.C. Munoz-Torres, S. Ogishima, A. Papakonstantinou, D. Piscia, N. Pontikos, N. Queralt-Rosinach, M. Roos, J. Sass, P.N. Schofield, D. Seelow, A. Siapos, D. Smedley, L.D. Smith, R. Steinhaus, J.C. Sundaramurthi, E.M. Swietlik, S. Thun, N.A. Vasilevsky, A.H. Wagner, J.L. Warner, C. Weiland, M. Axton, L. Babb, C.F. Boerkoel, B.P. Chaudhari, H.-L. Chin, M. Dumontier, N. Gazzaz, D.P. Hansen, H. Hochheiser, V.A. Kinsler, H. Lochmüller, A.R. Mankovich, G.I. Saunders, P.I. Sergouniotis, R. Thompson, A. Zankl, M.A. Haendel, P.N.R. and, The GA4GH Phenopacket schema defines a computable representation of clinical data, *Nature Biotechnol.* 40 (6) (2022) 817–820, <http://dx.doi.org/10.1038/s41587-022-01357-4>.
- [53] P.N. Robinson, 384 Phenopackets, 2020, <http://dx.doi.org/10.5281/zenodo.3905420>.
- [54] M. Papparoupa, L. Spineli, T. Framke, H. Ho, F. Schuppert, A. Gillissen, Pulmonary embolism in pneumonia: Still a diagnostic challenge? Results of a case-control study in 100 patients, *Dis. Markers* 2016 (2016) 1–8, <http://dx.doi.org/10.1155/2016/8682506>.
- [55] J. Bělohávek, V. Dytrych, A. Linhart, Pulmonary embolism, part I: Epidemiology, risk factors and risk stratification, pathophysiology, clinical presentation, diagnosis and nonthrombotic pulmonary embolism, *Exp. Clin. Cardiol.* 18 (2) (2013) 129–138.
- [56] L.T. Slater, W. Bradlow, R. Hoehndorf, D.F. Motti, S. Ball, G.V. Gkoutos, Koment: A semantic text mining framework, 2020, <http://dx.doi.org/10.1101/2020.08.04.233049>, bioRxiv 2020.08.04.233049.
- [57] C.D. Manning, M. Surdeanu, J. Bauer, J.R. Finkel, S. Bethard, D. McClosky, The Stanford CoreNLP Natural Language Processing Toolkit, in: Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations, 2014, pp. 55–60, <http://dx.doi.org/10.3115/v1/P14-5010>.
- [58] A. Smith, K. Leask, P. Tomlin, D. Donnai, A familial dysmorphic condition with hypotonia, seizures and precocious puberty, *Clin. Dysmorphol.* 17 (3) (2008) 161–164, <http://dx.doi.org/10.1097/mcd.0b013e328302f0c4>.
- [59] E.J. Bhoj, D. Li, M. Harr, S. Edvardson, O. Elpeleg, E. Chisholm, J. Juusola, G. Douglas, M.J. Guillen Sacoto, K. Siquier-Pernet, A. Saadi, C. Bole-Feyost, P. Nitschke, A. Narravula, M. Walke, M.B. Horner, D.-L. Day-Salvatore, P. Jayakar, S.A.S. Vergano, M.A. Tarnopolsky, M. Hegde, L. Collea, P. Crino, H. Hakonarson, Mutations in TBCK, encoding TBC1-domain-containing kinase, lead to a recognizable syndrome of intellectual disability and hypotonia, *Am. J. Hum. Genet.* 98 (4) (2016) 782–788, <http://dx.doi.org/10.1016/j.ajhg.2016.03.016>.
- [60] J.X. Chong, V. Caputo, I.G. Phelps, L. Stella, L. Worgan, J.C. Dempsey, A. Nguyen, V. Leuzzi, R. Webster, A. Pizzuti, C.T. Marvin, G.E. Ishak, S. Ardern-Holmes, Z. Richmond, University of Washington Center for Mendelian Genomics, M.J. Bamshad, X.R. Ortiz-Gonzalez, M. Tartaglia, M. Chopra, D. Doherty, Recessive inactivating mutations in TBCK, Encoding a rab GTPase-activating protein, cause severe infantile syndromic encephalopathy, *Am. J. Hum. Genet.* 98 (4) (2016) 772–781, <http://dx.doi.org/10.1016/j.ajhg.2016.01.016>.
- [61] R.J. Guerreiro, R. Brown, D. Dian, C. de Goede, J. Bras, S.E. Mole, Mutation of TBCK causes a rare recessive developmental disorder, *Neurol. Genetics* 2 (3) (2016) e76, <http://dx.doi.org/10.1212/NXG.0000000000000076>.
- [62] E. Zapata-Aldana, D.D. Kim, S. Remtulla, C. Prasad, C.-T. Nguyen, C. Campbell, Further delineation of TBCK - Infantile hypotonia with psychomotor retardation and characteristic facies type 3, *Eur. J. Med. Genet.* 62 (4) (2019) 273–277, <http://dx.doi.org/10.1016/j.ejmg.2018.08.004>.
- [63] S. Köhler, M. Gargano, N. Matentzoglou, L.C. Carmody, D. Lewis-Smith, N.A. Vasilevsky, D. Danis, G. Balagura, G. Baynam, A.M. Brower, T.J. Callahan, C.G. Chute, J.L. Est, P.D. Galer, S. Ganesan, M. Griese, M. Haimel, J. Pazmandi, M. Hanauer, N.L. Harris, M.J. Hartnett, M. Hastreiter, F. Hauck, Y. He, T. Jeske, H. Kearney, G. Kindle, C. Klein, K. Knoflach, R. Krause, D. Lagorce, J.A. McMurry, J.A. Miller, M.C. Munoz-Torres, R.L. Peters, C.K. Rapp, A.M. Rath, S.A. Rind, A.Z. Rosenberg, M.M. Segal, M.G. Seidel, D. Smedley, T. Talmy, Y. Thomas, S.A. Wiafe, J. Xian, Z. Yüksel, I. Helbig, C.J. Mungall, M.A. Haendel, P.N. Robinson, The Human Phenotype Ontology in 2021, *Nucleic Acids Res.* 49 (D1) (2021) D1207–D1217, <http://dx.doi.org/10.1093/nar/gkaa1043>.
- [64] D. Viasus, C. Garcia-Vidal, J.M. Cruzado, J. Adamuz, R. Verdaguier, F. Manresa, J. Dorca, F. Gudiol, J. Carratala, Epidemiology, clinical features and outcomes of pneumonia in patients with chronic kidney disease, *Nephrol. Dial. Transplant.* 26 (9) (2011) 2899–2906, <http://dx.doi.org/10.1093/ndt/gfq798>.
- [65] C. Kearon, Natural history of venous thromboembolism, *Circulation* 107 (23_suppl_1) (2003) <http://dx.doi.org/10.1161/01.cir.0000078464.82671.78>.
- [66] G. Yang, C.D. Staercke, W.C. Hooper, The effects of obesity on venous thromboembolism: A review, *Open J. Prevent. Med.* 02 (04) (2012) 499–509, <http://dx.doi.org/10.4236/ojpm.2012.24069>.
- [67] S.K. Kakkos, M. Gohel, N. Baekgaard, R. Bauersachs, S. Bellmunt-Montoya, S.A. Black, A.J. ten Cate-Hoek, I. Elalamy, F.K. Enzmann, G. Geroulakos, A. Gottsäter, B.J. Hunt, A. Mansilha, A.N. Nicolaidis, P.M. Sandset, G. Stansby, E.G. Committee, G.J. de Borst, F.B. Gonçalves, N. Chakfé, R. Hinchliffe, P. Kolh, I. Koncar, J.S. Lindholt, R. Tulamo, C.P. Twine, F. Vermassen, A. Wanhainen, D. reviewers, M.G.D. Maeseneer, A.J. Comerota, P. Gloviczki, M.J. Kruij, M. Monreal, P. Prandoni, M.V. de Ceniça, Editor's choice – European society for vascular surgery (ESVS) 2021 clinical practice guidelines on the management of venous thrombosis, *Eur. J. Vasc. Endovascular Surg.* 61 (1) (2021) 9–82, <http://dx.doi.org/10.1016/j.ejvs.2020.09.023>.
- [68] A.M. Alazami, N. Patel, H.E. Shamseldin, S. Anazi, M.S. Al-Dosari, F. Alzahrani, H. Hijazi, M. Alshammari, M.A. Aldahmesh, M.A. Salih, E. Faqeih, A. Alhashem, F.A. Bashiri, M. Al-Owain, A.Y. Kentab, S. Sogaty, S. Al Tala, M.H. Temsah, M. Tulbah, R.F. Aljelaify, S.A. Alshahwan, M.Z. Seidahmed, A.A. Alhadid, H. Aldhalaan, F. AlQallaf, W. Kurdi, M. Alfadhel, Z. Babay, M. Alsogheer, N. Kaya, Z.N. Al-Hassnan, G.M. Abdel-Salam, N. Al-Sannaa, F. Al Mutairi, H.Y. El Khashab, S. Bohlega, X. Jia, H.C. Nguyen, R. Hammami, N. Adly, J.Y. Mohamed, F. Abdulwahab, N. Ibrahim, E.A. Naim, B. Al-Younes, B.F. Meyer, M. Hashem, R. Shaheen, Y. Xiong, M. Abouelhoda, A.A. Aldeeri, D.M. Monies, F.S. Alkuraya, Accelerating novel candidate gene discovery in neurogenetic disorders via whole-exome sequencing of prescreened multiplex consanguineous families, *Cell Rep.* 10 (2015) 148–161, <http://dx.doi.org/10.1016/j.celrep.2014.12.015>.
- [69] J. Havrilla, M. Zhao, C. Liu, C. Weng, I. Helbig, E. Bhoj, K. Wang, Clinical phenotypic spectrum of 4095 individuals with Down syndrome from text mining of electronic health records, *Genes* 12 (2021) 1159, <http://dx.doi.org/10.3390/genes12081159>.