

# Searching for explanations of black-box classifiers in the space of semantic queries

Jason Liartis<sup>\*</sup>, Edmund Dervakos, Orfeas Menis-Mastromichalakis, Alexandros Chortaras and Giorgos Stamou

*Artificial Intelligence and Learning Systems Laboratory, School of Electrical and Computer Engineering, National Technical University of Athens, Greece*

*E-mails: jliartis@ails.ece.ntua.gr, eddiedervakos@islab.ntua.gr, menorf@ails.ece.ntua.gr, achort@cs.ntua.gr, gstam@cs.ntua.gr*

**Editors:** Roberto Confalonieri, Free University of Bozen-Bolzano, Faculty of Computer Science, Italy; Oliver Kutz, Free University of Bozen-Bolzano, Faculty of Computer Science, Italy; Diego Calvanese, Department of Computing Science, Umeå University, Sweden and Free University of Bozen-Bolzano, Faculty of Computer Science, Italy; Jose M. Alonso, University of Santiago de Compostela, CiTIUS, Spain; Shang-Ming Zhou, University of Plymouth, Faculty of Health, UK

**Solicited reviews:** anonymous reviewer; Daniele Porello, University of Genova, DAFIST, Italy; anonymous reviewer

**Abstract.** Deep learning models have achieved impressive performance in various tasks, but they are usually opaque with regards to their inner complex operation, obfuscating the reasons for which they make decisions. This opacity raises ethical and legal concerns regarding the real-life use of such models, especially in critical domains such as in medicine, and has led to the emergence of the eXplainable Artificial Intelligence (XAI) field of research, which aims to make the operation of opaque AI systems more comprehensible to humans. The problem of explaining a black-box classifier is often approached by feeding it data and observing its behaviour. In this work, we feed the classifier with data that are part of a knowledge graph, and describe the behaviour with rules that are expressed in the terminology of the knowledge graph, that is understandable by humans. We first theoretically investigate the problem to provide guarantees for the extracted rules and then we investigate the relation of "explanation rules for a specific class" with "semantic queries collecting from the knowledge graph the instances classified by the black-box classifier to this specific class". Thus we approach the problem of extracting explanation rules as a semantic query reverse engineering problem. We develop algorithms for solving this inverse problem as a heuristic search in the space of semantic queries and we evaluate the proposed algorithms on four simulated use-cases and discuss the results.

**Keywords:** Explainable AI (XAI), opaque machine learning classifiers, knowledge graphs, description logics, semantic query answering, reverse query answering, post-hoc explainability, explanation rules

## 1. Introduction

The opacity of deep learning models raises ethical and legal [1] concerns regarding the real-life use of such models, especially in critical domains such as medicine and judicial, and has led to the emergence of the eXplainable Artificial Intelligence (XAI) field of research, which aims to make the operation of opaque AI systems more comprehensible to humans [2, 3]. While many traditional machine learning models, such as decision trees, are inter-

---

<sup>\*</sup>Corresponding author. E-mail: jliartis@ails.ece.ntua.gr.

1     pretable by design, they typically perform worse than deep learning approaches for various tasks. Thus, in order     1  
2     to not sacrifice performance for the sake of transparency, a lot of research is focused on *post hoc* explainability, in     2  
3     which the model to be explained is treated as a black-box. Approaches to *post hoc* explainability vary with regard to     3  
4     data domain (images, text, tabular), form of explanations (rule-based, counterfactual, feature importance etc.), scope     4  
5     (global, local) [4] and application domain [5]. In this work we focus on global explanations, which attempt to explain     5  
6     the general function of a black-box regardless of data, as opposed to local explanations which attempt to explain     6  
7     the prediction of a classifier on a particular data sample. Specifically, we attempt to extract *rules* which simulate the     7  
8     behaviour of a black-box by considering semantic descriptions of samples, in addition to external knowledge. For     8  
9     example such a rule might be “Images depicting animals and house-hold items are classified as domestic animals”,     9  
10    where the semantic description of an image might be “This image depicts an elephant next to a chair” and external     10  
11    knowledge might contain information such as “elephants are animals” and “chairs are household items”. We do this     11  
12    by utilizing terminological, human-understandable knowledge expressed in the form of ontologies and knowledge     12  
13    graphs, taking advantage of the reasoning capabilities of the underpinning description logics [6]. Such extracted     13  
14    rules might be very useful for an end user to understand the reasons behind why an opaque model is making its     14  
15    decisions, especially when combined with other forms of explanations, such as local contrastive explanations [7].     15

16    Our approach to global rule-based explanations assumes that we are given a set of data samples with semantic     16  
17    descriptions linked to an external knowledge. The explanations will be presented in the terminology provided by     17  
18    the external knowledge. We call such a set of samples an *explanation dataset*. For example, a semantic description     18  
19    for an image might refer to the objects it depicts and relationships between them, such as scene graphs from visual     19  
20    genome [8], or COCO [9]. The external knowledge utilized could be an existing knowledge base, such as WordNet     20  
21    [10], ConceptNet [11], DBpedia [12], and even domain specific knowledge such as SNOMED-CT [13] for the     21  
22    medical domain, or it could be a knowledge base curated specifically for that explanation dataset. Given such a     22  
23    set of semantically enriched data, we then compute global rule-based explanations characterizing the output of a     23  
24    black-box classifier. We approach the problem as a query reverse engineering (QRE) problem by utilizing the close     24  
25    relationship between definite rules and conjunctive queries.     25

26    Utilizing external knowledge to boost transparency of opaque AI is an active research area which has produced     26  
27    important results in recent years [14–20]. Specifically, knowledge graphs (KG) [21] as a scalable common un-     27  
28    derstandable structured representation of a domain based on the way humans mentally perceive the world, have     28  
29    emerged as a promising complement or extension to machine learning approaches for achieving explainability. A     29  
30    particular aspect which might be improved by utilizing knowledge graphs, especially for generalized global ex-     30  
31    planations, is the form of the produced explanations. When the feature space of the classifier is sub-symbolic raw     31  
32    data, then providing explanations in terms of features might lead to unintuitive, or even misleading results [22–24].     32  
33    On the other hand, if there is underlying knowledge of the data, then explanations can be provided by using the     33  
34    terminology of the knowledge. For example, if a black-box classified every image depicting wild animals in the     34  
35    class *zoo*, a rule of the form “If an image depicts a giraffe or a zebra or . . . then it is classified as a *zoo*”, might be     35  
36    more intuitive than for example sets of pixel importances. Furthermore, by exploiting external knowledge, the form     36  
37    can be further condensed and lead to simpler explanations which are semantically identical, such as “If an image     37  
38    depicts a *wild animal* then it is classified as a *zoo*”. This approach to explainability mitigates some of the typical     38  
39    pitfalls of *post hoc* XAI, and transfers potential sources of erroneous or misleading explanations to the level of the     39  
40    data and the knowledge (instead of the explainability algorithm), which are easier to control in most settings. For     40  
41    instance, a typical pitfall of *post hoc* XAI mentioned in [24], is that explanations often do not make sense, or are     41  
42    incomplete (leaving out information to be understandable). In our approach, the explanations make as much sense     42  
43    as the terminology used in the knowledge graph, and are as informative as the explanation dataset.     43

44    There are many related rule-based explanation methods (both global and local) in recent literature. Many ap-     44  
45    proaches rely on statistics in order to generate lists of IF-THEN rules which mimic the behaviour of a classifier     45  
46    [25–29], while others extract rules in the form of decision trees [30, 31] or First-Order Logic expressions [32]. Rules     46  
47    have been argued to be the preferred method of presenting explanations to humans, as they are close to the way     47  
48    humans reason [4, 33, 34]. There are also related works that utilize ontologies for the purpose of explainability, and     48  
49    we build on these ideas in our work. In [19] the authors utilize ontologies, expressed in the Description Logic  $\mathcal{EL}$ , to     49  
50    enhance explanations of black-box classifiers produced in the form of decision trees. The authors conducted a user     50  
51    evaluation from which they concluded that users prefer decision trees extracted with the presence of an ontology     51

1 than those without. In [20] the authors present a framework quite similar to our own. They also approach the prob- 1  
 2 lem of extracting explanations as a query reverse engineering problem, but there is no emphasis on the data utilized 2  
 3 from producing the explanations compared to our notion of an explanation dataset. Their work is also preliminary 3  
 4 as they do not present any algorithms for extracting the desired queries. 4

5 Query Reverse Engineering (QRE) is the problem of reverse-engineering a query from its output [35]. This is 5  
 6 typically presented as a way for an end-user unfamiliar with a query language such as SQL to formulate a query 6  
 7 by simply providing examples of what they expect to see in the query results. This is also usually approached as an 7  
 8 interactive process; the user is continually providing more examples and marking answers of the query as desired 8  
 9 or undesired. QRE has been extensively studied for relational databases and the language of relational algebra or 9  
 10 SQL [36]. In this work we assume data is stored in a Knowledge Base and queries are formulated in the language 10  
 11 of Conjunctive Queries. We also assume that we know the exact desired output of the query, without the need 11  
 12 for a user to continually provide examples. Although conjunctive queries have the same expressivity as select- 12  
 13 project-join relational queries using only equality predicates, adapting existing works to fit our framework is not 13  
 14 easy or even desirable. Some common obstacles include i) the assumption of a dynamic labeling of individuals as 14  
 15 positive/negative by an end-user [37–39], ii) the (often implicit) exclusion of the self-join operation, which would 15  
 16 limit the expressivity of queries in ways undesirable to us [35, 37–39], and iii) foreign/primary key constraints [40]. 16

17 In the setting of conjunctive queries over Knowledge Bases, referred to as Semantic Query Answering or On- 17  
 18 tology Mediated Querying [41–45], QRE has been under-explored. Recent work focuses on establishing theoretic- 18  
 19 al results, such as the existence of solutions and the complexity of computing them [46–50]. The authors in [51] 19  
 20 present a few algorithms alongside complexity results but the algorithms are limited to tree-shaped queries. One 20  
 21 recent article [52], along with some theoretical results, presents two algorithms computing minimally complete and 21  
 22 maximally sound solutions to the closely related Query Separability problem. We consider these algorithms unsuit- 22  
 23 able in the context of explainability, since the solutions computed essentially amount to one large disjunction of 23  
 24 positive examples. In this work we focus on practical QRE over KBs and we present several algorithms which either 24  
 25 have polynomial complexity or have worst-case exponential complexity but can be efficient in practice. We employ 25  
 26 heuristics to limit the number of computations, we take care to limit the size of queries, since their intended use is 26  
 27 that of explanations presented to users, and we provide (approximate) solutions even when a exact solution does not 27  
 28 exist. 28

29 Our contributions: 29

- 30 – Following our previous work in the area [46], we here present a framework for extracting global rule-based 30
- 31 explanations of black-box classifiers, using exemplar items, external terminology and underlying knowledge 31
- 32 stored in a knowledge graph and defining the problem of explanation rule extraction as a semantic query reverse 32
- 33 engineering problem over the knowledge graph (see section 3). 33
- 34 – We propose practical algorithms which approximate the semantic query reverse engineering problem by using 34
- 35 heuristics, which we then use to generate explanations in the context of the proposed framework (see section 35
- 36 4). 36
- 37 – We implement the proposed algorithms and show results from experiments explaining image classifiers on 37
- 38 CLEVR-Hans3, Places365 and MNIST. We also compare our work with existing post-hoc, rule-based expla- 38
- 39 nation methods on baseline tabular data employing the Mushroom dataset (see section 5). 39
- 40 40

## 41 2. Background 41

### 42 2.1. Description Logics 42

43 Let  $\mathcal{V} = \langle \text{CN}, \text{RN}, \text{IN} \rangle$  be a *vocabulary*, where CN, RN, IN are mutually disjoint finite sets of *concept names* 43  
 44 (atomic concepts), *role names* (atomic roles) and *individual names*, respectively. Atomic concepts and roles can 44  
 45 be combined into complex concepts or roles using appropriate *constructors*. The set of available constructors for 45  
 46 constructing complex concepts and roles, along with their syntax and semantics defines a Description Logics (DL) 46  
 47 *dialect*. A list of common constructors is provided in Table 1, which are the constructors used by the *SHROIQ* 47  
 48 48  
 49 49  
 50 50  
 51 51

Table 1  
DL syntax and semantics.

Name	Syntax	Semantics
Top	$\top$	$\Delta^{\mathcal{I}}$
Bottom	$\perp$	$\emptyset$
Complement	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
Intersection	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
Union	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
Existential quantification	$\exists r.C$	$\{a \in \Delta^{\mathcal{I}} \mid \exists b. (a, b) \in r^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\}$
Value restriction	$\forall r.C$	$\{a \in \Delta^{\mathcal{I}} \mid \forall b. (a, b) \in r^{\mathcal{I}} \rightarrow b \in C^{\mathcal{I}}\}$
At-least number restriction	$\geq nr.C$	$\{a \in \Delta^{\mathcal{I}} \mid  \{b \in \Delta^{\mathcal{I}} \mid (a, b) \in r^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\}  \geq n\}$
At-most number restriction	$\leq nr.C$	$\{a \in \Delta^{\mathcal{I}} \mid  \{b \in \Delta^{\mathcal{I}} \mid (a, b) \in r^{\mathcal{I}} \wedge b \in C^{\mathcal{I}}\}  \leq n\}$
Nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
Inverse role	$r^{-}$	$\{(b, a) \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \mid (a, b) \in r^{\mathcal{I}}\}$
Role composition	$r \circ s$	$r^{\mathcal{I}} \circ s^{\mathcal{I}}$

Table 2  
Terminological and assertional axioms.

Name	Syntax	Semantics
Concept inclusion	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
Concept equivalence	$C \equiv D$	$C^{\mathcal{I}} = D^{\mathcal{I}}$
Role inclusion	$r \sqsubseteq s$	$r^{\mathcal{I}} \subseteq s^{\mathcal{I}}$
Role equivalence	$r \equiv s$	$r^{\mathcal{I}} = s^{\mathcal{I}}$
Concept assertion	$C(a)$	$a^{\mathcal{I}} \in C^{\mathcal{I}}$
Role assertion	$r(a, b)$	$(a, b)^{\mathcal{I}} \in r^{\mathcal{I}}$
Individual equality	$a = b$	$a^{\mathcal{I}} = b^{\mathcal{I}}$
Individual inequality	$a \neq b$	$a^{\mathcal{I}} \neq b^{\mathcal{I}}$

dialect, which underpins the Web Ontology Language OWL 2. Note that the table defines also the two special concept names  $\top$  and  $\perp$  that may belong to CN, representing the universal and empty concept, respectively.

Concepts and roles (either atomic or complex) are used to form terminological and assertional axioms, as shown in Table 2. Then, a *DL knowledge base*, simply a *knowledge base*, (KB) over a DL dialect  $\mathcal{L}$  and a vocabulary  $\mathcal{V}$  is a pair  $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ , where  $\mathcal{T}$  is a set of terminological axioms (also called TBox), and  $\mathcal{A}$  a set of assertional axioms (also called ABox) using the constructors of  $\mathcal{L}$  and elements of  $\mathcal{V}$ . We will call an ABox containing only assertions of the form  $C(a)$  and  $r(a, b)$ , where  $C \in \text{CN}$ ,  $r \in \text{RN}$ ,  $a, b \in \text{IN}$ , i.e. not involving complex expressions and (in)equality assertions, an *atomic ABox*.

The semantics of DL KBs are defined in the standard model-theoretic way using interpretations. In particular, given a non-empty domain  $\Delta$ , an interpretation  $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$  assigns a set  $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$  to each atomic concept  $C \in \text{CN}$ , a set of pairs  $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$  to each atomic role  $r \in \text{RN}$ , and an element  $a^{\mathcal{I}} \in \Delta$  to each individual  $a \in \text{IN}$ . Table 1 shows also the semantics of the listed constructors. An interpretation  $\mathcal{I}$  is a *model* of a KB  $\mathcal{K}$  iff it satisfies all axioms in  $\mathcal{T}$  and all assertions in  $\mathcal{A}$ , i.e. iff it satisfies the expression in the last column of Table 2 for the respective axioms.

The underlying DL dialect  $\mathcal{L}$  determines the expressivity of a knowledge base. Most classical DL dialects can be seen as fragments of first-order logic (FOL) by viewing atomic concepts and roles as unary and binary predicates respectively. In this paper we refer only to DL dialects that are fragments of first-order logic, and hence can be translated to first-order logic theories. We will denote the translation of such a  $\mathcal{K}$  to the respective first order logic theory by  $\text{FOL}(\mathcal{K})$ . E.g. the FOL translation of the DL concept  $\exists r.C$  is an expression of the form  $\exists x.(r(y, x) \wedge \phi_C(x))$ , where  $\phi_C(x)$  is the FOL translation of concept  $C$ . For more details see [6].

## 2.2. Conjunctive Queries

Given a vocabulary  $\mathcal{V}$ , a *conjunctive query* (simply, a *query*)  $q$  is an expression of the form  $\langle \{x_1, \dots, x_k\} \mid \exists y_1 \dots \exists y_l. (c_1 \wedge \dots \wedge c_n) \rangle$ , where  $k, l, n \geq 0$ ,  $x_i, y_i$  are variable names, the  $c_i$ s are distinct, each  $c_i$  is an atom  $C(u)$  or  $r(u, v)$ , where  $C \in \text{CN}$ ,  $r \in \text{RN}$ ,  $u, v \in \text{IN} \cup \{x_1, \dots, x_k\} \cup \{y_1, \dots, y_l\}$  and all  $x_i, y_i$  appear in at least one atom. The vector  $\langle x_1, \dots, x_k \rangle$  is the *head* of  $q$ , its elements are the *answer variables*, and  $\{c_1, \dots, c_n\}$  is the *body* of  $q$  ( $\text{body}(q)$ ). Generally, in the relevant literature, conjunctive queries are defined such that the body must be nonempty ( $n \geq 1$ ) but we allow empty bodies in the sense that we do not impose any constraints that the answers of the query must satisfy, i.e. all individuals are answers. As we will see, this is a natural extension that conforms to all other definitions we use from the existing literature, without needing any special adjustments. The set  $\text{var}(q)$  is the set of all variables appearing in  $q$ .

On conjunctive queries one can apply substitutions, which map variables to other variables or individual names. More formally, a *substitution*  $\theta$  is a finite set of mappings of the form  $\{z_1 \mapsto u_1, \dots, z_n \mapsto u_n\}$  where  $z_i$  is a variable name,  $u_i$  is either an individual name or a variable name distinct from all  $z_i$ s, and the variables  $z_i$  are distinct. A substitution may contain at most one mapping for a variable. The conjunctive query  $q\theta$  resulting from applying  $\theta$  on  $q$  is obtained from  $q$  by replacing all occurrences of  $z$  in  $q$  by  $u$  for each  $z \mapsto u \in \theta$ . Any variables not mapped by  $\theta$  remain unchanged.

Since in this paper we are interested in classifying single items, we focus on queries having a *single* answer variable. For convenience, we will also assume that all arguments of all  $c_i$  in a query are variables, and we call such queries *single answer variable queries*, or *SAV queries* for short. For simplicity, in the rest of the paper, by saying query we will mean SAV query, and we will write a (SAV) query  $q$  as  $\{c_1, \dots, c_n\}$ , considering always  $x$  as the answer variable, essentially identifying the query by its body. This will allow us to treat queries as sets, and write e.g.  $q_1 \cup q_2$ ; this query represents the query that has the same head as  $q_1$  and  $q_2$  (i.e.  $\langle x \rangle$ ), and body the union of the bodies of  $q_1$  and  $q_2$ . Similarly, we will write e.g.  $q_1 \subseteq q_2$ , meaning  $\text{body}(q_1) \subseteq \text{body}(q_2)$ , and  $c \in q$  meaning  $c \in \text{body}(q)$ . Following the above assumptions, all definitions that follow will be stated only for SAV queries, although more general formulations for general conjunctive queries might exist.

Given a KB  $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ , a SAV query  $q$  and an interpretation  $\mathcal{I}$  of  $\mathcal{K}$ , a *match* for  $q$  is a mapping  $\pi : \text{var}(q) \rightarrow \Delta^{\mathcal{I}}$  such that  $\pi(u) \in C^{\mathcal{I}}$  for all  $C(u) \in q$ , and  $(\pi(u), \pi(v)) \in r^{\mathcal{I}}$  for all  $r(u, v) \in q$ . Then,  $a \in \text{IN}$  is an *answer* for  $q$  in  $\mathcal{I}$  if there is a match  $\pi$  for  $q$  such that  $\pi(x) = a^{\mathcal{I}}$ . For queries with empty bodies it obviously holds that  $\text{ans}(q, \mathcal{I}) = \text{IN}$ . Finally,  $a \in \text{IN}$  is a *certain answer* for  $q$  over  $\mathcal{K}$  if in every model  $\mathcal{I}$  of  $\mathcal{K}$  there is a match  $\pi$  for  $q$  such that  $\pi(x) = a^{\mathcal{I}}$ . The set of certain answers to  $q$  is denoted by  $\text{cert}(q, \mathcal{K})$ . Naturally, when the query body is empty it holds that  $\text{cert}(q, \mathcal{K}) = \text{IN}$ . Because obtaining the set of certain answers of a query over some KB typically requires reasoning with the axioms of the underlying TBox, such queries are characterized in the literature as *semantic queries*.

Let  $\mathcal{K}$  be a knowledge base over a vocabulary  $\mathcal{V}$ , and  $\mathcal{Q}$  the (possibly infinite) set of all SAV queries over  $\mathcal{V}$ . We can partially order  $\mathcal{Q}$  using *query subsumption*: A SAV query  $q_2$  *subsumes* a SAV query  $q_1$  (we write  $q_1 \leq_S q_2$ ) iff there is a substitution  $\theta$  s.t.  $q_2\theta \subseteq q_1$  and  $\theta$  leaves variable  $x$  unchanged. If  $q_1, q_2$  are mutually subsumed, they are *syntactically equivalent* ( $q_1 \equiv_S q_2$ ).  $q \leq_S q'$  implies  $\text{cert}(q, \mathcal{K}) \subseteq \text{cert}(q', \mathcal{K})$ , since a match  $\pi$  for the variables of  $q$  can be composed with  $\theta$  to produce a match  $\pi\theta$  for the variables of  $q'$ . Let  $q, q'$  be queries s.t.  $q' \subseteq q$ . If  $q'$  is a minimal subset of  $q$  s.t.  $q' \equiv_S q$ , then  $q'$  is a *condensation* of  $q$ . If that minimal  $q'$  is the same as  $q$ , then  $q$  is *condensed* [53]. Intuitively, syntactically equivalent queries always have the same certain answers, and a condensation of some syntactically equivalent queries is the most compact query (not containing redundant atoms) that is syntactically equivalent to the rest.

Given the queries  $q_1, q_2, \dots, q_n$ , a query least common subsumer  $\text{QLCS}(q_1, q_2, \dots, q_n)$  of them is defined as a query  $q$  for which  $q_1, q_2, \dots, q_n \leq_S q$ , and for all  $q'$  s.t.  $q_1, q_2, \dots, q_n \leq_S q'$  we have  $q \leq_S q'$ . The query least common subsumer can be seen as the most specific generalization of  $q_1, q_2, \dots, q_n$ , and it is unique up to syntactical equivalence. We should note that this notion of query least common subsumer is different from the usual notion of least common subsumer of concepts which has been extensively studied for various description logic expressivities [54–57].

### 2.3. Graphs

A SAV query  $q$  can be viewed as the directed labeled graph  $\langle V, E, \ell_V, \ell_E \rangle$  (a *query graph*), where  $V = \text{var}(q)$  is the set of nodes,  $E = \{\langle u, v \rangle \mid r(u, v) \in q\} \subseteq V \times V$  is the set of edges,  $\ell_V : V \rightarrow 2^{\text{CN}}$  with  $\ell_V(u) = \{C \mid C(u) \in q\}$  is the node labeling function, and  $\ell_E : E \rightarrow 2^{\text{RN}}$  with  $\ell_E(\langle u, v \rangle) = \{r \mid r(u, v) \in q\}$  is the edge labeling function. The answer variable is not explicitly identified since it is assumed to be always  $x$ . A SAV query is *connected*, if the corresponding query graph is connected.

In the following, it will be useful to also treat atomic ABoxes as graphs. Similarly to the case of SAV queries, an atomic ABox  $\mathcal{A}$  can be represented as the graph  $\langle V, E, \ell_V, \ell_E \rangle$  (an *ABox graph*), where  $V \subseteq \text{IN}$  is the set of nodes, in particular the subset of all individual names occurring in  $\mathcal{A}$ ,  $E = \{\langle a, b \rangle \mid r(a, b) \in \mathcal{A}\} \subseteq V \times V$  is the set of labeled edges,  $\ell_V : V \rightarrow 2^{\text{CN}}$  with  $\ell_V(a) = \{C \mid C(a) \in \mathcal{A}\}$  is the node labeling function, and  $\ell_E : E \rightarrow 2^{\text{RN}}$  with  $\ell_E(\langle a, b \rangle) = \{r \mid r(a, b) \in \mathcal{A}\}$  is the edge labeling function.

Given two graphs  $G_1 = \langle V_1, E_1, \ell_{V_1}, \ell_{E_1} \rangle$ ,  $G_2 = \langle V_2, E_2, \ell_{V_2}, \ell_{E_2} \rangle$ , a homomorphism  $h : G_1 \rightarrow G_2$  is defined as a function from  $V_1$  to  $V_2$  that preserves edges and labels. More specifically it is such that: i) if  $\langle a, b \rangle \in E_1$  then  $\langle h(a), h(b) \rangle \in E_2$ , ii)  $\ell_{V_1}(a) \subseteq \ell_{V_2}(h(a))$ , and iii)  $\ell_{E_1}(\langle a, b \rangle) \subseteq \ell_{E_2}(\langle h(a), h(b) \rangle)$ . If there exists a homomorphism from  $G_1$  to  $G_2$ , we will write for simplicity  $G_1 \rightarrow G_2$ . When  $G_1$  and  $G_2$  are query graphs we will make the additional assumption that  $h$  preserves the answer variable, i.e.  $h(x) = x$ . If  $h$  is a bijection whose inverse is also a homomorphism, then  $h$  is an isomorphism. It is easy to see that the query graph of  $q_1$  is homomorphic to the query graph of  $q_2$  iff  $q_2 \leq_S q_1$ .

### 2.4. Rules

A (definite) *rule* is a FOL expression of the form  $\forall x_1 \dots \forall x_n (c_1, \dots, c_m \Rightarrow c_0)$ , where  $c_i$  are atoms and  $x_1, \dots, x_n$  are all the variables appearing in the several  $c_i$ . The atoms  $c_1, \dots, c_m$  are the *body* of the rule, and  $c_0$  its *head*. A rule over a vocabulary  $\mathcal{V} = \langle \text{CN}, \text{RN}, \text{IN} \rangle$  is a rule where each  $c_i$  is either  $C(u)$ , where  $C \in \text{CN}$ , or  $r(u, v)$ , where  $r \in \text{RN}$ . Assuming that  $c_0$  is of the form  $D(x)$ , where  $D \in \text{CN}$ , and that  $x$  appears in the body of such a rule, we will say that the rule is *connected*, if its body, seen as a SAV query is connected. A rule is usually written as  $c_1, \dots, c_m \rightarrow c_0$ .

### 2.5. Classifiers

A classifier is viewed as a function  $F : \mathcal{D} \rightarrow \mathcal{C}$ , where  $\mathcal{D}$  is the classifier's domain (ex. images, audio, text), and  $\mathcal{C}$  is a set of class names (ex. "Dog", "Cat").

## 3. Framework

### 3.1. A motivating example

Integration of artificial intelligence methods and tools with biomedical and health informatics is a promising area, in which technologies for explaining machine learning classifiers will play a major role. In the context of the COVID-19 pandemic for example, black-box machine learning audio classifiers have been developed, which, given audio of a patient's cough, predict whether the person should seek medical advice or not [58]. In order to develop trust and use these classifiers in practice, it is important to explain their decisions, i.e. to provide convincing answers to the question "Why does the machine learning classifier suggest to *seek medical advice*?". There are post hoc explanation methods (both global and local) that try to answer this question in terms of the input of the black-box classifier, which in this case is audio signals. Although this information could be useful for AI engineers and developers, it is not understandable to most medical experts and end users, since audio signals themselves are obscure sub-symbolic data. Thus, it is difficult to convincingly meet the explainability requirements and develop the necessary trust to utilize the black-box classifier in practice, unless explanations are expressed in the terminology used by the medical experts (using terms like "sore throat", "dry cough" etc).

1 In the above context, suppose we have a dataset of audio signals of coughs which have been characterized by 1  
2 medical professionals by using standardized clinical terms, such as “Loose Cough”, “Dry Cough”, “Dyspnoea”, in 2  
3 addition to a knowledge base in which these terms and relationships between them are defined, such as SNOMED- 3  
4 CT [13]. For example, consider such a dataset with coughs from five patients p1, p2, p3, p4, p5 (obviously in practice 4  
5 we may need a much more extended set of patients) with characterizations from the medical experts: “p1 has a sore 5  
6 throat”, “p2 has dyspnoea”, “p3 has a sore throat and dyspnoea”, “p4 has a sore throat and a dry cough”, “p5 has a 6  
7 sore throat and a loose cough”. We also have available relationships between these terms as defined in SNOMED-CT 7  
8 like “Loose Cough is Cough”, “Dry Cough is Cough”, “Cough is Lung Finding”, and “Dyspnoea is Lung Finding”. 8

9 Now assume that a black-box classifier predicts that p3, p4, and p5 should seek medical advice, while p1 and 9  
10 p2 should not. Then we can say that: on this dataset, if a patient has a sore throat and a lung finding then it is 10  
11 classified positively by the specific classifier, i.e. the classifier suggests "seek medical advice". Depending on the 11  
12 characteristics of the dataset itself and its ability to cover interesting examples, such an extracted rule could help the 12  
13 medical professional understand why the black-box is making decisions in order to build necessary trust, but also it 13  
14 could help an AI engineer improve the model’s performance by indicating potential biases. Importantly, even though 14  
15 we cannot know that the black-box classifier is actually using the semantics that are used to describe the data, this 15  
16 approach gives a user control over what information they want to examine the classifier for. For example, a system 16  
17 developer and a medical professional would use different semantic annotations because explanations are useful to 17  
18 them at different levels of abstraction. In a more concrete example, in our experiments on the Visual Genome dataset 18  
19 (Sec. 5.3) we employ the WordNet general-purpose ontology [10], but another end user could easily swap it for 19  
20 another general-purpose ontology such as ConceptNet [59]. Importantly, this plug-and-play approach allows for the 20  
21 improvement of the terminology used in the explanations as the external ontologies improve over time. Another 21  
22 indication of the usefulness of ontologies comes from [19] in which the authors examine the understandability of 22  
23 explanations extracted with the aid of an ontology. Conducting a user study, the authors conclude that the users find 23  
24 explanations that utilize domain knowledge provided by the ontology more understandable. 24

25 The utilization of such datasets and knowledge is the main strength of the proposed framework, but it can also be 25  
26 its main weakness. Specifically, erroneous labeling during the construction of the dataset could lead to misleading 26  
27 explanations, as could biases in the data. For example, if the length of the recordings of the coughs were also 27  
28 available to us, and every instance of “Sore Throat” also had a short length recording, then we would not be able to 28  
29 distinguish between these two characteristics, and the resulting explanations would not be useful. 29

### 30 3.2. Explaining opaque machine learning classifiers 31

32 Explanation of opaque machine learning classifiers is based on a dataset that we call *Explanation Dataset* (see 33  
34 Fig. 1), containing *exemplar patterns*, that are actually characteristic examples from the set of elements that the 34  
35 unknown classifier takes as input. Machine learning classifiers usually take as input element features (like the cough 35  
36 audio signal mentioned in the motivating example). The explanation dataset additionally contains a semantic de- 36  
37 scription of the exemplars in terms that are understandable by humans (like "dry cough" mentioned in the motivating 37  
38 example). Taking the output of the unknown classifier (the classification class) for all the exemplars, we construct 38  
39 the *Exemplar Data Classification* information (see Fig. 1) thus we know the exact set of exemplars that are classi- 39  
40 fied by the unknown classifier to a specific class (like the "seek medical advice" class mentioned in the motivating 40  
41 example). Using the knowledge represented in the *Exemplar Data Semantic Description* (see Fig. 1), we define the 41  
42 following *reverse semantic query answering problem*: "given a set of exemplars and their semantic description find 42  
43 *intuitive and understandable* semantic queries that have as certain answers *exactly* this set out of all the exemplars 43  
44 of the explanation dataset". The specific problem is interesting, with certain difficulties and computationally very 44  
45 demanding [43, 45, 46, 51, 60]. Here, by extending a method presented in [61] we present an *Explainer* (see Fig. 1) 45  
46 that tries to solve this problem, following a *Semantic Query Heuristic Search* method, that searches in the *Semantic 46  
47 Query Space* (the set containing all queries that have non-empty certain answer set) for queries that are solutions to 47  
48 the above problem. 48

49 Introductory definitions and interesting theoretical results concerning the above approach are presented in [61]. 49  
50 Here, we reproduce some of them and introduce others, in order to develop the necessary framework for presenting 50  
51 the proposed method. 51

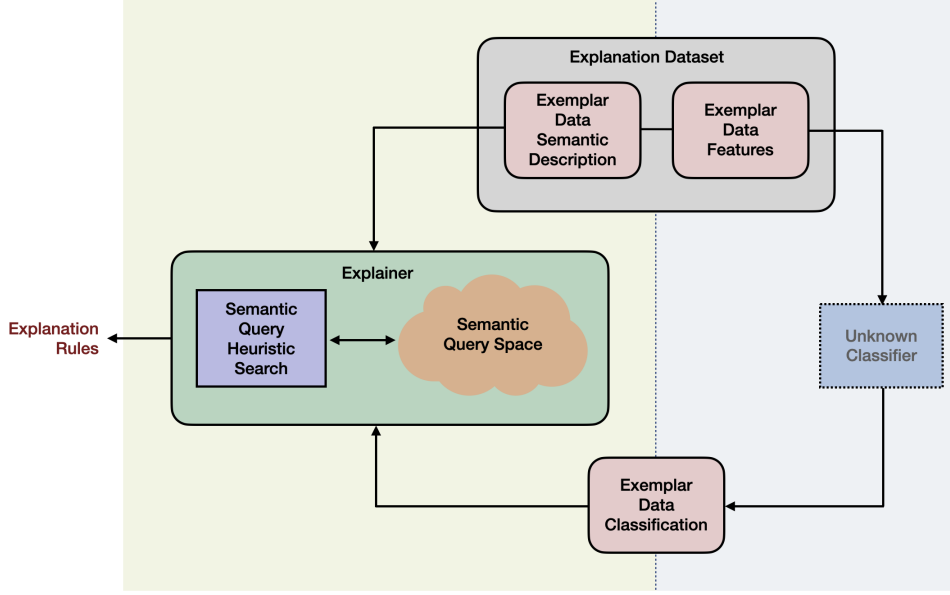


Fig. 1. A framework for explaining opaque machine learning classifiers

A defining aspect is that the rule explanations are provided in terms of a specific vocabulary. To do this in practice, we require a set of items (exemplar data) which can: a) be fed to a black-box classifier and b) have a semantic description using the desired terminology. As mentioned before, here we consider that: a) the exemplar data has for its items all the information that the unknown classifier needs in order to classify them (the necessary *features*), and b) the semantic data descriptions are expressed as Description Logics knowledge bases (see Fig. 1).

**Definition 1** ([46]). Let  $\mathcal{D}$  be a domain of item feature data,  $\mathcal{C}$  a set of classes, and  $\mathcal{V} = \langle \text{IN}, \text{CN}, \text{RN} \rangle$  a vocabulary such that  $\mathcal{C} \cup \{\text{Exemplar}\} \subseteq \text{CN}$ . Let also  $\text{EN} \subseteq \text{IN}$  be a set of exemplars. An explanation dataset  $\mathcal{E}$  in terms of  $\mathcal{D}$ ,  $\mathcal{C}$  and  $\mathcal{V}$  is a tuple  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$ , where  $\mathcal{M} : \text{EN} \rightarrow \mathcal{D}$  is a mapping from the exemplars to the item feature data, and  $\mathcal{S} = \langle \mathcal{T}, \mathcal{A} \rangle$  is a DL knowledge base over  $\mathcal{V}$  such that  $\text{Exemplar}(a) \in \mathcal{A}$  iff  $a \in \text{EN}$ , the elements of  $\mathcal{C}$  do not appear in  $\mathcal{S}$ , and  $\text{Exemplar}$  and the elements of  $\text{EN}$  do not appear in  $\mathcal{T}$ .

Intuitively, an explanation dataset contains a set of exemplar data (i.e. characteristic items in  $\mathcal{D}$  which can be fed to the unknown classifier) semantically described in terms of a specific vocabulary  $\mathcal{V}$ ; the semantic descriptions are in the knowledge base  $\mathcal{S}$ . Each exemplar data item is represented in  $\mathcal{S}$  by an individual name; these individual names make up the set of exemplars  $\text{EN}$ , and each one of them is mapped to the corresponding exemplar data item by  $\mathcal{M}$ . Because the knowledge encoded in  $\mathcal{S}$  may involve also other individuals, the Exemplar concept exists to identify exactly those individuals that correspond to exemplar data within the knowledge base. Given a classifier  $F : \mathcal{D} \rightarrow \mathcal{C}$  and a set of individuals  $\mathcal{I} \subseteq \text{EN}$ , the positive set (pos-set) of  $F$  on  $\mathcal{I}$  for class  $C \in \mathcal{C}$  is  $\text{pos}(F, \mathcal{I}, C) = \{a \in \mathcal{I} : F(\mathcal{M}(a)) = C\}$ . Based on the classifier's prediction on the exemplar data for a class (i.e. the pos-set) we can produce explanation rules by grouping them according to their semantic description in the explanation dataset.

**Definition 2** ([46]). Let  $F : \mathcal{D} \rightarrow \mathcal{C}$  be a classifier,  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$  an explanation dataset in terms of a domain  $\mathcal{D}$  and a set of classes  $\mathcal{C}$ , and an appropriate vocabulary  $\mathcal{V} = \langle \text{CN}, \text{RN}, \text{IN} \rangle$  (with  $\mathcal{C} \subseteq \text{CN}$ ). Given a concept  $C \in \mathcal{C}$ , a connected rule of the form

$$\text{Exemplar}(x), c_1, c_2, \dots, c_n \rightarrow C(x)$$

where  $c_i$  is an atom  $D(u)$  or  $r(u, v)$ , where  $D \in \text{CN}$ ,  $r \in \text{RN}$ , and  $u, v$  are variables, is an explanation rule of  $F$  for class  $C$  over  $\mathcal{E}$ . We denote the above rule by  $\rho(F, \mathcal{E}, C)$ , or simply by  $\rho$  whenever the context is clear. We may also



omit  $\text{Exemplar}(x)$  from the body, since it is a conjunct of any explanation rule. The explanation rule  $\rho$  is correct if and only if

$$\text{FOL}(\mathcal{S} \cup \{\text{Exemplar} \sqsubseteq \bigsqcup_{a \in \text{EN}} \{a\}\} \cup \{C(a) \mid a \in \text{pos}(F, \text{EN}, C)\}) \models \rho$$

where  $\text{FOL}(\mathcal{K})$  is the first-order logic translation of the DL knowledge base  $\mathcal{K}$ .

Note that in the above, the argument of  $\text{FOL}(\cdot)$  is a KB, i.e. a set of DL axioms, in particular all axioms originally in  $\mathcal{S}$  plus the technical axiom  $\text{Exemplar} \sqsubseteq \bigsqcup_{a \in \text{EN}} \{a\}$ , to force the interpretation of  $\text{Exemplar}$  to be contained in the interpretation of the elements of  $\text{EN}$ , and the assertions  $C(a)$ . (Note that because  $\text{EN}$  is a set of individuals, the expression  $\bigsqcup_{a \in \text{EN}} \{a\}$  makes use of the union and nominal DL constructors (one-of constructor) to form a concept by enumerating its members.)

Explanation rules describe sufficient conditions (the body of the rule) for an item to be classified in the class indicated at the head of the rule by the classifier under investigation. The correctness of a rule indicates that the rule covers every  $a \in \text{EN}$ , meaning that for each exemplar of  $\mathcal{S}$ , either the body of the rule is not true, or both the body and the head of the rule are true. Of course we cannot know if the classifier is using the particular semantics, however, depending on the explanation dataset, explanation rules can provide useful information that can indicate potential biases of the classifier, that are expressed in the desired terminology.

**Example 1.** Suppose we have the problem described in the motivating example of Section 3.1 with black-box classifiers predicting whether a person should seek medical advice based on audios of their cough, and that we are creating an explanation dataset in order to explain the respective classifiers. The vocabulary used should be designed so that it will produce meaningful explanations to the end-user, which in our case would probably be a doctor or another professional of the medical domain. In this case, it should contain concepts for the different medical terms like the findings (cough, sore throat, etc.), and according to the definition of the explanation dataset, the class categories (seek medical advice, or not) and the concept  $\text{Exemplar}$  as concept names (CN). Additionally, a role linking patients to the respective findings should exist in the role names (RN), and the patients as well as the findings themselves would be the individual names (IN). Following this, we create the vocabulary ( $\mathcal{V}$ ) as shown below:

$$\text{IN} = \{p1, p2, p3, p4, p5, s1, s2, s3, s4, s5, s6, s7, s8\}$$

$$\text{CN} = \{\text{DryCough}, \text{LooseCough}, \text{Cough}, \text{SoreThroat}, \text{LungFinding}, \text{Finding}, \text{Dyspnoea}, \text{MedicalAdvice}, \text{NoMedicalAdvice}, \text{Exemplar}\}$$

$$\text{RN} = \{\text{hasFinding}\}$$

Having the domain  $\mathcal{D}$  (audio signals), the set of classes  $\mathcal{C}$  ( $\text{MedicalAdvice}$ ,  $\text{NoMedicalAdvice}$ ), and the vocabulary  $\mathcal{V}$  we can now define the explanation dataset  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$ . The set of exemplars ( $\text{EN}$ ) in our case contains the patient individuals of  $\text{IN}$ , so  $\text{EN} = \{p1, p2, p3, p4, p5\}$ . The mapping  $\mathcal{M}$  of the explanation dataset links these exemplars to the respective audio of each patient. The only thing that is missing from our explanation dataset is the knowledge base  $\mathcal{S}$  consisting of an  $\text{ABox}$   $\mathcal{A}$  and a  $\text{TBox}$   $\mathcal{T}$ . The  $\text{ABox}$  contains information regarding the patient audio characterizations from the medical professionals (“p1 has a sore throat”, “p2 has dyspnoea”, “p3 has a sore throat and dyspnoea”, “p4 has a sore throat and a dry cough”, “p5 has a sore throat and a loose cough”) as well as the assertions regarding the exemplar status of individuals, while the  $\text{TBox}$  contains relationships between the medical terms as defined in  $\text{SNOMED-CT}$ , as shown below:

$$\begin{aligned} \mathcal{A} = \{ & \text{Exemplar}(p1), \text{Exemplar}(p2), \text{Exemplar}(p3), \text{Exemplar}(p4), \text{Exemplar}(p5), \text{hasFinding}(p1, s1), \\ & \text{hasFinding}(p2, s2), \text{hasFinding}(p3, s3), \text{hasFinding}(p3, s4), \text{hasFinding}(p4, s5), \text{hasFinding}(p4, s6), \\ & \text{hasFinding}(p5, s7), \text{hasFinding}(p5, s8), \text{SoreThroat}(s1), \text{Dyspnoea}(s2), \text{SoreThroat}(s3), \\ & \text{Dyspnoea}(s4), \text{SoreThroat}(s5), \text{DryCough}(s6), \text{SoreThroat}(s7), \text{LooseCough}(s8) \} \end{aligned}$$

$$\mathcal{T} = \{\text{LooseCough} \sqsubseteq \text{Cough}, \text{DryCough} \sqsubseteq \text{Cough}, \text{Cough} \sqsubseteq \text{LungFinding}, \text{LungFinding} \sqsubseteq \text{Finding}, \\ \text{Dyspnoea} \sqsubseteq \text{LungFinding}, \text{SoreThroat} \sqsubseteq \text{Finding}\}$$

Now suppose that a black-box classifier  $F$  predicts that p3, p4, and p5 should seek medical advice, while p1 and p2 do not need to (same as the motivating example of Section 3.1). The explanation rule

$$\rho_1 : \text{Exemplar}(x), \text{hasFinding}(x, y), \text{SoreThroat}(y), \text{hasFinding}(x, z), \text{LungFinding}(z) \rightarrow \text{MedicalAdvice}(x)$$

for that classifier based on the explanation dataset  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$  is a correct rule, as well as the explanation rule

$$\rho_2 : \text{Exemplar}(x), \text{hasFinding}(x, y), \text{Cough}(y) \rightarrow \text{MedicalAdvice}(x),$$

while the rules

$$\rho_3 : \text{Exemplar}(x), \text{hasFinding}(x, y), \text{SoreThroat}(y) \rightarrow \text{MedicalAdvice}(x) \quad \text{and}$$

$$\rho_4 : \text{Exemplar}(x), \text{hasFinding}(x, y), \text{Dyspnea}(y) \rightarrow \text{MedicalAdvice}(x)$$

are not correct, since  $\text{FOL}(\mathcal{S}') \models \rho_1$  and  $\text{FOL}(\mathcal{S}') \models \rho_2$ , but  $\text{FOL}(\mathcal{S}') \not\models \rho_3$  and  $\text{FOL}(\mathcal{S}') \not\models \rho_4$ , where  $\mathcal{S}' = \mathcal{S} \cup \{\text{Exemplar} \sqsubseteq \{a \mid a \in \text{EN}\} \cup \{\text{MedicalAdvice}(a) \mid a \in \text{pos}(F, \text{EN}, \text{MedicalAdvice})\}\}$ .

As mentioned in Section 2, a SAV query is an expression of the form  $\{c_1, \dots, c_n\}$ , an expression that resembles the body of explanation rules. By treating the bodies of explanation rules as queries, the problem of computing explanations can be solved as a query reverse engineering problem [46].

**Definition 3** ([46]). Let  $F : \mathcal{D} \rightarrow \mathcal{C}$  be a classifier,  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$  an explanation dataset in terms of  $\mathcal{D}$ ,  $\mathcal{C}$  and an appropriate vocabulary  $\mathcal{V}$ , and  $\rho(F, \mathcal{E}, C) : \text{Exemplar}(x), c_1, c_2, \dots, c_n \rightarrow C(x)$  an explanation rule. The SAV query

$$q_\rho \doteq \{\text{Exemplar}(x), c_1, c_2, \dots, c_n\}$$

is the explanation rule query of explanation rule  $\rho$ .

This query reverse engineering problem follows the *Query by Example* paradigm or QbE. This term refers to reverse engineering queries that contain some positive examples in their answer set but not any negative examples and is widely used in the related literature [20, 36, 47–50, 62]. We give a short formal definition of QbE adapted from [48].

**Definition 4** (Query by Example). Given a DL knowledge base  $\mathcal{S} = \langle \mathcal{T}, \mathcal{A} \rangle$ , two sets  $S^+, S^- \subseteq \text{IN}$  and a query language  $\mathcal{Q}$ . Is there a query  $q \in \mathcal{Q}$  such that  $S^+ \subseteq \text{cert}(q, \mathcal{S})$  and  $S^- \cap \text{cert}(q, \mathcal{S}) = \emptyset$ ?

In our case,  $S^+$  would be  $\text{pos}(F, \text{EN}, C)$  and  $S^- = \text{EN} \setminus \text{pos}(F, \text{EN}, C)$ , while  $\mathcal{Q}$ , the language of conjunctive queries.

The relationship between the properties of explanation rules and the respective queries allows us to detect and compute correct rules based on the certain answers of the respective explanation rule queries, as shown in Theorem 1.

**Theorem 1** ([46]). Let  $F : \mathcal{D} \rightarrow \mathcal{C}$  be a classifier,  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$  an explanation dataset in terms of  $\mathcal{D}$ ,  $\mathcal{C}$  and an appropriate vocabulary  $\mathcal{V}$ ,  $\rho(F, \mathcal{E}, C) : \text{Exemplar}(x), c_1, c_2, \dots, c_n \rightarrow C(x)$  an explanation rule, and  $q_\rho$  the explanation rule query of  $\rho$ . The explanation rule  $\rho$  is correct if and only if

$$\text{cert}(q_\rho, \mathcal{S}) \subseteq \text{pos}(F, \text{EN}, C)$$

*Proof.* Let  $\mathcal{S} = \langle \mathcal{T}, \mathcal{A} \rangle$ . Because by definition  $\text{Exemplar}(a) \in \mathcal{A}$  iff  $a \in \text{EN}$  and  $\text{Exemplar}$  does not appear anywhere in  $\mathcal{T}$ , we have

$$\text{cert}(q_\rho, \mathcal{S}) = \text{cert}(\{\text{Exemplar}, c_1, \dots, c_n\}, \langle \mathcal{T}, \mathcal{A} \rangle)$$

$$\begin{aligned}
&= \text{EN} \cap \text{cert}(\{c_1, \dots, c_n\}, \langle \mathcal{T}, \mathcal{A} \rangle) \\
&= \text{cert} \left( \left\{ \text{Exemplar} \right\}, \left\langle \left\{ \text{Exemplar} \sqsubseteq \bigsqcup_{a \in \text{EN}} \{a\} \right\}, \mathcal{A} \right\rangle \right) \cap \text{cert}(\{c_1, \dots, c_n\}, \langle \mathcal{T}, \mathcal{A} \rangle) \\
&= \text{cert} \left( \left\{ \text{Exemplar}, c_1, \dots, c_n \right\}, \left\langle \mathcal{T} \cup \left\{ \text{Exemplar} \sqsubseteq \bigsqcup_{a \in \text{EN}} \{a\} \right\}, \mathcal{A} \right\rangle \right) \\
&= \text{cert} \left( q_\rho, \left\langle \mathcal{T} \cup \left\{ \text{Exemplar} \sqsubseteq \bigsqcup_{a \in \text{EN}} \{a\} \right\}, \mathcal{A} \right\rangle \right)
\end{aligned}$$

Because by definition  $C$  does not appear anywhere in  $\mathcal{S}$ , we have also that  $\text{cert}(q_\rho, \mathcal{S}) = \text{cert}(q_\rho, \mathcal{S}')$ , where  $\mathcal{S}' = \mathcal{S} \cup \{ \text{Exemplar} \sqsubseteq \bigsqcup_{a \in \text{EN}} \{a\} \} \cup \{ C(a) \mid a \in \text{pos}(F, \text{EN}, C) \}$ , since the assertions  $C(a)$  are not involved neither in the query nor in  $\mathcal{S}$  and hence have no effect.

By definition of a certain answer,  $e \in \text{cert}(q, \mathcal{K})$  iff for every model  $\mathcal{I}$  of  $\mathcal{K}$  there is a match  $\pi$  s.t.  $\pi(x) = e^{\mathcal{I}}$  and  $\pi(u) \in D^{\mathcal{I}}$  for all  $D(u) \in q$  and  $(\pi(u), \pi(v)) \in r^{\mathcal{I}}$  for all  $r(u, v) \in q$ .

Assume that  $\rho$  is correct and let  $e \in \text{cert}(q_\rho, \mathcal{S})$ . We have proved that also  $e \in \text{cert}(q_\rho, \mathcal{S}')$ . Because  $\rho$  is correct, by Def. 3 it follows that every model  $\mathcal{I}$  of  $\mathcal{S}'$  is also a model of  $\rho$ . Because the body of  $q_\rho$  is the same as the body of  $\rho$ ,  $\pi$  makes true both the body of  $\rho$  and the head of  $\rho$ , which is  $C(x)$ , hence  $e^{\mathcal{I}} \in C^{\mathcal{I}}$ . It follows that  $C(e)$  is true in  $\mathcal{I}$ . But the only assertions of the form  $C(e)$  in  $\mathcal{S}'$  are the assertions  $\{ C(a) \mid a \in \text{pos}(F, \text{EN}, C) \}$ , thus  $e \in \text{pos}(F, \text{EN}, C)$ .

For the inverse, assume that  $\text{cert}(q_\rho, \mathcal{S}) \subseteq \text{pos}(F, \text{EN}, C)$ , equivalently  $\text{cert}(q_\rho, \mathcal{S}') \subseteq \text{pos}(F, \text{EN}, C)$ . Thus if  $e \in \text{cert}(q_\rho, \mathcal{S})$  then  $e^{\mathcal{I}} \in C^{\mathcal{I}}$ . Since this holds for every model  $\mathcal{I}$  of  $\mathcal{S}'$  and the body of  $q_\rho$  is the same as the body of  $\rho$ , it follows that  $\mathcal{I}$  is also a model of  $\rho$ , i.e.  $\rho$  is correct.  $\square$

Theorem 1 shows a useful property of the certain answers of the explanation rule query of a correct rule ( $\text{cert}(q, \mathcal{S}) \subseteq \text{pos}(F, \text{EN}, C)$ ) that can be utilized in order to identify and produce correct rules. Intuitively, an explanation rule is correct, if all of the certain answers of the respective explanation rule query are mapped by  $\mathcal{M}$  to data which is classified in the class indicated at the head of the rule. However, it is obvious that according to the above we can have correct rules that barely approximate the behaviour of the classifier (e.g. an explanation rule query with only one certain answer that is in the pos-set of the classifier), while other correct rules might exactly match the output of the classifier (e.g. a query  $q$  for which  $\text{cert}(q, \mathcal{S}) = \text{pos}(F, \text{EN}, C)$ ). Thus, it is useful to define a *recall* metric for explanation rule queries by comparing the set of certain answers with the pos-set of a class  $C$ :

$$\text{recall}(q, \mathcal{E}, C) = \frac{|\text{cert}(q, \mathcal{S}) \cap \text{pos}(F, \text{EN}, C)|}{|\text{pos}(F, \text{EN}, C)|},$$

assuming that  $\text{pos}(F, \text{EN}, C) \neq \emptyset$ .

**Example 2.** Continuing Example 1, we can create the explanation rule queries of the respective explanation rules of the example as follows:  $q_1(x) = \{ \text{Exemplar}(x), \text{hasFinding}(x, y), \text{SoreThroat}(y), \text{hasFinding}(x, z), \text{LungFinding}(z) \}$  as the explanation rule query of  $\rho_1$ ,  $q_2(x) = \{ \text{Exemplar}(x), \text{hasFinding}(x, y), \text{Cough}(y) \}$  as the explanation rule query of  $\rho_2$ ,  $q_3(x) = \{ \text{Exemplar}(x), \text{hasFinding}(x, y), \text{SoreThroat}(y) \}$  as the explanation rule query of  $\rho_3$ , and  $q_4(x) = \{ \text{Exemplar}(x), \text{hasFinding}(x, y), \text{Dyspnoea}(y) \}$  as the explanation rule query of  $\rho_4$ .

For the above queries we can retrieve their certain answers over our knowledge base  $\mathcal{S}$ , and get  $\text{cert}(q_1, \mathcal{S}) = \{ \text{p3}, \text{p4}, \text{p5} \}$ ,  $\text{cert}(q_2, \mathcal{S}) = \{ \text{p4}, \text{p5} \}$ ,  $\text{cert}(q_3, \mathcal{S}) = \{ \text{p1}, \text{p3}, \text{p4}, \text{p5} \}$ , and  $\text{cert}(q_4, \mathcal{S}) = \{ \text{p2}, \text{p3} \}$ .

With respect to the classifier  $F$  of Example 1, for which  $\text{pos}(F, \text{EN}, \text{MedicalAdvice}) = \{ \text{p3}, \text{p4}, \text{p5} \}$ , we see, as the theorem states, that for the correct rules  $\rho_1$  and  $\rho_2$  it holds that  $\text{cert}(q_1, \mathcal{S}) \subseteq \text{pos}(F, \text{EN}, \text{MedicalAdvice})$  and  $\text{cert}(q_2, \mathcal{S}) \subseteq \text{pos}(F, \text{EN}, \text{MedicalAdvice})$ , while for rules  $\rho_3$  and  $\rho_4$  that are not correct, it holds that  $\text{cert}(q_3, \mathcal{S}) \not\subseteq \text{pos}(F, \text{EN}, \text{MedicalAdvice})$ , and  $\text{cert}(q_4, \mathcal{S}) \not\subseteq \text{pos}(F, \text{EN}, \text{MedicalAdvice})$ , respectively.

The explanation framework described in this section provides the necessary expressivity to formulate accurate rules, using a desired terminology, even for complex problems [46]. However, some limitations of the framework, like only working with correct rules, can be a significant drawback for explanation methods built on top of that. An explanation rule query might not be correct due to the existence of individuals in the set of certain answers which are not in the pos-set. By viewing these individuals as exceptions to a rule, we are able to provide as an explanation a rule that is not correct, along with the exceptions which would make it correct if they were omitted from the explanation dataset; the exceptions could provide useful information to an end-user about the classifier under investigation. Thus, we extend the existing framework by introducing correct explanation rules with exceptions, as follows:

**Definition 5.** Let  $F : \mathcal{D} \rightarrow \mathcal{C}$  be a classifier,  $\mathcal{E} = \langle \mathcal{M}, \mathcal{S} \rangle$  an explanation dataset in terms of  $\mathcal{D}$ ,  $\mathcal{C}$  where  $\mathcal{S}$  is a knowledge base  $\mathcal{S} = \langle \mathcal{A}, \mathcal{T} \rangle$ ,  $\text{EN}$  the set of exemplars of  $\mathcal{E}$ , and let  $\text{EX}$  be a subset of  $\text{EN}$ . An explanation rule  $\rho(F, \mathcal{E}, C)$  is correct with exceptions  $\text{EX}$  for class  $C$  if the rule  $\rho(F, \mathcal{E}', C)$  is correct for class  $C$ , where  $\mathcal{E}' = \langle \mathcal{M}, \mathcal{S}' \rangle$ , and  $\mathcal{S}'$  is the knowledge base  $\mathcal{S}' = \langle \mathcal{A}', \mathcal{T} \rangle$ , and  $\mathcal{A}' = \mathcal{A} \setminus \{\text{Exemplar}(a) \mid a \in \text{EX}\}$ .

Since we allow exceptions to explanation rules, it is useful to define a measure of precision of the corresponding explanation rule queries as

$$\text{precision}(q, \mathcal{E}, C) = \frac{|\text{cert}(q, \mathcal{S}) \cap \text{pos}(F, \text{EN}, C)|}{|\text{cert}(q, \mathcal{S})|}.$$

if  $\text{cert}(q, \mathcal{S}) \neq \emptyset$  and  $\text{precision}(q, \mathcal{E}, C) = 0$  otherwise.

Obviously, if the precision of a rule query is 1, then it represents a correct rule, otherwise it is correct with exceptions. Furthermore, we can use the Jaccard similarity between the set of certain answers of the explanation rule query and the pos-set, as a generic measure which combines recall and precision to compare the two sets of interest as:

$$\text{degree}(q, \mathcal{E}, C) = \frac{|\text{cert}(q, \mathcal{S}) \cap \text{pos}(F, \text{EN}, C)|}{|\text{cert}(q, \mathcal{S}) \cup \text{pos}(F, \text{EN}, C)|}.$$

**Example 3.** The rules  $\rho_3$  and  $\rho_4$  of Example 2 that are not correct; they are correct with exceptions. Table 3 shows the precision, recall, and degree metrics of the explanation rule queries of Example 2 along with the exceptions  $\text{EX}$  of the respective rules.

Table 3  
Metrics and Exceptions of the example Explanation Rules and the respective Explanation Rule Queries.

Rule	Query	Recall	Precision	Degree	Exceptions (EX)
$\rho_1$	$q_1$	1.0	1.0	1.0	{}
$\rho_2$	$q_2$	0.67	1.0	0.67	{}
$\rho_3$	$q_3$	1.0	0.75	0.75	{p1}
$\rho_4$	$q_4$	0.33	0.5	0.25	{p2}

It can often be useful to present a set of rules to the end user instead of a single one. This could be, for example, because the pos-set cannot be captured by a single correct rule but only by a set of correct rules. This is a strategy commonly employed by other rule based systems such as those we employed in Sec. 5 to compare our algorithms to, namely RuleMatrix [26] and Skope-Rules [25]. The metrics already defined for a single query can be expanded to a set of queries  $Q = \{q_1, q_2, \dots, q_n\}$ , by simply expanding the definition of certain answers for a set of queries to  $\text{cert}(Q, \mathcal{S}) = \bigcup_{i=1}^n \text{cert}(q_i, \mathcal{S})$ .

As we have previously mentioned, we do not have any formal guarantees that the semantics expressed in the explanations extracted by the reverse engineering process are linked to the ones used by the classifier. In [24] the author proposes the term ‘‘summary of predictions’’ for approaches such as ours, that are independent of the classifier features, and the explanations are produced by summarizing the predictions of the classifier. We nevertheless use

the term explanation to be consistent with the existing literature. The measures we defined in this section provide some confidence about the existence of a link between the semantics of the classifier and those of the explanations. Assuming that the explanations produced are of high quality (high precision, low recall, short in length) and that the explanation dataset is constructed in a way that takes into account possible biases, it is reasonable to assume that there is at least some correlation between the semantics that the explanation presents with those that the classifier is using.

Unfortunately, as the semantic descriptions become richer, it is more likely that a query that perfectly separates the positive and negative individuals exists since the space of queries that can be expressed using the vocabulary of the descriptions becomes larger. This could lead to rules with high precision (and possibly high degree) that do not hold for new examples which could be perceived as a form of overfitting. This resembles the curse of dimensionality of machine learning, where the plethora of features leads to overfitting models when there is a lack of sufficient data. It would be helpful to have some statistical measures after the reverse engineering process that express how likely it is that the query(ies) perform well by coincidence, rather than due to a link between the semantics expressed in the explanation dataset and the ones used by the classifier. Computing such measures is far from trivial due to the complex structure of the Query Space [46]. In this work, we employ some alternatives in our experiments. Namely, using a holdout set in order to check how well the rules produced generalize (Mushrooms experiment, section 5.1), and employing other XAI methods to cross-reference the explanations produced (Visual Genome experiment, section 5.3). We consider a thorough examination of such methods to be out of the scope of this work.

#### 4. Computation of Explanations

From Section 3, we understand that the problem that we try to solve is closely related to the query reverse engineering problem, since we need to compute queries given a set of individuals. However, since in most cases there is not a single query that fits our needs (have as certain answers the pos-set of the classifier), we need to find (out of all the semantic queries that have a specific certain answer set) a set of queries that accurately describe the classifier under investigation. Therefore, the problem can also be seen as a heuristic search problem. The duality of rules and queries within our framework, reduces the search of correct rules (with exceptions) to the search of queries that contain elements of the pos-set in their certain answers. Reverse engineering queries for subsets of EN is challenging for the following reasons:

- The subsets of EN for which there exists a correct rule query ( $\{I \mid I \subseteq \text{EN}, \text{there exists } q \text{ s.t. } \text{cert}(q, \mathcal{S}) = I\}$ ) can potentially be exponentially many ( $2^{|\text{EN}|}$ ).
- The Query Space i.e. the set containing all queries that have non-empty certain answer set ( $\{q \mid \text{cert}(q, \mathcal{S}) \cap \text{EN} \neq \emptyset\}$ ) can potentially be infinite [46].
- For any subset  $I$  of EN, the number of queries s.t.  $\text{cert}(q, \mathcal{S}) = I$  can be zero, positive or infinite.
- Computing the certain answers of arbitrary queries can be exponentially slow, so it is computationally prohibitive to evaluate each query under consideration while exploring the query space.

The difficulties described above are addressed in the following ways:

- In this paper, we only consider knowledge bases of which the TBox can be eliminated (such as RL [63]; see also the last paragraph of this section). This enables us to create finite queries that contain all the necessary conjuncts to fully describe each individual. We are then able to merge those queries to create descriptions of successively larger subsets of EN.
- We do not directly explore the subsets of EN for which there exists a correct rule query, and the computation of certain answers is not required for the algorithms. Instead, we explore the Query Space blindly, but heuristically. We create queries that are guaranteed to be within the Query Space and are also guaranteed to contain heuristically selected individuals in their certain answers, without knowing their exact certain answers. The heuristic we employ, aids us in selecting similar individuals to merge. Intuitively, this helps us create queries that introduce as few as possible unwanted certain answers.

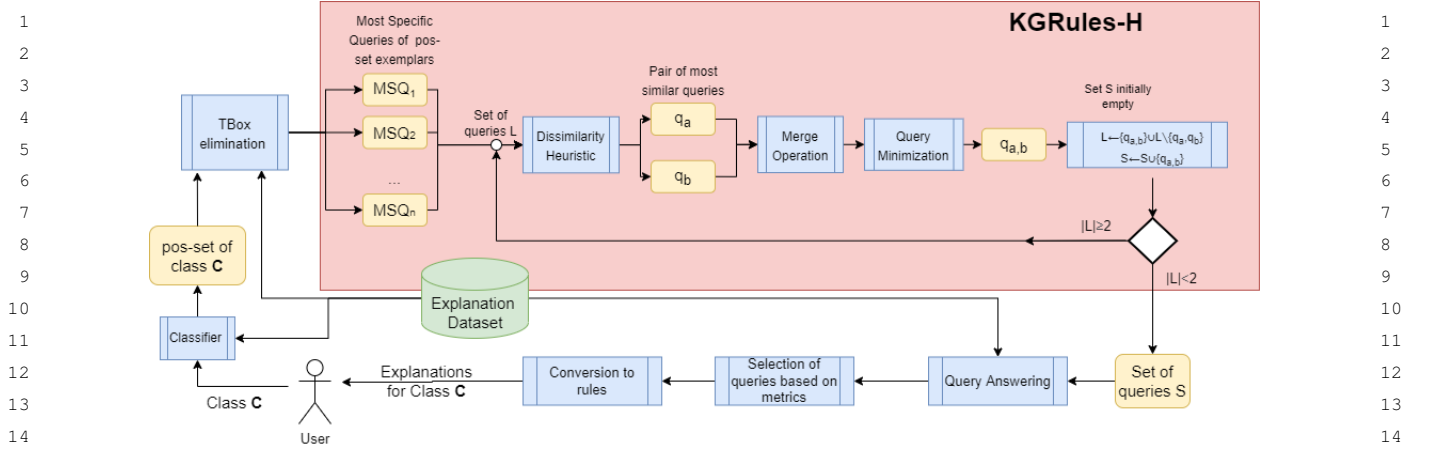


Fig. 2. Visualization of how KGRules-H is integrated into our framework.

- We are not concerned with the entire set of queries with non-empty sets of certain answers, but only with queries which have specific characteristics in order to be used as explanations. Specifically, the queries have to be short in length, with no redundant conjuncts, have as certain answers elements of the pos-set of the class under investigation, and as few others as possible.
- One of the proposed algorithms guarantees that given a set  $I$ , if a query  $q$  exists s.t.  $\text{cert}(q, \mathcal{S}) = I$ , then it will find at least one such query. If there do not exist such queries then, since the computation of certain answers is not involved in the algorithms, the result of the heuristic search will be a “good guess” of queries which have similar certain answer sets to  $I$ . If there exist infinite such queries, then we do not have a guarantee that we have found the shortest, most understandable one, however the proposed algorithms take care to create queries with few variables (see Alg. 2 and section 4.3.2).

In the following we describe the proposed algorithms for computing explanations. The core algorithm, which is outlined as Alg. 1 and we call KGRules-H, takes as input an atomic ABox  $\mathcal{A}$  and a set of individual names  $I$ , and produces a list of queries. It is assumed that both  $\mathcal{A}$  and  $I$  are obtained from an explanation dataset  $\mathcal{E}$ . In particular,  $I$  is a subset of the respective EN corresponding to a pos-set of a classifier  $F$  to be explained for some class  $C$ , i.e.  $\text{pos}(F, \text{EN}, C)$ , and  $\mathcal{A}$  is an atomic ABox containing all the available knowledge about the individuals in  $I$  encoded in the knowledge base of  $\mathcal{E}$ . The output queries are intended to serve as explanation rule queries for class  $C$ . An outline of how KGRules-H is used to generate explanations is shown in Fig. 2.

---

### Algorithm 1: KGRULES-H

---

**Input:** An atomic ABox  $\mathcal{A}$  and a set of individual names  $I$ .

**Output:** A list of queries  $S$ .

```

1  $S \leftarrow []$ 
2  $L \leftarrow \{\text{MSQ}(a, \mathcal{A}) \mid a \in I\}$ 
3 while  $|L| \geq 2$  do
4    $q_A, q_B \leftarrow \arg \min_{q, q' \in L, q \neq q'} \text{QueryDissimilarity}(q, q')$ 
5    $q \leftarrow \text{Merge}(q_A, q_B)$ 
6    $L \leftarrow (L \setminus \{q_A, q_B\}) \cup \{q\}$ 
7   append  $q$  to  $S$ 
8 end
9 return  $S$ 

```

---

The algorithm starts by initializing an empty list of queries  $S$ , and by creating an initial description for each individual in  $I$  in the form of a most specific query (MSQ). A detailed definition of a MSQ is given in Section

4.1; intuitively, an MSQ of an individual  $a$  for an ABox  $\mathcal{A}$  is a SAV query  $q$  that captures the maximum possible information about  $a$  and is such that  $a \in \text{cert}(\text{MSQ}(q, \mathcal{A}))$ . Given these descriptions, one for each individual, which are added in a set  $L$ , the algorithm then tries to combine the elements of  $L$  in order to generate more general descriptions. In particular, at each iteration of the while loop, it selects two queries from  $L$  and merges them. The queries to be merged are selected based on their dissimilarity; two most similar queries are selected and merged. The intuition is that by merging relatively similar queries, the resulting queries will continue to be “as specific as possible” since they will generalize out the limited dissimilarities of the original queries. In Section 4.2 we discuss how dissimilarity is estimated, and in Section 4.3 we describe two alternative approaches for merging queries. Given two queries  $q_A, q_B$  that have been selected as most similar, the algorithm merges them by constructing a new SAV query  $q = \text{Merge}(q_A, q_B)$  such that  $\text{cert}(q, \mathcal{A}) \supseteq \text{cert}(q_A, \mathcal{A}) \cup \text{cert}(q_B, \mathcal{A})$ . The newly created query replaces the ones it was merged from in  $L$ , and is also appended to  $S$ . Once the queries to be merged have been exhausted, Alg. 1 terminates by returning the list  $S$  which will contain  $|I| - 1$  SAV queries, the results of each merging step in order of creation. Thus, the queries can be considered to be in  $S$  in some decreasing order of “specificness”, although the actual order depends on the order the elements in  $L$  are considered.

If the queries in  $S$  have subsets of  $I$  as their certain answers, given that we have assumed that  $I$  is the pos-set of a classifier for some class  $C$ , the queries can be treated as candidate explanation rule queries for  $C$ . In this case, the queries can be interpreted as explanation rule queries, converted to the respective explanation rules, and presented as explanations. In general, however, there is no guarantee that the certain answers of a merged query will be a subset of  $I$ . In this case, which is the typical case, the computed explanation rules will be rules with exceptions. The queries produced will differ with respect to the metrics introduced in Section 3. Specifically, the queries produced early on, which are the result of few merges are expected to have low recall but high precision, with the inverse holding for queries produced in the later iterations. The importance placed on each metric can be different for each end user, therefore there is no single best query produced by Alg. 1. Instead, metrics for each query should be produced and presented to the end user to let them rank the explanations according to their preferences. This is the methodology we have employed in Sec. 5. For each experiment conducted we present a selection of queries, each performing better for a different metric.

In Section 4.3 we prove some optimality results for one of the merging methods, the query least common subsumer (QLCS). In particular, if there is a correct explanation rule (without exceptions), then we are guaranteed to find the corresponding explanation rule query, using the QLCS. This optimality property does not hold for the Greedy Matching operation defined in section 4.3.2, nor for Alg. 3.

As mentioned above, Alg. 1 works on the assumption that all available knowledge about the relevant individuals  $I$  is encoded in a (finite) atomic ABox  $\mathcal{A}$ . This is essential for the computation of the MSQs. Given that the knowledge base associated with an explanation dataset is in general of the form  $\langle \mathcal{T}, \mathcal{A} \rangle$ , this assumption means that if the original knowledge involves a non-empty TBox  $\mathcal{T}$ , it has to be eliminated before applying Alg. 1. In particular, *TBox elimination* is the process of expanding a given ABox by applying the axioms of a given TBox, so that the TBox is eventually not needed, in the sense that all certain answers of the original knowledge can now be obtained only from the expanded ABox. Because TBox elimination is not always possible, this poses certain restrictions on the applicability of Alg. 1, namely that if the original knowledge of the explanation dataset is indeed of the form  $\langle \mathcal{T}, \mathcal{A} \rangle$  with  $\mathcal{T} \neq \emptyset$ , it should be possible, as outlined above, to be transformed through TBox elimination to an equivalent w.r.t. query answering finite assertional-only knowledge base  $\mathcal{A}'$ , such that  $\text{cert}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \text{cert}(q, \langle \emptyset, \mathcal{A}' \rangle)$  for any  $q$ . For knowledge bases that this is possible, the standard approach for TBox elimination is *materialization* and is typically performed by first encoding the axioms in  $\mathcal{T}$  as a set of inference rules generating ABox assertions, and then iteratively applying them on the knowledge base until no more assertions can be generated. Materialization is possible, e.g. for Description Logic Programs and the Horn-*SHIQ* DL dialect. Finite materialization may not be possible even for low expressivity DL dialects, such as DL-Lite. [64, 65]

#### 4.1. Most Specific Queries

As mentioned above, intuitively, a *most specific query* (MSQ) of an individual  $a$  for an atomic ABox  $\mathcal{A}$  encodes the maximum possible information about  $a$  provided by  $\mathcal{A}$ . It is defined as a least query in terms of subsumption

which has  $a$  as a certain answer; in particular, a query  $q$  is a MSQ of an individual  $a$  for an atomic ABox  $\mathcal{A}$  iff  $a \in \text{cert}(q, \mathcal{A})$  and for all  $q'$  such that  $a \in \text{cert}(q', \mathcal{A})$  we have  $q \leq_s q'$ .

Following the subsumption properties, MSQs are unique up to syntactical equivalence. In Alg. 1 it is assumed that  $\text{MSQ}(a, \mathcal{A})$  represents one such MSQ of  $a$  for  $\mathcal{A}$ . Since  $\mathcal{A}$  is finite, a MSQ is easy to compute. This can be done by viewing  $\mathcal{A}$  as a graph, taking the connected component of that graph that contains node  $a$ , and converting it into a graph representing a query by replacing all nodes (which in the original graph represent individuals) by arbitrary, distinct variables, taking care to replace node  $a$  by variable  $x$ .

**Theorem 2.** *Let  $\mathcal{A}$  be a finite atomic ABox over a vocabulary  $\langle \text{CN}, \text{RN}, \text{IN} \rangle$ ,  $a$  an individual in  $\text{IN}$ , and  $\text{conn}(a, \mathcal{A})$  the connected component of the ABox graph of  $\mathcal{A}$  containing  $a$ . Let also  $q$  be a SAV query with query graph  $G$ . If there exists an isomorphism  $f : G \rightarrow \text{conn}(a, \mathcal{A})$  such that  $f(x) = a$ , then  $q$  is an MSQ of  $a$  for  $\mathcal{A}$ .*

*Proof.* Let  $\mathcal{I}$  be the interpretation such that  $\Delta^{\mathcal{I}} = \text{IN}$ ,  $a^{\mathcal{I}} = a$  for all  $a \in \text{IN}$ ,  $C^{\mathcal{I}} = \{a \mid C(a) \in \mathcal{A}\}$ , for all  $C \in \text{CN}$ , and  $r^{\mathcal{I}} = \{(a, b) \mid r(a, b) \in \mathcal{A}\}$  for all  $r \in \text{RN}$ . Clearly,  $\mathcal{I}$  is a model of  $\mathcal{A}$ , and it is easy to see that  $\text{cert}(q, \mathcal{A}) = \text{ans}_{\mathcal{I}}(q, \mathcal{A})$  for any query  $q$ , where  $\text{ans}_{\mathcal{I}}(q, \mathcal{A})$  are the answers to  $q$  under the interpretation  $\mathcal{I}$ . Because of this identity, we can, without loss of generality, restrict ourselves to this  $\mathcal{I}$  w.r.t.  $\text{cert}(q, \mathcal{A})$ .

Since  $f$  exists, we can construct for  $q$  the match  $\pi : \text{var}(q) \rightarrow \text{IN}$ , such that  $\pi(u) = f(u)^{\mathcal{I}}$ . It follows that  $\pi(x) = a$ , and so  $a \in \text{cert}(q, \mathcal{A})$ . Let  $q'$  be a query with query graph  $G'$  such that  $a \in \text{cert}(q', \mathcal{A})$ . Since  $a$  is a certain answer of  $q'$ ,  $G'$  must be homomorphic to  $\mathcal{A}$  by a homomorphism  $g$  such that  $g(x) = a$ . Since homomorphisms preserve edges and  $q'$  is connected, the image of  $G'$  under  $g$  is a subgraph of  $\text{conn}(a, \mathcal{A})$ . Then  $h = g \circ f^{-1}$  is a homomorphism from  $G'$  to  $G$  with  $h(x) = x$ . Thus,  $q \leq_s q'$ . □

Since  $q \leq_s q'$  implies that  $\text{cert}(q, \mathcal{A}) \subseteq \text{cert}(q', \mathcal{A})$ , the MSQ of  $a$  has as few as possible certain answers other than  $a$ . This is desired since it implies that the initial queries in Alg. 1 do not have any exceptions (certain answers not in the pos-set), unless this cannot be avoided. Therefore, the iterative query merging process of constructing explanations starts from an optimal standpoint.

**Example 4.** *Continuing the previous examples, Fig. 3 shows the graphs of the MSQs of p3 and p4, constructed by materializing the TBox, and then extracting the connected components of p3 and p4 from the new ABox.*




Fig. 3. The MSQs of p3 and p4 represented as graphs

For use with Alg. 1, we denote a call to a concrete function implementing the above described approach for obtaining  $\text{MSQ}(a, \mathcal{A})$  by  $\text{GRAPHMSQ}(a, \mathcal{A})$ . We should note that the result of  $\text{GRAPHMSQ}(a, \mathcal{A})$  will be a MSQ, but in general it will not be condensed, and thus may contain redundant atoms and variables.



## 4.2. Query Dissimilarity

At each iteration, Alg. 1 selects two queries to merge in order to produce a more general description covering both queries. To make the selection, we use a heuristic which is meant to express how dissimilar two queries are, so that at each iteration the two most similar queries are selected and merged, with the purpose of generating an as least generic as possible more general description. Given two queries  $q_1, q_2$  with respective graph representations  $G_1 = (V_1, E_1, \ell_{V_1}, \ell_{E_1})$  and  $G_2 = (V_2, E_2, \ell_{V_2}, \ell_{E_2})$ , we define the *query dissimilarity* heuristic between  $q_1$  and  $q_2$  as follows:

$$\text{QueryDissimilarity}(q_1, q_2) = \sum_{v_1 \in V_1} \min_{v_2 \in V_2} \text{diss}_{q_1 q_2}(v_1, v_2) + \sum_{v_2 \in V_2} \min_{v_1 \in V_1} \text{diss}_{q_2 q_1}(v_2, v_1)$$

where

$$\begin{aligned} \text{diss}_{q_1 q_2}(v_1, v_2) &= |L_1(v_1) \setminus L_2(v_2)| \\ &+ \sum_{r \in R} \{\max(\text{indegree}_{G_1}^r(v_1) - \text{indegree}_{G_2}^r(v_2), 0) + \max(\text{outdegree}_{G_1}^r(v_1) - \text{outdegree}_{G_2}^r(v_2), 0)\}, \end{aligned}$$

$R$  is the set of all role names appearing in the edge labels of the two graphs, and  $\text{indegree}_G^r(v)$  ( $\text{outdegree}_G^r(v)$ ) is the number of incoming (outcoming) edges  $e$  in node  $v$  of graph  $G$  with  $r \in \ell(e)$ .

The intuition behind the above dissimilarity measure is that the graphs of queries which are dissimilar consist of nodes with dissimilar labels connected in dissimilar ways. Intuitively, we expect such queries to have dissimilar sets of certain answers, although there is no guarantee that this will always be the case. In order to have a heuristic that can be computed efficiently, we do not examine complex ways in which the nodes may be interconnected, but only examine the local structure of the nodes by comparing their indegrees and outdegrees. The way in which we compare the nodes of two query graphs is optimistic; we compare each node with its best possible counterpart—the node of the other graph which it is the most similar to. Note that  $\text{diss}_{q_1 q_2}(v_1, v_2)$  does not equal  $\text{diss}_{q_2 q_1}(v_2, v_1)$ . The first quantity expresses how many conjuncts of  $q_1$  containing  $v_1$  do not appear in  $q_2$  containing  $v_2$ . This is best illustrated using a short example.

**Example 5.** Let  $q_1 = \{C(x), D(x), r(x, y_{11}), r(y_{12}, x), s(x, y_{11})\}$ ,  $q_2 = \{C(x), E(x), r(x, y_{21}), r(x, y_{23})\}$  and assume that we are comparing variable  $x$  of  $q_1$  with variable  $x$  of  $q_2$ . As far as concepts are concerned,  $C(x)$  appears in both  $q_1$  and  $q_2$ , while  $D(x)$  only appears in  $q_1$ , so one concept conjunct is missing from  $q_2$ . Moreover,  $x$  appears in one “outgoing”  $r$  conjunct in  $q_1$  and in two outgoing  $r$  conjuncts in  $q_2$ , so no outgoing conjuncts are missing from  $q_2$ . Also,  $x$  appears in one “incoming”  $r$  conjunct in  $q_1$  but in no “incoming”  $r$  in  $q_2$  conjuncts, so one more conjunct is missing. Finally, the  $s$  conjunct of  $x$  is missing from  $q_2$ , giving us a total of three missing conjuncts from  $q_2$ , therefore  $\text{diss}_{q_1 q_2}(x_1, x_2) = 3$ .

Using an efficient representation of the queries, it is easy to see that  $\text{QueryDissimilarity}(q_1, q_2)$  can be computed in time  $O(|\text{var}(q_1)| \cdot |\text{var}(q_2)|)$ .

## 4.3. Query Merging

The next step in Alg. 1 is the merging of the two most similar queries that have been selected. In particular, given a query  $q_A$  for some individuals  $I_A$ , i.e. a query such that  $I_A \subseteq \text{cert}(q_A, \mathcal{A})$ , and a query  $q_B$  for some individuals  $I_B$ , i.e. a query such that  $I_B \subseteq \text{cert}(q_B, \mathcal{A})$ , the purpose is to merge the two queries to produce a more general query  $q$  for  $I_A \cup I_B$ . The necessary condition that such query must satisfy is  $I_A \cup I_B \subseteq \text{cert}(q, \mathcal{A})$ . In the context of Alg. 1 the queries  $q_A, q_B$  may be MSQs of some individuals, or results of previous query merges.




Fig. 4. The Kronecker product of the MSQs of p3 and p4

#### 4.3.1. Query Least Common Subsumer

Our first approach for merging queries is using the query least common subsumer (QLCS). In the next few paragraphs we will present how to construct such a query from two SAV queries which will make apparent that the QLCS of two SAV queries always exists. As mentioned in Section 2, a QLCS is a most specific generalization of two queries. By choosing to use  $QLCS(q_A, q_B)$  (since QLCS is unique up to syntactical equivalence, we assume that  $QLCS(q_A, q_B)$  represents one such QLCS for  $q_A$  and  $q_B$ ) in the place of  $Merge(q_A, q_B)$  in Alg. 1, and assuming that the input set of individuals is  $I = \{a_1, \dots, a_n\}$ , it should be obvious that the last query computed by Alg. 1 will be  $q = QLCS(MSQ(a_1, \mathcal{A}), MSQ(a_2, \mathcal{A}), \dots, MSQ(a_n, \mathcal{A}))$ . For any query  $q'$  such that  $\text{cert}(q', \mathcal{A}) \supseteq I$ , it will by definition hold that  $MSQ(a_i, \mathcal{A}) \leq_S q'$ ,  $i = 1, \dots, n$ . Therefore,  $q \leq_S q'$ , which implies that  $\text{cert}(q, \mathcal{A}) \subseteq \text{cert}(q', \mathcal{A})$ . This means that if there is an exact explanation rule query of  $I$ , that will be a QLCS of its MSQs, while if there is not, the computed QLCS explanation rule query will have the fewest possible exceptions. Thus, the main advantage of using the QLCS is that it guarantees “optimality” of the computed explanations.

To compute a QLCS we use an extension of the Kronecker product of graphs to labeled graphs. Given two labeled graphs  $G_1 = (V_1, E_1, \ell_{V_1}, \ell_{E_1})$  and  $G_2 = (V_2, E_2, \ell_{V_2}, \ell_{E_2})$ , which in our case will represent queries, their Kronecker product  $G = G_1 \times G_2$  is the graph  $G = (V, E, \ell_V, \ell_E)$ , where  $V = V_1 \times V_2$ ,  $E = \{ \langle \langle u_1, v_1 \rangle, \langle u_2, v_2 \rangle \rangle \mid \langle u_1, u_2 \rangle \in E_1, \langle v_1, v_2 \rangle \in E_2, \text{ and } \ell_{E_1}(\langle u_1, u_2 \rangle) \cap \ell_{E_2}(\langle v_1, v_2 \rangle) \neq \emptyset \} \subseteq V \times V$ ,  $\ell_V : V \rightarrow 2^{\text{CN}}$  with  $\ell_V(\langle v_1, v_2 \rangle) = \ell_{V_1}(v_1) \cap \ell_{V_2}(v_2)$ , and  $\ell_E : E \rightarrow 2^{\text{RN}}$  with  $\ell_E(\langle \langle u_1, v_1 \rangle, \langle u_2, v_2 \rangle \rangle) = \ell_{E_1}(\langle u_1, u_2 \rangle) \cap \ell_{E_2}(\langle v_1, v_2 \rangle)$ .

**Example 6.** Fig. 4 shows the Kronecker product of the query graphs of Example 4.

As with the Kronecker product of unlabeled graphs, for any graph  $H$ , we have that  $H \rightarrow G_1, G_2$  if and only if  $H \rightarrow G_1 \times G_2$  [66]. Since we are interested in graphs representing SAV queries, we can assume that  $H, G_1, G_2$  are connected. Assuming that  $G_1, G_2$  represent two SAV queries  $q_1, q_2$  (with answer variable  $x$ ), it will specifically hold that  $H \rightarrow G_1, G_2$  if and only if  $H \rightarrow \text{conn}(\langle x, x \rangle, G_1 \times G_2)$ , where  $\text{conn}(\langle x, x \rangle, G_1 \times G_2)$  is the connected component of  $G_1 \times G_2$  containing the pairs of answer variables node  $\langle x, x \rangle$ . Let  $q$  be the query whose query graph is the connected component of  $G_1 \times G_2$  containing node  $\langle x, x \rangle$ , with node  $\langle x, x \rangle$  replaced by  $x$  (the answer variable of the new query) and all other nodes of that connected component replaced by distinct arbitrary variable names other than  $x$ . Since we allow SAV queries to have empty bodies, even if the node  $\langle x, x \rangle$  is not labeled with any atoms it corresponds to a valid query that obviously subsumes  $q_1$  and  $q_2$ . Because homomorphisms between query graphs correspond to subsumption relations between the respective queries, and vice versa, it will hold that  $q_1, q_2 \leq_S q$  and for any other query  $q'$  s.t.  $q_1, q_2 \leq_S q'$  it will hold that  $G_{q'} \rightarrow G_1, G_2$ , therefore  $G_{q'} \rightarrow G_1 \times G_2$  and therefore  $q \leq_S q'$ . Computing  $G_1 \times G_2$  is therefore a concrete way to construct the QLCS.

For use with Alg. 1, we denote a call to the concrete function implementing that approach for computing  $QLCS(q_1, q_2)$  by  $KRONECKERQLCS(q_1, q_2)$ . The time complexity of  $KRONECKERQLCS(q_1, q_2)$  is  $O(|\text{var}(q_1)|^2 \cdot |\text{var}(q_2)|^2)$ , and the Kronecker product of the graphs  $G_1, G_2$ , contains  $|V_1| \cdot |V_2|$  nodes. Therefore, the final query constructed by Alg. 1 using  $KRONECKERQLCS(q_1, q_2)$  to implement  $Merge(q_1, q_2)$  will have  $O(m^n)$  variables, where  $m = \max_{i=1, \dots, n} |\text{var}(q_i)|$  is the maximum variable count of the MSQs, and  $n = |I|$  is the number of individuals in  $I$ . Constructing this query is by far the most expensive operation of the algorithm and dominates the running time with a complexity of  $O(m^{2n})$ . This makes the use of the above procedure for computing QLCS's prohibitive




Fig. 5. The QLCS of the MSQs of p3 and p4 produced by KRONECKERQLCS represented as a graph

for merging queries. However, the query obtained using the Kronecker product is in general not condensed, and in general will contain many redundant atoms and variables.

**Example 7.** Fig. 5 shows the QLCS of the MSQs of Example 4, extracted from the Kronecker product of Example 6. It is obvious that the nodes with label  $\{\text{Finding}\}$  are redundant due to the presence of the nodes with label  $\{\text{SoreThroat, Finding}\}$  and  $\{\text{LungFinding, Finding}\}$ .

Removing the redundant parts of a query is essential not only for reducing the running time of the algorithm. Since these queries are intended to be shown to humans as explanations, ensuring that they are compact is imperative to improve comprehensibility. As mentioned in Section 2, the operation that compacts a query by creating a syntactically equivalent query by removing all its redundant parts is condensation. However, condensing a query is coNP-complete [53]. For this reason, we utilize Alg. 2, an approximation algorithm which removes redundant conjuncts and variables without a guarantee of producing a fully condensed query.

Alg. 2 iterates through the variables of the input query, and checks if deleting one of them is equivalent to unifying it with another one. In particular, at each iteration of the main loop, the algorithm attempts to find a variable suitable for deletion. If none is found, the loop terminates. The inner loop iterates through all pairs of variables. At each iteration, variable  $v'$  is deleted if unifying it with  $v$ , by replacing all instances of  $v'$  in the query with  $v$ , produces no new conjuncts. By unifying variable  $v'$  with  $v$ , all conjuncts of the form  $C(v')$  become  $C(v)$  and all conjuncts of the forms  $r(v', v'')$ ,  $r(v'', v')$ ,  $r(v', v')$  become  $r(v, v'')$ ,  $r(v'', v)$ ,  $r(v, v)$ , respectively. If these conjuncts are already present, then removing them is equivalent to unifying  $v'$  with  $v$ . Alg. 2 is correct because removing a variable from query  $q$  produces a query that subsumes  $q$ , while unifying two variables of  $q$  produces a query that is subsumed by  $q$ . Therefore, Alg. 2 produces syntactically equivalent queries.

---

**Algorithm 2:** APPROXQUERYMINIMIZE
 

---

**Input:** A query represented as a graph  $q = \langle V, E, \ell_V, \ell_E \rangle$ .

**Output:** An approximately minimized query  $q'$ , also represented as a graph.

```

1 do
2    $q' \leftarrow q$ 
3   foreach pair  $(v, v'), v, v' \in V, v \neq v'$  do
4     if  $\ell_V(v') \subseteq \ell_V(v)$  and
5        $\langle v', v'' \rangle \in E \Rightarrow (\langle v, v'' \rangle \in E, \ell_E(v', v'') \subseteq \ell_E(v, v''))$ ,  $v'' \neq v'$  and
6        $\langle v'', v' \rangle \in E \Rightarrow (\langle v'', v \rangle \in E, \ell_E(v'', v') \subseteq \ell_E(v'', v))$ ,  $v'' \neq v'$  and
7        $\langle v', v' \rangle \in E \Rightarrow (\langle v, v \rangle \in E, \ell_E(v', v') \subseteq \ell_E(v, v))$  then
8       Delete variable  $v'$  from  $q$ .
9     end
10  end
11 while  $q' \neq q$ 
12 return  $q'$ 
  
```

---

**Example 8.** We can showcase an example of Alg. 2 detecting an extraneous variable in the QLCS of  $q_1$  and  $q_2$  from Example 7, running the main loop for  $v = y_1$  and  $v' = y_2$ . The condition at line 3 first checks if the label of node  $v'$ ,  $\{\text{Finding}\}$ , is a subset of the label of node  $v$ ,  $\{\text{SoreThroat, Finding}\}$ , which is true. Then for every incoming edge to  $v'$  it checks if there is an edge incoming to  $v$  with the same label and origin. They both have one incoming edge from node  $x$  labeled with  $\text{hasFinding}$  so the second condition will also evaluate to true. Node  $v'$  has no incoming edges, so the third condition will also evaluate to true. Therefore, Alg. 2 will detect node  $v'$  as extraneous and delete it. Running the entire algorithm for this query will result in the query represented in Fig. 6 as a graph.




Fig. 6. The minimized graph of the QLCS of Fig. 5

The main loop of Alg. 2 is executed at most  $\text{var}(q)$  times, since at each loop either a variable is deleted or the loop terminates. The inner loop checks all pairs of variables ( $O(|\text{var}(q)|^2)$ ), and the if condition requires  $O(\text{var}(q))$  set comparisons of size at most  $|\text{CN}|$ , and  $2|\text{RN}|$  comparisons of rows and columns of adjacency matrices. Treating  $|\text{CN}|$  and  $|\text{RN}|$  as constants, the complexity of Algorithm 2 is  $O(|\text{var}(q)|^4)$ .

Given the above, our first practical implementation of Alg. 1, uses  $\text{GRAPHMSQ}(a, \mathcal{A})$  to implement  $\text{MSQ}(a, \mathcal{A})$ , and  $\text{APPROXQUERYMINIMIZE}(\text{KRONECKERQLCS}(q_A, q_B))$  to implement  $\text{Merge}(q_A, q_B)$ .

Regardless of the approximate query minimization described above, even if full query condensation could be performed, there is no guarantee that condensing the queries computed by Alg. 1 using QLCS as the merge operation, have meaningfully smaller condensations. This led us to two different approaches for further dealing with the rapidly growing queries QLCS produces.

The simplest approach to address this problem is to reject any queries produced by  $\text{KRONECKERQLCS}(q_1, q_2)$  that, even after minimization, have variable counts higher than a pre-selected threshold. This strategy essentially introduces a weaker version of Alg. 1, without the guarantees of optimality, but with polynomial running time, which is outlined in Alg. 3, which we call KGRules-HT.

Alg. 3 is the same as Alg. 1 except that  $q$  is not added to  $L$  and  $S$  unless its variable count is less than or equal to an input threshold,  $t$ . This simple change reduces the complexity of the algorithm to polynomial in terms of  $|I|$  and  $t$ . It should be noted that setting a very low threshold for  $t$  could potentially lead to all queries being rejected and the algorithm returning an empty set. The effects of  $t$  on the queries produced can vary significantly for different inputs and therefore different explanation datasets, thus  $t$  should be adjusted experimentally. It is also assumed that all MSQ's created when initializing  $L$  have less than  $t$  variables.

If we let  $n = |I|$ , calculating the query dissimilarity for all pairs of queries, costs  $O(n^2 t^2)$  operations. Using memoization limits the cost of calculating the dissimilarity through all iterations to  $O(n^2 t^2)$ . Each iteration of the main loop involves  $O(n^2)$  operations for selecting  $q_A, q_B$  and  $O(t^4)$  operations for calculating their QLCS. Thus, in total, Alg. 3 has a running time of  $O(n^3 + n^2 t^2 + n t^4)$ .

Our second approach to overcome the problem posed by the rapidly growing size of the queries produced by QLCS was to consider a different merge operation, which is described in the following section.

#### 4.3.2. Greedy Matching

As already mentioned, the above described procedure for merging queries by computing a QLCS, often introduces too many variables that Alg. 2 can't minimize effectively. A QLCS of  $q_1, q_2$ , if condensed, is optimal at producing a query that combines their certain answers while introducing few as possible new variables. However, the new

**Algorithm 3: KGrules-HT****Input:** An atomic ABox  $\mathcal{A}$ , a set of individual names  $I$  and a threshold  $t$ .**Output:** A set of queries  $S$ .

---

```

1   $S \leftarrow \emptyset$ 
2   $L \leftarrow \{\text{GRAPHMSQ}(a, \mathcal{A}) \mid a \in I\}$ 
3  while  $|L| \geq 2$  do
4       $q_A, q_B \leftarrow \arg \min_{q, q' \in L, q \neq q'} \text{QueryDissimilarity}(q, q')$ 
5       $q \leftarrow \text{APPROXQUERYMINIMIZE}(\text{KRONECKERQLCS}(q_A, q_B))$ 
6       $L \leftarrow (L \setminus \{q_A, q_B\})$ 
7      if  $|\text{var}(q)| \leq t$  then
8           $L \leftarrow L \cup \{q\}$ 
9          append  $q$  to  $S$ 
10     end
11 end
12 return  $S$ 

```

---

variables that are introduced may still be too many. In practice, we may be able to afford some leniency in terms of certain answers and the loss of optimality that the QLCS provides, so we can explore alternative methods of unifying queries without the cost of increased variables.

One such method, is finding a *common subquery*, i.e. finding the common conjuncts of two queries  $q_1, q_2$ . Assuming that  $|\text{var}(q_1)| \geq |\text{var}(q_2)|$  and that the variables of  $q_1, q_2$  are renamed apart except for the answer variable  $x$ , a *matching*  $\theta$  is a 1-1 mapping from the variables of  $q_2$  to the variables of  $q_1$ . The *common subquery*  $q$  is formed by renaming the variables of  $q_2$  according to the matching, and keeping the conjuncts also found in  $q_1$ , i.e.  $q = q_1 \cap q_2 \theta$ .

**Example 9.** *Continuing the previous examples, we will find common subqueries of the MSQs of p3 and p4. We can match the lung finding of p4 with the lung finding of p3 as well as their sore throat findings. This would be done by renaming  $y_{21}$  to  $y_{11}$  and  $y_{22}$  to  $y_{12}$  and then keeping the common conjuncts resulting in the query  $q(x) = \{\text{Exemplar}(x), \text{hasFinding}(x, y_{11}), \text{SoreThroat}(y_{11}), \text{Finding}(y_{11}), \text{hasFinding}(x, y_{12}), \text{LungFinding}(y_{12}), \text{Finding}(y_{12})\}$ . In our approach we evaluate the quality of matchings based on the number of conjuncts in the query they induce, which in this case is 7.*

Finding the maximum common subquery belongs to a larger family of problems of finding maximal common substructures in graphs, such as the mapping problem for processors [67], structural resemblance in proteins [68] and the maximal common substructure match (MCSSM) [69]. Our problem can be expressed as a weighted version of most of these problems, since they only seek to maximize the number of nodes in the common substructure, which, in our case, corresponds to the number of variables in the resulting subquery. Since we want to maximize the number of common conjuncts, we could assign weights to variable matchings ( $y$  corresponds to  $z$ ) due to concept conjuncts, and to pairs of variable matchings ( $y$  corresponds to  $z$  and  $y'$  corresponds to  $z'$ ) due to role conjuncts. This is an instance of the general graph matching problem. These problems are NP-hard and are therefore usually solved with approximation algorithms [70, 71]. With Alg. 4, we introduce our own method for approximately solving this problem since we need to impose the additional restriction that the resulting query must be connected. Any conjuncts in the resulting query that aren't connected to the answer variable do not influence the query's certain answers, so we consider them extraneous.

Alg. 4 initializes an empty variable renaming  $\theta$ .  $V$  and  $U$  contain the variables that have not been matched yet. At each iteration of the main loop, the algorithm attempts to match one of the variables of  $q_2$  with one of the variables of  $q_1$ . Only matches that conserve the connectedness of the induced query are considered (set  $S$ ). If there are no such matches the main loop terminates. Otherwise, the match that adds the largest number of conjuncts to the induced query is selected. This match is added to  $\theta$  and the corresponding variables are removed from  $V$  and  $U$ . Finally, the query  $q$  which consists of the common conjuncts of  $q_1$  and  $q_2$  is constructed, by renaming the variables of  $q_2$  according to  $\theta$  and keeping the conjuncts that also appear in  $q_1$ .

**Algorithm 4:** GREEDYCOMMONCONJUNCTS

---

**Input:** Two queries  $q_1, q_2$ , with query graphs  $G_{q_1} = \langle V_1, E_1, \ell_{V_1}, \ell_{E_1} \rangle, G_{q_2} = \langle V_2, E_2, \ell_{V_2}, \ell_{E_2} \rangle$ .

**Output:** A query consisting of common conjuncts of  $q_1$  and  $q_2$ .

---

```

1  if  $|\text{var}(q_1)| < |\text{var}(q_2)|$  then
2  |   Swap  $q_1, q_2$ 
3  end
4   $\theta \leftarrow \{\}$ 
5   $U \leftarrow \text{var}(q_1) \setminus \{x\}$ 
6   $V \leftarrow \text{var}(q_2) \setminus \{x\}$ 
7  while true do
8  |    $S \leftarrow \{z \mapsto y \mid z \in V, y \in U, \langle z, z' \rangle \in E_1, \langle y, y' \rangle \in E_2, \ell_{E_1}(\langle z, z' \rangle) \cap \ell_{E_2}(\langle y, y' \rangle) \neq \emptyset, z' \mapsto y' \in \theta\}$ 
9  |    $\cup \{z \mapsto y \mid z \in V, y \in U, \langle z', z \rangle \in E_1, \langle y', y \rangle \in E_2, \ell_{E_1}(\langle z', z \rangle) \cap \ell_{E_2}(\langle y', y \rangle) \neq \emptyset, z' \mapsto y' \in \theta\}$ 
10 |   if  $S \neq \emptyset$  then
11 |   |    $\hat{z} \mapsto \hat{y} \leftarrow \arg \max_{z \mapsto y \in S} \{|q_1 \cap q_2(\theta \cup \{z \mapsto y\})| - |q_1 \cap q_2\theta|\}$ 
12 |   |    $\theta \leftarrow \theta \cup \{\hat{z} \mapsto \hat{y}\}$ 
13 |   |    $V \leftarrow V \setminus \{\hat{z}\}$ 
14 |   |    $U \leftarrow U \setminus \{\hat{y}\}$ 
15 |   else
16 |   |   break
17 |   end
18 end
19  $q \leftarrow q_1 \cap q_2\theta$ 
20 return  $q$ 

```

---

An efficient implementation of Alg. 4 will use a max-priority queue to select at each iteration the match that adds the largest number of conjuncts to the induced query. The priority queue should contain an element for each pair of variables that can be matched, with the priority of the element being either 0 if the pair is not in  $S$ , and otherwise equal to the number of conjuncts it would add to the query. At each iteration some priorities may need to be updated. The time complexity of priority queues varies depending on their implementation. Implementations such as Strict Fibonacci Heaps have lower time complexity but perform worse in practice than simpler implementations with higher time complexities such as Binary Heaps. A Strict Fibonacci Heap implementation would have a time complexity of  $O(n^2m^2)$ , where  $n = |\text{var}(q_1)|$ ,  $m = |\text{var}(q_2)|$ , while a Binary Heap implementation  $O(n^2m^2 \log(nm))$ .

## 5. Experiments

We evaluated the proposed algorithms and framework by generating explanations for various classifiers, computing the metrics presented in Section 3, comparing with other rule-based approaches where possible, and discussing the quality and usefulness of the results.<sup>1</sup> Specifically, we explored four use-cases: a) we are given a tabular dataset which serves as both the features of the classifier and the explanation dataset, b) we are given raw data along with curated semantic descriptions (ABox) but without a TBox, c) we are given curated semantic descriptions, along with a TBox, and d) we are only given raw data, so semantic descriptions need to be constructed automatically. The first use-case facilitated comparison of the proposed KGRules-H algorithm with other rule-based XAI approaches from literature, the second and third use-cases were more suitable for validating the usefulness of the proposed framework in its intended setting, while the fourth explored how KGRules-H could be utilized in a scenario in which semantic data descriptions are not available. For the first use-case we experimented on the Mushroom<sup>2</sup> dataset which

---

<sup>1</sup>All code and data are available at <https://github.com/ails-lab/kgrules-h>

<sup>2</sup><https://archive.ics.uci.edu/ml/datasets/mushroom>

1 involves only categorical features. For the second use-case we used the CLEVR-Hans3 dataset [72] which consists 1  
 2 of images of 3D geometric shapes in addition to rich metadata which we use as semantic data descriptions. For the 2  
 3 third we constructed an explanation dataset from a subset of the Visual Genome [8] to explain a classifier trained 3  
 4 on Places365 [73] for scene classification, while for the fourth we used MNIST [74] which contains images of 4  
 5 handwritten digits, and no metadata . 5

6 The components needed for each experiment, were: a) a black-box classifier for which we provide rule-based ex- 6  
 7 planations, b) an explanation dataset with semantic descriptions of exemplar data using the appropriate terminology 7  
 8 and c) reasoning capabilities for semantic query answering, in order to evaluate the generated rules. As classifiers 8  
 9 we chose widely used neural networks, which are provided as default models in most deep learning frameworks. 9

10 For constructing explanation datasets in practice, we identified two general approaches: a) the curated approach, 10  
 11 and b) the automated information extraction approach. For the manual curated approach, the semantic descriptions 11  
 12 of exemplar data were provided. In an ideal scenario, curated explanation datasets are created by domain experts, 12  
 13 providing semantic descriptions which are meaningful for the task, and by using the appropriate terminology, lead 13  
 14 to meaningful rules as explanations. In our experiments, we simulated such a manually curated dataset by using Visual 14  
 15 Genome, which provides semantic descriptions for each image present in the dataset, using terminology linked to 15  
 16 WordNet [10]. Of course, using human labor for the creation of explanation datasets in real-world applications would 16  
 17 be expensive, thus we also experimented with automatically generating semantic descriptions for exemplar data. 17  
 18 Specifically, for the automated information extraction approach we used domain specific, robust, feature extraction 18  
 19 techniques and then provided semantic descriptions in terms of the extracted features. In these experiments, we 19  
 20 automatically generated semantic descriptions for images in MNIST, by using ridge detection, and then describing 20  
 21 each image as a set of intersecting lines. 21

22 For acquiring certain answers of semantic queries, which requires reasoning, we set up repositories on GraphDB.<sup>3</sup> 22  
 23 For the case of the Mushrooms dataset we used the certain answers to measure fidelity, number of rules and average 23  
 24 rule length and compared our approach with RuleMatrix [26] which implements scalable bayesian rule lists [25], 24  
 25 Skope-Rules<sup>4</sup> and the closely related KGrules[46]. For CLEVR-Hans3, we mainly used the certain answers to 25  
 26 explore whether our explainability framework can detect the foreknown bias of the classifier, while for Visual 26  
 27 Genome we attempt to discover biases, by observing the best generated rule-queries with respect to precision, recall, 27  
 28 and degree as defined in Section 3. Finally, for MNIST we observed quality and usability related properties of 28  
 29 generated rule queries and their exceptions. 29  
 30

### 31 5.1. Mushrooms 31

32  
 33 The purpose of the Mushroom experiment was to compare our results with other rule-based approaches from the 33  
 34 literature. Since other approaches mostly provide explanation rules in terms of the feature space, the explanation 34  
 35 dataset was created containing only this information. We should note that this was only done for comparison's sake, 35  
 36 and is not the intended use-case for the proposed framework, in which there would exist semantic descriptions which 36  
 37 cannot be necessarily represented in tabular form, and possibly a TBox. 37  
 38

#### 39 5.1.1. Explanation Dataset 39

40 The Mushroom dataset contains 8124 rows, each with 22 categorical features of 2 to 9 possible values. We 40  
 41 randomly chose subsets of up to 4000 rows to serve as the exemplars of explanation datasets. The vocabulary 41  
 42  $\langle \text{CN}, \text{RN}, \text{IN} \rangle$  used by the explanation datasets consisted of: an individual name for each index of a row included in 42  
 43 the explanation dataset (in IN), and a concept name for each combination of categorical feature and value (in CN), 43  
 44 giving us a total of  $|\text{CN}| = 123$ . In this case, the set of exemplars (EN) coincided with the set of individual names. 44  
 45 Then, for each feature and for each row, we added an assertion to the ABox, representing the value of the feature 45  
 46 for the corresponding row. Thus, the knowledge base of an explanation dataset and its vocabulary had the form: 46  
 47

$$48 \text{IN} = \{\text{row}_1, \text{row}_2 \dots \text{row}_n\} 48$$

49  
 50 <sup>3</sup><https://graphdb.ontotext.com/> 50

51 <sup>4</sup><https://github.com/scikit-learn-contrib/Skope-Rules> 51

1 CN = {CapShapeBell, CapShapeConical, CapShapeConvex, . . . , HabitatGrass, HabitatLeaves, . . . }

2 RN =  $\emptyset$

3  $\mathcal{A} = \{\text{CapShapeFlat}(\text{row}_1), \text{CapSurfaceFibrous}(\text{row}_1), \dots, \text{HabitatPaths}(\text{row}_n)\}$

### 6 5.1.2. Setting

7 For this set of experiments, we used a simple multi-layer perceptron with one hidden layer as the black-box  
8 classifier. The classifier achieved 100% accuracy on train and test set in all experiments. After training, we generated  
9 explanations using Alg. 1 (in the case of tabular data, both merge operations outlined in section 4.3 are reduced  
10 to a simple intersection of the concept assertions of the merged individuals), along with three other rule-based  
11 approaches from related literature, and compared them.

12 We split the dataset into four parts: 1) a training set, used to train the classifier, 2) a test set, used to evaluate  
13 the classifier, 3) an explanation-train set, used to generate explanation rules, and 4) an explanation-test set, used  
14 to evaluate the generated explanation rules. When running KGrules-H and KGrules, the explanation dataset was  
15 constructed from the explanation-train set. We experimented by changing the size of this dataset, from 100 to 4000  
16 rows, and observed the effect it had on the explanation rules. On the explanation-test set, we measured the fidelity  
17 of the generated rules which is defined as the proportion of items for which classifier and explainer agree. We also  
18 measured the number of generated rules and average rule length. We used the proposed KGrules-H algorithm to  
19 generate explanations, and we also generated rules (on the same data and classifier) with RuleMatrix, Skope-rules  
20 and the closely related KGrules.

21 To compare with the other methods, which return a set of rules at their output, in this experiment we only consid-  
22 ered correct rules we generated. To choose which rules to consider (what would be shown to a user), from the set  
23 of all correct rule-queries generated, we greedily chose queries starting with the one that had the highest count of  
24 certain answers on the explanation-train set, and then iteratively adding queries that provided the highest count of  
25 certain answers, not provided by any of the previously chosen queries. This is not necessarily the optimal strategy  
26 of rule selection for showing to a user (it never considers rules with exceptions), and we plan to explore alternatives  
27 in future work.

28 Finally, for all methods except for the related KGrules we measured running-time using same runtime on Google  
29 Colab<sup>5</sup>, by using the “%timeit”<sup>6</sup> magic command with default parameters. KGrules was not compatible with  
30 this benchmarking test, since it is implemented in Java as opposed to Python which is the implementation language  
31 provided for the other methods. In addition, KGrules implements an exhaustive exponential algorithm, so it is  
32 expected to be much slower than all other methods.

### 34 5.1.3. Results

35 The results of the comparative study are shown in Table 4. A first observation is that for small explanation datasets,  
36 the proposed approach did not perform as well as the other methods regarding fidelity, while for large ones it even  
37 outperformed them. This could be because for small explanation datasets, when the exemplars are chosen randomly,  
38 there are not enough individuals, and variety in the MSQs of these individuals, for the algorithm to generalize by  
39 merging their MSQs. This is hinted also from the average rule length, which is longer for both KGrules-H and  
40 KGrules for explanation dataset sizes 100 and 200, which indicates less general queries. Conversely, for explanation  
41 datasets of 600 exemplars and larger, the proposed approach performs similarly, in terms of fidelity, with related  
42 methods. Regarding running time, KGrules-H is the fastest except for the case of the largest explanation dataset, for  
43 which Skope-rules is faster. However, in this case Skope-rules’s performance suffers with respect to fidelity, whereas  
44 RuleMatrix and KGrules-H achieve perfect results. Furthermore, our proposed approach seems to generate longer  
45 rules than all other methods which on the one hand means that they are more informative, though on the other hand  
46 they are potentially less understandable by a user. This highlights the disadvantages constructing queries using the  
47 MSQs as a starting point. This was validated upon closer investigation, where we saw that the rules generated from  
48 the small explanation datasets were more specific than needed, as the low value of fidelity was due only to false  
49

50 <sup>5</sup><https://colab.research.google.com/>

51 <sup>6</sup><https://ipython.readthedocs.io/en/stable/interactive/magics.html#magic-timeit>



negatives, and there were no false positives. Detecting redundant conjuncts might be done more efficiently for rules concerning tabular data, but for general rules, on which our approach is based, this task is very computationally demanding.

Size	Method	Fidelity	Nr. of Rules	Avg. Length	Time (seconds)
100	KGrules-H	92.70%	4	14	0.15
	KGrules	<b>97.56%</b>	11	5	-
	RuleMatrix	94.53%	3	2	3.65
	Skope-rules	97.01%	3	2	1.29
200	KGrules-H	93.40%	8	14	0.18
	KGrules	98.37%	11	5	-
	RuleMatrix	97.78%	4	2	3.65
	Skope-rules	<b>98.49%</b>	4	2	1.47
600	KGrules-H	<b>99.60%</b>	10	13.3	0.61
	KGrules	99.41%	13	4	-
	RuleMatrix	99.43%	6	1	7.69
	Skope-rules	98.52%	4	2	1.58
1000	KGrules-H	<b>100%</b>	9	14.2	1.67
	KGrules	99.58%	14	6.57	-
	RuleMatrix	99.90%	6	1.33	11.1
	Skope-rules	98.50%	4	2.25	2.95
4000	KGrules-H	<b>100%</b>	10	14.2	34.30
	KGrules	99.72%	16	5.81	-
	RuleMatrix	<b>100%</b>	7	1.43	43.00
	Skope-rules	96.85%	2	2	3.01

Table 4

Performance on the Mushroom dataset.

## 5.2. CLEVR-Hans3

The second set of experiments involved CLEVR-Hans3 [72], which is an image classification dataset designed to evaluate algorithms that detect and fix biases of classifiers. This dataset provides sufficient and reliable information to create an explanation dataset, while the given train-test split contains intentional biases, which was ideal as a grounds for experimentation, as we could observe the extent to which the proposed explanation rules can detect them.

### 5.2.1. Explanation Dataset

We created an explanation dataset  $\mathcal{E}$  using the test set of CLEVR-Hans3, consisting of 750 images for each class. By utilizing the descriptions of images provided as metadata, we defined a vocabulary  $\langle \text{CN}, \text{RN}, \text{IN} \rangle$ , in which there was an individual name for each image and for each depicted object (in IN), a concept name for each size, color, shape and material in addition to three concept names (Class1, Class2, Class3) corresponding to the set of classes  $\mathcal{C}$  and two indicative concept names Image and Object, as concept names (in CN), and the role name (in RN) contains which was used to link images to objects it depicts. We used the (unique) names of the image files of the test set as names for the exemplars (EN) of our explanation dataset, so the mapping  $\mathcal{M}$  was straightforward (mapping the file name to the respective image). We then created the knowledge base  $\mathcal{S}$  over this vocabulary, with the ABox containing the semantic description of all images and the respective objects, and the TBox containing certain rather




Fig. 7. Image samples of the three classes of CLEVR-Hans3 along with the class rules and the confounding factors in parentheses.

trivial inclusion axioms. The sets IN, CN, RN and the TBox of our knowledge base and the respective vocabulary were the following:

$$\text{IN} = \{\text{image1.png}, \text{object1}_1, \text{object2}_1, \dots, \text{objectN}_1, \text{image2.png}, \text{object1}_2, \text{object2}_2, \dots, \text{image2250.png}, \text{object1}_{2250}, \dots, \text{objectM}_{2250}\}$$

$$\text{CN} = \{\text{Image}, \text{Object}, \text{Cube}, \text{Cylinder}, \text{Sphere}, \text{Metal}, \text{Rubber}, \text{Blue}, \text{Brown}, \text{Cyan}, \text{Gray}, \text{Green}, \text{Purple}, \text{Red}, \text{Yellow}, \text{Large}, \text{Small}, \text{Class1}, \text{Class2}, \text{Class3}\}$$

$$\text{RN} = \{\text{contains}\}$$

$$\mathcal{T} = \{C \sqsubseteq \text{Object} \mid C \notin \{\text{Image}, \text{Object}, \text{Class1}, \text{Class2}, \text{Class3}\}\}.$$

### 5.2.2. Setting

For the experiments on CLEVR-Hans3 we used the same ResNet34 [75] classifier and training procedure as those used by the creators of the dataset in [72]. The performance of the classifier is shown in Table 5 and so is the confusion matrix which summarizes the predictions of the classifier indicating the differences between the actual and the predicted classes of the test samples. As expected, the classifier has lower values on some metrics regarding the first two classes, and this is attributed to the confounding factors and not the quality of the classifier, since it achieved 99.4% accuracy in the validation set. After training the classifier, we acquired its predictions on the test set and generated explanations for each class with Alg. 1 with QLCS as the merge operation. Afterwards, utilizing the reasoning capabilities of GraphDB, we loaded the knowledge base  $\mathcal{S}$  of our explanation dataset, and obtained the certain answers of the corresponding explanation rule queries in order to evaluate the produced explanation rules. The evaluation was based on the metrics mentioned in Section 3, essentially comparing the certain answers of the explanation rule queries, with the corresponding exemplars of the pos-set of the classifier.

### 5.2.3. Results

The explanation rules generated for the ResNet34 classifier using KGrules-H and QLCS as the merge operation, as outlined in section 4.3.1, are shown in Table 6, where we show the rule, the value of each metric and the numbers of positive individuals. The term positive individuals refers to the certain answers of the respective explanation rule query that are also elements of the pos-set (they are classified in the respective class).

In our representation of explanation rule queries in Tables 6,7 we have omitted the answer variable  $x$ , along with all conjuncts of the form  $x$  contains  $y$  and conjuncts of the form  $\text{Object}(y)$ , for brevity. In addition, the rules in the Table are not formally written, to make more visually clear the characteristics of the objects involved. For example, the rule of the first row (Best precision for class 1) would formally be written  $\text{Exemplar}(x), \text{contains}(x, y_1), \text{contains}(x, y_2), \text{contains}(x, y_3), \text{Large}(y_1), \text{Cube}(y_1), \text{Gray}(y_1), \text{Large}(y_2), \text{Cylinder}(y_2), \text{Large}(y_3), \text{Metal}(y_3) \rightarrow \text{Class1}(x)$ .

The algorithm found a correct rule (precision = 1) for each class, in addition to a rule query with recall = 1, whose certain answers are a superset of the positive set. The best degree was achieved for class 3, which lacks a confounding factor, meaning the classifier is not expected to be biased. Correct rule queries are of particular interest since they can be translated into guaranteed IF-THEN rules which the classifier follows on the particular dataset. For instance the highest recall correct rule query for class 1 is translated into the rule "If the image contains a Large Gray Cube, a Large Cylinder and a Large Metal Object then it is classified to class 1.". This rule clearly shows the bias of the classifier, since it is the description of the class with the added confounding factor (the Large Cube is Gray). Similarly the (not correct) rule query with recall = 1 for the same class can be translated into the rule "If the image does not contain a Large Cube then it is not classified to class 1", since the set of certain answers is a super set of the positive set. We observed that correct rule queries tend to be more specific than others, with the most general rules with exceptions being those with recall = 1. Other rules which were correct with exceptions, tended to lie somewhere in the middle with respect to how general or specific they are, but they were the ones which lead to the highest values of degree. By observing these results, we concluded that in practice, a set of rules, both correct and with exceptions, can give us a very clear picture of what the black-box classifier is doing. However, in order to not overwhelm an end-user with a large number of rules, we should develop a strategy to select which rules to show to the user. Here, as opposed to the Mushroom experiment (5.1) the strategy we used was to show the highest recall, highest precision and highest degree rules, along with their exceptions if any, but as mentioned, we plan to explore additional strategies in the future, such as showing disjunctions of correct rules.

It is interesting to note that the rule query with recall = 1 produced for class 1 contained a Large Cube but not a Large Cylinder, which is also in the description of the class. This shows that in the training process the classifier learned to pay more attention to the presence of cubes rather than the presence of cylinders. The elements of the highest recall correct rule that differ from the true description of class 1 can be a great starting point for a closer inspection of the classifier. We expected the presence of a Gray Cube from the confounding factor introduced in the training and validation sets, but in a real world scenario similar insights can be reached by inspecting the queries. In our case, we further inquired the role that the Gray Cube and the Large Metal Object play in the correct rule by removing either of them from the query and examining its behavior. In Table 7 we can see that the gray color was essential for the correct rule while the Large Metal Object was not, and in fact its removal improved the rule and returned almost the entire class.

Another result that piqued our attention was the highest degree explanation for class 3 which is the actual rule that describes this class. This explanation was not a correct rule, since it had two exceptions, which we can also see in the confusion matrix of the classifier and we were interested to examine what sets these two individuals apart. We found that both of these individuals are answers to the query "y1 is Large, Gray, Cube". This showed us once again the great effect the confounding factor of class 1 had on the classifier.

Our overall results show that the classifier tended to emphasize low level information such as color and shape and ignored higher level information such as texture and the combined presence of multiple objects. This was the

Table 5  
Performance of the ResNet34 model on CLEVR-Hans3.

True label	Test set metrics			Confusion matrix		
	Precision	Recall	F1-score	Class 1	Class 2	Class 3
Class 1	0.94	0.16	0.27	118	511	121
Class 2	0.59	0.98	0.54	5	736	9
Class 3	0.85	1.00	0.92	2	0	748

Table 6  
Optimal explanations with regard to the three metrics on CLEVR-Hans3.

Metric	Explanation Rules	Precision	Recall	Degree	Positives
Class 1					
Best Precision	y1 is Large, Cube, Gray. y2 is Large, Cylinder. y3 is Large, Metal.	1.00	0.66	0.66	83
Best Recall	y1 is Large, Cube.	0.09	1.00	0.09	125
Best Degree	y1 is Large, Cube, Gray. y2 is Large, Cylinder. y3 is Large, Metal.	1.00	0.66	0.66	83
Class 2					
Best Precision	y1 is Small, Sphere. y2 is Large, Rubber. y3 is Small, Metal, Cube. y4 is Small, Brown. y5 is Small, Rubber, Cylinder.	1.00	0.09	0.09	116
Best Recall	y1 is Cube.	0.63	1.00	0.63	1247
Best Degree	y1 is Metal, Cube. y2 is Small, Metal.	0.78	0.8	0.65	1005
Class 3					
Best Precision	y1 is Metal, Blue. y2 is Large, Blue, Sphere. y3 is Yellow, Small, Sphere. y4 is Small, Rubber. y5 is Metal, Sphere.	1.00	0.42	0.42	365
Best Recall	y1 is Large. y2 is Sphere.	0.42	1.00	0.42	878
Best Degree	y1 is Yellow, Small, Sphere. y2 is Large, Blue, Sphere.	0.99	0.85	0.85	748

reason why the confounding factor of class 1 had an important effect to the way images were classified, while the confounding factor of class 2 seemed to have had a much smaller one. Furthermore, the added bias made the classifier reject class 1 images, which however had to be classified to one of the other two classes (no class was not an option). Therefore one of the other classes had to be “polluted” by samples which were not confidently classified to a class. This motivates us to expand the framework in the future to work with more informative sets than the pos-set, such as elements which were classified with high confidence, and false and true, negatives and positives

### 5.3. Visual Genome and Places

For the third set of experiments, we produced explanations of an image classifier trained on the Places365 dataset for scene classification. To do this, we constructed an explanation dataset from a subset of the visual genome, which includes a *scene-graph* for each image. Nodes of scene-graphs represent depicted objects, and edges represent

Table 7

Two modified versions of the class 1 correct rule produced by removing conjuncts.

Query	Positives	Negatives
y1 is Large, Cube. y2 is Large, Cylinder. y3 is Large, Metal.	108	547
y1 is Large, Cube, Gray. y2 is Large, Cylinder.	93	0

relationships between them, while both edges and nodes are labeled with WordNet synsets. This allows us to encode the provided scene-graphs as an ABox, and to use the hyponym-hypernym hierarchy of WordNet as a TBox.

### 5.3.1. Explanation Dataset

To construct the explanation dataset, we first selected the two most confused classes based on the confusion matrix of the classifier on the Places365 test set. These were “Desert Sand” and “Desert Road”. Then, we acquired the predictions of the classifier on the entirety of the visual genome dataset, and kept as exemplars the images for which the top prediction was one of the aforementioned classes. There were 273 such images. We defined a vocabulary  $\langle \text{CN}, \text{RN}, \text{IN} \rangle$ , where the set of individual names IN contained a unique name for each image and for each depicted object, based on their visual genome IDs, while the mapping function was a dictionary linking image IDs (EN) to the corresponding image files. The size of IN was 6462. We initialized the set of concept names CN with two concepts: Image and Object and then added to CN all WordNet synsets that appear as node labels, in addition to all of their hypernyms. The size of CN was 1155. Similarly, the set of role names (RN) was initialized with the role depicts, which links images to their depicted objects, and then extended with the set of WordNet synsets that appear as edge labels in the selected images’ scene graphs, in addition to all of their hypernyms, totaling 226 role names. Using this vocabulary we constructed a knowledge base  $\langle \mathcal{T}, \mathcal{A} \rangle$ . The ABox  $\mathcal{A}$  was initialized with assertions of the form  $\text{Image}(\text{image})$ ,  $\text{depicts}(\text{image}, \text{object})$ , encoding information about which images depict which objects. Then, based on the labels of the nodes of the scene-graphs of images we introduced concept assertions of the form  $\text{SomeWordNetConcept}(\text{object})$  and based on the labels of the edges of scene-graphs we introduced role assertions of the form  $\text{SomeWordNetRole}(\text{objectA}, \text{objectB})$ . Finally, the TBox  $\mathcal{T}$  contained hierarchies of concepts and roles, based on the hypernym-hyponym relation from WordNet, encoded as  $\text{hyponym} \sqsubseteq \text{hypernym}$ .

$$\text{IN} = \{\text{image1.png}, \text{object1}_1, \text{object2}_1, \dots, \text{objectN}_1, \text{image2.png}, \text{object1}_2, \text{object2}_2, \dots, \text{image2250.png}, \text{object1}_{2250}, \dots, \text{objectM}_{2250}\}$$

$$\text{CN} = \{\text{Image}, \text{Object}, \text{motorcycle.n.01}, \text{car.n.01}, \text{motor\_vehicle.n.01}, \dots\}$$

$$\text{RN} = \{\text{depicts}, \text{drink.v.01}, \text{eat.v.02}, \text{consume.v.02} \dots\}$$

$$\mathcal{T} = \{\text{motorcycle.n.01} \sqsubseteq \text{motor\_vehicle.n.01}, \text{car.n.01} \sqsubseteq \text{motor\_vehicle.n.01} \dots, \text{drink.v.02} \sqsubseteq \text{consume.v.02}, \text{eat.v.02} \sqsubseteq \text{consume.v.02} \dots\}.$$

### 5.3.2. Setting

As a black-box classifier in this set of experiments we used the *ResNet50* classifier<sup>7</sup> provided by the official GitHub repository<sup>8</sup> for models pretrained on Places365, which classifies images to 365 different classes.<sup>9</sup> The top1 error is 44.82% and the top5 error is 14.71%. As mentioned, we used the confusion matrix to select the two most confused classes to generate explanations for, which were “Desert Sand” and “Desert Road”. As a merge operation for Alg. 1 we used GreedyCommonConjuncts (Alg. 4), as the QLCS produced queries which could not be sufficiently minimized. Similarly with before, we loaded the knowledge of the explanation dataset in GraphDB, and after acquiring the certain answers of queries generated by Alg. 1, we validated their quality with respect to the

<sup>7</sup>PyTorch model: [http://places2.csail.mit.edu/models\\_places365/resnet50\\_places365.pth.tar](http://places2.csail.mit.edu/models_places365/resnet50_places365.pth.tar)

<sup>8</sup><https://github.com/CSAILVision/places365>

<sup>9</sup>[https://github.com/zhoubolei/places\\_devkit/blob/master/categories\\_places365.txt](https://github.com/zhoubolei/places_devkit/blob/master/categories_places365.txt)

precision, recall and degree metrics. After observing the best generated queries, we sought to further verify their accuracy for describing the black-box classifier, by looking into the black box to find out which parts of the image contributed to the prediction. We used a technique to generate heatmaps of pixel importance called *Score-CAM*[76]. *Score-CAM* is a white-box method for producing heatmaps that works by selectively masking areas of the input image and determining their importance to the prediction of the classifier.

### 5.3.3. Results

The best rule queries for each metric for each of the two classes is shown in Table 8. For simplicity of presentation, we have omitted conjuncts of the form *Exemplar(x)*, *Image(x)*, *depicts(x,y)* and *Object(y)*, where  $x$  is the answer variable. For both classes the generated queries have some unexpected conjuncts despite decent performance with respect to the three metrics. For example the conjunct *giraffe.n.01(y3)* appearing in the best correct rule for “Desert Road”, and the best degree rule for the same class being simply *Image(x)*, *depicts(x,y1)*, *animal.n.01(y1)*. Furthermore, for “Desert Sand”, the concept *road.n.01* stands out. The concept *communication.n.02* has the hyponyms *sign.n.02* and *written\_communication.n.01* in the WordNet hierarchy which could refer to license plates and traffic signs. The second best query in terms of precision was *Image(x)*, *depicts(x,y1)*, *depicts(x,y2)*, *depicts(x,y3)*. *motor\_vehicle.n.01(y1)*, *sky.n.01(y2)*, *road.n.01(y3)*, *along.r.01(y1,y3)*. Again the concept *motor\_vehicle.n.01* stands out.

In Fig. 8 we present the saliency maps of three images from Visual Genome classified as “Desert Road” on the left and three classified as “Desert Sand” on the right. The saliency maps support the associations the queries highlighted, namely that images containing animals, especially giraffes, get classified as “Desert Road”, while images with roads and vehicles get classified as “Desert Sand”. Not all images exhibited such strong association. In Fig. 9 the classifier seems to be mostly influenced by the ground which is covered in dry grass. Given that the second highest value in the confusion matrix of this classifier was for the pair “Desert Sand”, “Desert Vegetation” we conjecture that some mistakes were made during the training of this classifier or several images in the training set of the Places365 dataset are mislabeled. The classifier may have been fed images that should be described as “Desert Road” but with the target label being “Desert Sand” and images that should be described as “Desert Vegetation” but with the target label being “Desert Road”. This would explain the weird associations exhibited and the conjuncts appearing in the explanations. The association of giraffes with the label “Desert Road” can be explained with mislabeled images depicting desert vegetation. This is still a wrongful association since giraffes can also populate grasslands such as in the top left image of Fig. 8.

We would like to point out that our method of producing explanations discovered these erroneous associations without asking from the end-user to manually inspect the heatmaps of a large number of images, rather than to read the best performing query for each one of our metrics. We believe that this provides a better user experience since it requires less effort for the initial discovery of suspect associations, which can then be thoroughly tested by examining other explanations such as heatmaps. Our method is also model-agnostic, while methods such as *Score-CAM* require access to the layers within a convolutional neural network.

## 5.4. MNIST

For the fourth and final set of experiments we used MNIST, which is a dataset containing grayscale images of handwritten digits [74]. It is a very popular dataset for machine learning research, and even though classification on MNIST is essentially a solved problem, many recent explainability approaches experiment on this dataset (for example [77]). For us, MNIST was ideal for experimenting with automatic explanation dataset generation by using traditional feature extraction from computer vision. An extension to this approach would be using more complex information extraction from images, such as object detection or scene graph generation, for applying the explainability framework to explain generic image classifiers. This however is left for future work.

### 5.4.1. Explanation Dataset

For creating the explanation dataset for MNIST, we manually selected a combination of 250 images from the test set, including both typical and unusual exemplars for each digit. The unusual exemplars were chosen following the mushroom experiment (5.1), in which we saw that small explanation datasets do not facilitate good explanation rules

Table 8

Optimal explanations with regard to the three metrics using the VG explanation dataset

Metric	Explanation Rules	Precision	Recall	Degree	Positives
Desert Road					
Best Precision	y1 is field.n.01.	1.00	0.12	0.12	16
	y2 is natural_object.n.01.				
	y3 is giraffe.n.01.				
	y4 is body_part.n.01.				
	y5 is woody_plant.n.01.				
Best Recall	y1 is organism.n.01	0.54	1.00	0.54	139
Best Degree	y1 is animal.n.01	0.72	0.84	0.64	118
Desert Sand					
Best Precision	y1 is instrumentality.n.03	1.00	0.12	0.12	16
	y2 is road.n.01				
	y3 is communication.n.02				
	y4 is sky.n.01				
	y5 is tree.n.01				
Best Recall	y1 is physical_entity.n.01	0.49	1.00	0.49	134
Best Degree	y1 is instrumentality.n.01	0.92	0.56	0.56	76
	y2 is road.n.01				

when the exemplars are chosen randomly, so we aimed for variety of semantic descriptions. In addition, the unusual exemplars tended to be misclassified, and we wanted to see how their presence would impact the explanations.

Since there was no semantic information available that could be used to construct an explanation dataset, we automatically extracted descriptions of the images, by using feature extraction methods. Specifically, the images were described as a collection of intersecting lines, varying in angle, length and location within the image. These lines were detected using the technique of ridge detection [78]. The angles of the lines were quantized to 0, 45, 90 or 135 degrees, and the images were split into 3 horizontal (top, middle, bottom) and 3 vertical (left, center, right) zones which define 9 areas (top left, top center, . . . , bottom right). For each line we noted the areas it passes through. In Fig. 10 we show an example of an MNIST image, along with the results of the aforementioned information extraction procedure using ridge detection.

Based on the selected images and the extracted information, we created our explanation dataset  $\mathcal{E}$ . We constructed a vocabulary  $\langle \text{CN}, \text{RN}, \text{IN} \rangle$ , with an individual for each image and each line therein as individual names (IN), the concepts defining the angle, location and length of each line, two indicative concepts Image and Line, as well as ten concepts (one for each digit) corresponding to the set of classes  $\mathcal{C}$ , as concept names (CN), and the roles contains with domain Image, and range Line, indicating the existence of a line in a specific image, and the symmetric intersects with both domain and range Line, indicating that two lines intersect each other, as the role names (RN). To define the mapping  $\mathcal{M}$  of the explanation dataset, after extracting the images and separating them into digits, we numbered them and named the Exemplars (and the corresponding file names) according to their digit and index. This also makes the images easy to retrieve. We then created the knowledge base  $\mathcal{S}$  over this vocabulary, with the ABox containing the semantic description of all exemplar images and the respective lines, and the TBox containing certain rather trivial inclusion axioms. The vocabulary IN, CN, RN and the Tbox of our knowledge base are the following:

$$\text{IN} = \{\text{test\_zero1}, \text{test\_zero1\_line0}, \dots, \text{test\_zero1\_line7}, \text{test\_zero6}, \text{test\_zero6\_line0}, \dots, \text{test\_nine979},$$

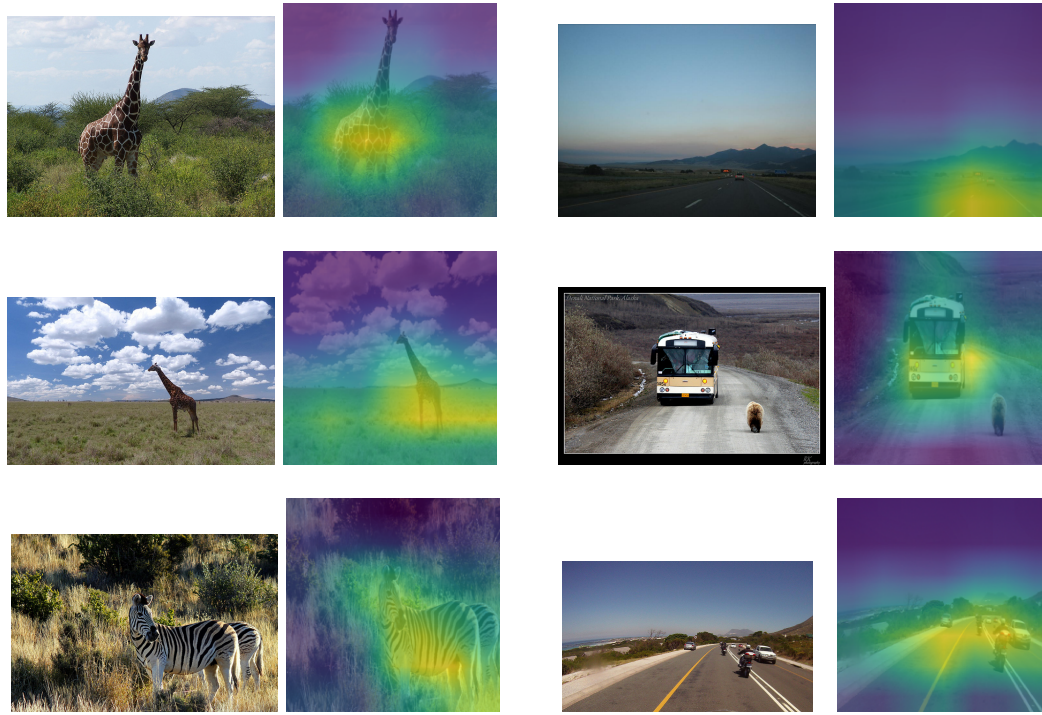


Fig. 8. Images from the Visual Genome dataset and their saliency maps produced by Score-CAM. Images on the left were classified as “Desert Road” and images on the right as “Desert Sand”.



Fig. 9. An image classified as “Desert Road” exhibiting weaker association to the presence of a giraffe.

$\text{test\_nine979\_line7, \dots, lineM}_{250}$

$\text{CN} = \{\text{Image, Line, Line0deg, Line45deg, Line90deg, Line135deg, TopLeft, TopCenter, TopRight, MidLeft, MidCenter, MidRight, BotLeft, BotCenter, BotRight, Short, Medium, Long}\}$

$\text{RN} = \{\text{contains, intersects}\}$ .

$\mathcal{T} = \{C \sqsubseteq \text{Line} \mid C \notin \{\text{Image, Line}\}\}$ .

#### 5.4.2. Setting

For MNIST we used the example neural network provided by PyTorch<sup>10</sup> as the classifier to be explained. The classifier achieved 99.8% accuracy on the training set and 99.2% on the test set. On the explanation dataset, the accuracy is 73%, and a confusion matrix for the classifier on the explanation dataset is shown in Table 11. The per-

<sup>10</sup><https://github.com/pytorch/examples/tree/master/mnist>



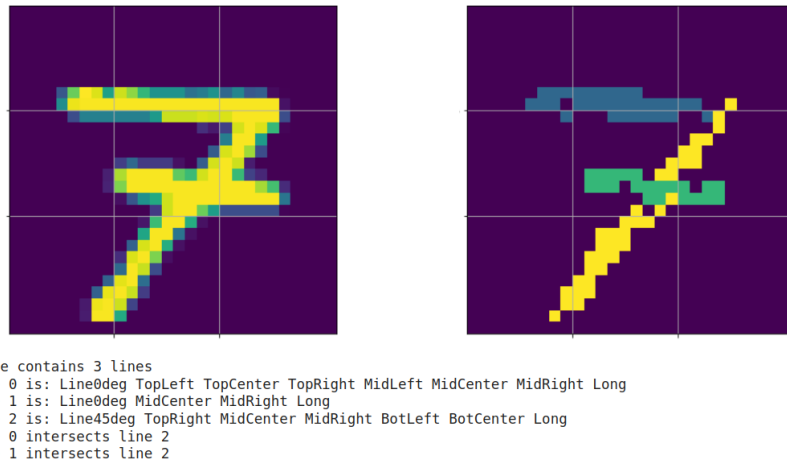


Fig. 10. An example of a digit, the results of ridge detection, and the corresponding description.

Confusion Matrix on the Explanation Dataset




Fig. 11. Confusion matrix of the classifier on the explanation dataset for MNIST

formance of the classifier is poor when compared to the whole test set, due to the fact that we have included several manually selected exemplars, which are unusually drawn, but still valid as digits (based on our own judgement). Similarly to CLEVR-Hans3, we generated explanations for the predictions of the classifier on the exemplars using KGrules-HT (Alg. 3) and KGrules-H (Alg. 1) with Alg. 4 as the merge operation, and loaded the explanation dataset in GraphDB for acquiring certain answers. We also experimented with the QLCS as the merge operation, but the resulting queries mostly contained a large number of variables which could not be effectively minimized, which could be due to the complex connections between variables with a symmetric role.

### 5.4.3. Results

For MNIST there does not exist a ground truth semantic description for each class, as was the case for CLEVR-Hans3, nor is there a pre-determined bias of the classifier, thus we could not easily measure our framework’s useful-

ness in this regard. Instead, since the explanation dataset was constructed automatically, we explored quality related features of the generated explanations.

For all digits the algorithm produced at least one correct rule (precision = 1) and a rule with exceptions with recall = 1. The highest degree of rule queries for each digit are shown as a bar-plot in Fig. 12a. In general, the values of the metric seem low, with the exception of digit 0, which would indicate that the algorithms did not find a single rule which approximates the pos-set to a high degree. For some of the digits, including 0, the highest degree rule is also a correct rule. For closer inspection, we show the best degree rule query for digit 0 which is the highest, and for digit 5 which is the lowest.




Fig. 12. Metrics of generated rule queries for MNIST

The explanation rule for digit 0 involved six lines, as indicated by the conjuncts  $\text{contains}(x, y_1)$ ,  $\text{contains}(x, y_2)$ ,  $\text{contains}(x, y_3)$ ,  $\text{contains}(x, y_4)$ ,  $\text{contains}(x, y_5)$ ,  $\text{contains}(x, y_6)$ . For five of the six lines, the explanation rule query included their location in the image, indicated by the conjuncts  $\text{TopCenter}(y_1)$ ,  $\text{BotRight}(y_2)$ ,  $\text{BotCenter}(y_2)$ ,  $\text{MidRight}(y_3)$ ,  $\text{TopRight}(y_5)$ ,  $\text{BotCenter}(y_6)$ . For all six lines the explanation rule included information about their orientation, indicated by the conjuncts  $\text{Line45deg}(y_1)$ ,  $\text{Line45deg}(y_2)$ ,  $\text{Line90deg}(y_3)$ ,  $\text{Line90deg}(y_4)$ ,  $\text{Line135deg}(y_5)$ ,  $\text{Line135deg}(y_6)$ . Finally, the rule-query included the following conjuncts which show which lines intersect each other  $\text{intersects}(y_1, y_4)$ ,  $\text{intersects}(y_2, y_3)$ ,  $\text{intersects}(y_3, y_5)$ ,  $\text{intersects}(y_4, y_6)$ . A rule query with so many conjuncts could potentially be difficult for a user to decipher, so in this case we found it useful to visualize the rules. The above rule is visualized in Fig. 13a. The visualization is clear and intuitive as an explanation for digits classified as zeroes, however visualization of rules will not be possible in all applications. This shows the importance of taking under consideration understandability when designing explanation pipelines, which in our case depends mostly on the vocabulary and expressivity of the underlying explanation dataset. In this case, the vocabulary used was itself somewhat obscure for users (sets of intersecting lines are not easy to understand by reading a rule), which could have been mitigated if the explanation dataset had been curated by humans and not created automatically. In this particular use-case it was not a problem however, since the visualization of rules was easy in most cases.

The highest degree explanation rule for digit 5, which was the lowest out of the best of all digits, again involved six lines indicated by the conjuncts  $\text{contains}(x, y_1)$ ,  $\text{contains}(x, y_2)$ ,  $\text{contains}(x, y_3)$ ,  $\text{contains}(x, y_4)$ ,  $\text{contains}(x, y_5)$ ,  $\text{contains}(x, y_6)$ . This time however, only three lines had information about their location in the image, indicated by the conjuncts  $\text{BotCenter}(y_2)$ ,  $\text{BotCenter}(y_4)$ ,  $\text{MidCenter}(y_5)$ , and five lines had information about their orientation  $\text{Line0deg}(y_1)$ ,  $\text{Line0deg}(y_2)$ ,  $\text{Line45deg}(y_3)$ ,  $\text{Line45deg}(y_4)$ ,  $\text{Line135deg}(y_6)$ . Furthermore, this rule-query included information about the size of two lines, indicated by the conjuncts  $\text{Medium}(y_4)$ ,  $\text{Short}(y_6)$ . Finally, as with before, we get a set of conjuncts showing which lines intersect each other:  $\text{intersects}(y_1, y_3)$ ,  $\text{intersects}(y_2, y_4)$ ,  $\text{intersects}(y_3, y_5)$ . This rule query is not easy to understand and it is even difficult to visualize, since there is not enough information about the location of each line, thus it is not actually usable. This was expected to an extent, due to the low value of the degree metric, but again highlights the importance of taking usability under consideration when choosing which rule-queries to show to a user.

Regarding correct rules, the algorithm produced several for each digit. Since the sets of certain answers of correct rule queries are subsets of the pos-set of each class, we measured the per class fidelity of the disjunction of all

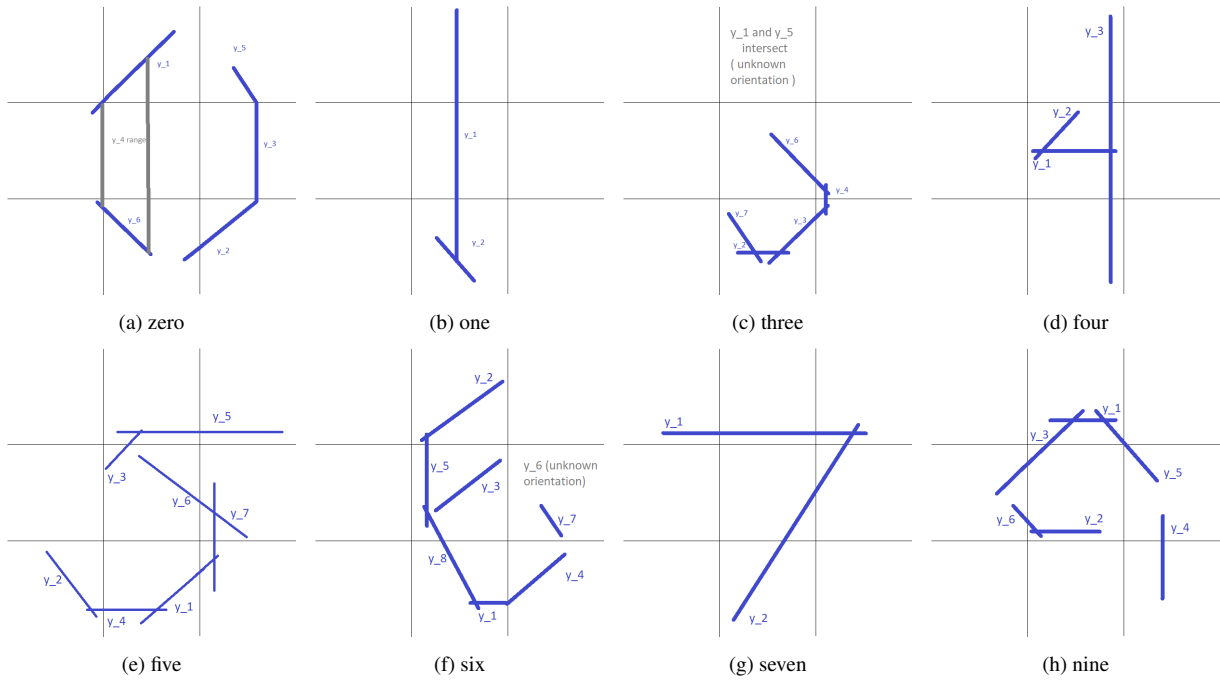


Fig. 13. Visualizations of best recall correct rules for digits




Fig. 14. Misclassified digits that follow the best recall correct rules.

correct rules, as if giving a user a rule-set, similarly to the Mushroom experiment (5.1). In Fig. 12c we show as a bar-plot the fidelity for each class. With the exception of digit one, the pos-sets of all digits were sufficiently covered by the set of correct rules. The failure for digit 1 was expected, since the descriptions of the exemplars classified as 1 contain few lines (for example consisting of a single large line in the middle) which tend to be part of descriptions of other digits as well (all digits could be drawn in a way in which there is a single line in the middle). This is a drawback of the open world assumption of DLs since we cannot guarantee the non-existence of lines that are not provided in the descriptions. The open world assumption is still desirable since it allows for incomplete descriptions of exemplars. In cases such as the medical motivating example used throughout this paper, a missing finding such as “Dyspnoea” does not always imply that the patient does not suffer from dyspnoea. It could also be a symptom that has not been detected or has been overlooked.

The highest recall of a single correct rule for each digit is shown as a bar-plot in Fig. 12b. Since correct rules are easily translated into IF-THEN rules, we expected them to be more informative than the highest degree ones, which requires looking at the exceptions to gain a clearer understanding of the rule. We investigate closer by analyzing the best correct rule for each digit.

1 For digit 0, the best correct rule rule was the same as the highest degree rule presented previously. In Fig. 14a 1  
 2 we provide an example of a six misclassified as a 0, which follows this correct rule. Comparing the misclassified 6 2  
 3 with the visualizations for rules of the digits 0 (Fig. 13a) and 6 (Fig. 13f) we can see that this 6 might have been 3  
 4 misclassified as a 0 because the closed loop part of the digit reaches the top of the image. According to the correct 4  
 5 rule for 0, an image that contains two vertical semicircles in the left and right sides of the image is classified as a 0, 5  
 6 and because of this peculiarity in the drawing of the misclassified six, the image (Fig. 13a) obeys this rule. 6

7 The best correct rule for digit 1 had the lowest recall out of all correct rules for other digits, which means that it 7  
 8 returned a small subset of the positives. Specifically, this rule returned only two of the 30 individuals classified as the 8  
 9 digit 1. It was still usable as an explanation however, as it only involved two lines contains( $x, y_1$ ), contains( $x, y_2$ ), 9  
 10 both of which were thoroughly described with regard to their location BotCenter( $y_1$ ), MidCenter( $y_1$ ), TopCenter( $y_1$ ), 10  
 11 BotCenter( $y_2$ ), their orientation Line90deg( $y_1$ ), Line135deg( $y_2$ ), their length Long( $y_1$ ), Short( $y_2$ ) and the fact that 11  
 12 they intersect intersects( $y_1, y_2$ ). This rule is visualized in Fig. 13b. 12

13 For digit 2, the best correct rule query returned nine out of the 25 positives, three of which were misclassified 13  
 14 by the classifier. This rule involved three lines, of which two had conjuncts indicating their location BotRight( $y_1$ ), 14  
 15 BotLeft( $y_3$ ), BotCenter( $y_3$ ), MidCenter( $y_3$ ), only one line contained information about its orientation and size 15  
 16 Line45deg( $y_3$ ), Long( $y_3$ ), and the only other information was that intersects( $y_2, y_3$ ). Note that for  $y_2$  the only avail- 16  
 17 able information was that it intersects  $y_3$ . This query is difficult to visualize, due to the missing information about 17  
 18 two of the three lines, however it is still useful as an explanation. Specifically,  $y_3$  represents a long diagonal line 18  
 19 from the bottom left to the middle of the image, which is a characteristic of only the digits 2 and 7. Additionally, 19  
 20 there is a line of any orientation in the bottom right of the image, which would differentiate it from a typical 7, and 20  
 21 another line which intersects the long diagonal at any position. As is apparent also in the confusion matrix of the 21  
 22 classifier on the explanation dataset (Table 11), the black-box often mixed up sevens with twos, and this explanation 22  
 23 rule returns one of the sevens which is misclassified as a two, shown in Fig. 14d. This digit is not typically drawn, 23  
 24 and from the rule query the information we get is that it might have been misclassified because of the existence of a 24  
 25 line at the bottom right. 25

26 To investigate closer, the next correct rule which we analyze is that of highest recall for digit 7. This query returned 26  
 27 only three of the 24 images which were classified as sevens, and all three were correct predictions by the classifier. 27  
 28 The rule involved two intersecting lines (intersects( $y_1, y_2$ )), of which the first is described as Line0deg( $y_1$ ), Long( $y_1$ ) 28  
 29 TopLeft( $y_1$ ), TopCenter( $y_1$ ), TopRight( $y_1$ ) which is clearly the characteristic top part of the digit, while the second 29  
 30 line is described as Line45deg( $y_2$ ), BotCenter( $y_2$ ), MidCenter( $y_2$ ), MidRight( $y_2$ ), Long( $y_2$ ), which is the diagonal 30  
 31 part of the digit. The description of the diagonal has a different description than the diagonal line which was part of 31  
 32 the rule for digit 2. Instead of BotLeft( $y$ ), the rule contains a conjunct MidRight( $y$ ). This could be another hint as to 32  
 33 why the 7 shown in Fig. 14d was not classified as correctly, as the digit appears to be leaning slightly to the right, 33  
 34 which makes the diagonal pass through BotLeft, which is in the description of the diagonal for a digit 2 instead of 34  
 35 MidRight which is for digit 7. To conclude if this is the case however would require investigating more rules, since 35  
 36 the one presented covers only a subset of exemplars classified as digit 7. 36

37 For digit 3, the best correct explanation rule returned five of the 26 individuals which were classified as 37  
 38 3, including one misclassified 8. This rule involved seven different lines, thus it was not expected to be 38  
 39 understandable by a user at a glance. However, there was plentiful information for each line in the rule, 39  
 40 which made it possible to visualize. Specifically, regarding the location of the lines, the rule query con- 40  
 41 tained the conjuncts TopCenter( $y_1$ ), BotCenter( $y_2$ ), BotCenter( $y_3$ ), BotRight( $y_4$ ), MidRight( $y_4$ ), TopCenter( $y_5$ ), 41  
 42 MidCenter( $y_6$ ), BotCenter( $y_7$ ). Regarding orientation, five of the seven lines had relevant conjuncts: Line0deg( $y_2$ ), 42  
 43 Line45deg( $y_3$ ), Line90deg( $y_4$ ), Line135deg( $y_6$ ), Line135deg( $y_7$ ). Additionally, three lines had information about 43  
 44 their size Short( $y_4$ ), Short( $y_5$ ), Medium( $y_6$ ). Finally, the explanation rule contained conjuncts showing which lines 44  
 45 intersect each other: intersects( $y_1, y_5$ ), intersects( $y_2, y_3$ ), intersects( $y_2, y_7$ ), intersects( $y_3, y_4$ ), intersects( $y_4, y_6$ ). 45  
 46 This rule-query is visualized in Fig. 13c. An interesting aspect of this explanation, is that the two lines which are at 46  
 47 the top center ( $y_1$  and  $y_5$ ) do not have a specified orientation, while the other four lines are involved in more con- 47  
 48 juncts in the explanation rule query and are described in more detail. This could be an indication of the importance 48  
 49 of these lines for a digit to be classified as a 3. However, these lines could also be a part of other digits such as 5, 49  
 50 which is the next digit which we analyze. 50  
 51 51

1 For the digit 5, the best correct rule query returned four of the 31 positives for the class of which one 1  
 2 is a misclassified 3. It is a very specific query involving seven lines, all of which are described regarding 2  
 3 their orientation  $\text{Line45deg}(y_1)$ ,  $\text{Line135deg}(y_2)$ ,  $\text{Line45deg}(y_3)$ ,  $\text{Line0deg}(y_4)$ ,  $\text{Line0deg}(y_5)$ ,  $\text{Line135deg}(y_6)$ , 3  
 4  $\text{Line90deg}(y_7)$ , and their location  $\text{BotCenter}(y_1)$ ,  $\text{BotLeft}(y_2)$ ,  $\text{BotCenter}(y_2)$ ,  $\text{MidCenter}(y_3)$ ,  $\text{BotCenter}(y_4)$ , 4  
 5  $\text{TopRight}(y_5)$ ,  $\text{TopCenter}(y_5)$ ,  $\text{MidRight}(y_6)$ ,  $\text{MidCenter}(y_6)$ ,  $\text{MidRight}(y_7)$ ,  $\text{BotRight}(y_7)$ . There was no informa- 5  
 6 tion about lines' sizes, and there are five line intersections  $\text{intersects}(y_1, y_4)$ ,  $\text{intersects}(y_1, y_7)$ ,  $\text{intersects}(y_2, y_4)$ , 6  
 7  $\text{intersects}(y_3, y_5)$ ,  $\text{intersects}(y_6, y_7)$ . This query is visualized in Fig. 13e. An interesting aspect of this rule-query is 7  
 8 the fact that it contains a misclassified 3 in its set of certain answers, specifically the three shown in Fig. 14c. In 8  
 9 the context of the proposed framework, the 3 is misclassified because it obeys the correct rule for the digit 5. From 9  
 10 comparing the visualization of the correct rule for the digit 5, and the misclassified 3, we can see that the digit obeys 10  
 11 the rule because the top part of the three is vertically compressed, making it less distinguishable from a 5. This 11  
 12 clearly shows us a potential flaw of the classifier. 12

13 For the digit 4, the best correct rule query returned five of the 22 positives all of which were correct pre- 13  
 14 dictions. The query involved three lines which were all well described regarding their orientation  $\text{Line0deg}(y_1)$ , 14  
 15  $\text{Line45deg}(y_2)$ ,  $\text{Line90deg}(y_3)$ , and their location  $\text{MidCenter}(y_1)$ ,  $\text{MidCenter}(y_2)$ ,  $\text{BotCenter}(y_3)$ ,  $\text{MidCenter}(y_3)$ , 15  
 16  $\text{TopCenter}(y_3)$ . Two lines were also described with respect to their size  $\text{Medium}(y_1)$ ,  $\text{Long}(y_3)$ , and there were two 16  
 17 intersections of lines  $\text{intersects}(y_1, y_2)$ ,  $\text{intersects}(y_1, y_3)$ . This rule is visualized in Fig. 13d. This is a straight- 17  
 18 forward description of the digit four, and as expected it covers only true positives. 18

19 For the digit 6, the best resulting correct rule involved the most variables (each representing a line) out of all 19  
 20 correct rules. It returned four of the 18 positives for the class all of which were correct predictions. Of the eight lines 20  
 21 described in the query, seven had information about their orientation  $\text{Line0deg}(y_1)$ ,  $\text{Line45deg}(y_2)$ ,  $\text{Line45deg}(y_3)$ , 21  
 22  $\text{Line45deg}(y_4)$ ,  $\text{Line90deg}(y_5)$ ,  $\text{Line135deg}(y_7)$ ,  $\text{Line135deg}(y_8)$ . All lines were described with respect to their 22  
 23 position  $\text{BotCenter}(y_1)$ ,  $\text{TopCenter}(y_2)$ ,  $\text{MidCenter}(y_3)$ ,  $\text{BotRight}(y_4)$ ,  $\text{MidCenter}(y_5)$ ,  $\text{MidRight}(y_6)$ ,  $\text{MidRight}(y_7)$ , 23  
 24  $\text{BotCenter}(y_8)$ ,  $\text{MidCenter}(y_8)$ . Additionally four lines had a determined size  $\text{Short}(y_1)$ ,  $\text{Short}(y_6)$ ,  $\text{Short}(y_7)$ , 24  
 25  $\text{Medium}(y_8)$ , and there were five intersections of lines  $\text{intersects}(y_1, y_4)$ ,  $\text{intersects}(y_1, y_8)$ ,  $\text{intersects}(y_2, y_5)$ , 25  
 26  $\text{intersects}(y_5, y_8)$ ,  $\text{intersects}(y_6, y_7)$ . This rule is visualized in Fig. 13f, it is a straight-forward description of a digit 6. 26  
 27

28 For digit 8, the best correct rule query returned four of the 20 positives, all of which are classified correctly. It 27  
 29 involved seven lines, of which five were described with respect to their orientation  $\text{Line45deg}(y_2)$ ,  $\text{Line90deg}(y_3)$ , 28  
 30  $\text{Line135deg}(y_5)$ ,  $\text{Line45deg}(y_6)$ ,  $\text{Line0deg}(y_7)$ , six regarding their location  $\text{MidCenter}(y_1)$ ,  $\text{BotCenter}(y_2)$ , 29  
 31  $\text{BotCenter}(y_3)$ ,  $\text{MidCenter}(y_3)$ ,  $\text{BotCenter}(y_4)$ ,  $\text{TopCenter}(y_6)$ ,  $\text{TopCenter}(y_7)$ , and five regarding their size: 30  
 32  $\text{Medium}(y_1)$ ,  $\text{Short}(y_3)$ ,  $\text{Short}(y_4)$ ,  $\text{Short}(y_5)$ ,  $\text{Short}(y_7)$ . Finally, the rule query involved four intersections of lines 31  
 33  $\text{intersects}(y_1, y_3)$ ,  $\text{intersects}(y_2, y_4)$ ,  $\text{intersects}(y_5, y_7)$ ,  $\text{intersects}(y_6, y_7)$ . This rule is difficult to visualize due to the 32  
 34 missing information (only four of the seven lines have information about both their location and orientation), and 33  
 35 thus is not really useful as an explanation. 34

36 Finally, the best correct rule query for digit 9 returned five of the 15 positives, of which all were correct 35  
 37 predictions by the classifier. It involved six lines which were all thoroughly described regarding their orienta- 36  
 38 tion  $\text{Line0deg}(y_1)$ ,  $\text{Line0deg}(y_2)$ ,  $\text{Line45deg}(y_3)$ ,  $\text{Line90deg}(y_4)$ ,  $\text{Line135deg}(y_5)$ ,  $\text{Line135deg}(y_6)$ , and their 37  
 39 location  $\text{TopCenter}(y_1)$ ,  $\text{MidCenter}(y_2)$ ,  $\text{TopCenter}(y_3)$ ,  $\text{MidCenter}(y_3)$ ,  $\text{MidLeft}(y_3)$ ,  $\text{MidRight}(y_4)$ ,  $\text{BotRight}(y_4)$ , 38  
 40  $\text{TopCenter}(y_5)$ ,  $\text{MidRight}(y_5)$ ,  $\text{MidLeft}(y_6)$ . Additionally, two lines were described regarding their size  $\text{Medium}(y_3)$ , 39  
 41  $\text{Short}(y_6)$ . The query also contained three conjuncts which indicated intersections of lines  $\text{intersects}(y_1, y_3)$ , 40  
 42  $\text{intersects}(y_1, y_5)$ ,  $\text{intersects}(y_2, y_6)$ . This query is visualized in Fig. 13h, it is a straight-forward description of a 41  
 43 digit 9. 42

### 43 5.5. Discussion 44

45 The proposed approach performed similarly with the state-of-the-art on tabular data, was able to detect biases in 46  
 47 the CLEVR-Hans case, detect flaws of the classifier in the Visual Genome case, and provided meaningful expla- 47  
 48 nations even in the MNIST case. In this framework, the resulting explanations depend (almost exclusively) on the 48  
 49 properties of the explanation dataset. In an ideal scenario, end-users trust the explanation dataset, the information 49  
 50 it provides about the exemplars and the terminology it uses. It is like a quiz, or an exam, for the black-box, which 50  
 51 is carefully curated by domain experts. This scenario was simulated in the CLEVR-Hans3 and Visual Genome use- 51

1 cases in which the set of rules produced by the proposed algorithms clearly showed in which cases the black-box 1  
2 classifies items in specific classes, highlighting potential biases acquired during training. The framework is also 2  
3 useful when the explanation dataset is created automatically by leveraging traditional feature extraction, as is shown 3  
4 in the MNIST use-case. In this case, we found the resulting queries to be less understandable than before, which 4  
5 stems mainly from the vocabulary used, since sets of intersecting lines are not easily understandable unless they are 5  
6 visualized. They are also subjectively less trustworthy, since there are usually flaws with most automatic informa- 6  
7 tion extraction procedures. However, since sets of correct rules sufficiently covered the sets of individuals, and rules 7  
8 with exceptions achieved decent performance regarding precision, recall and degree, if an end-user invested time 8  
9 and effort to analyze the resulting rules, they could get a more clear picture about what the black-box is doing. 9

10 We also found interesting the comparison of correct rules with those with exceptions. Correct rules are, in general, 10  
11 more specific than others, as they always have a subset of the pos-set as certain answers. This means that, even 11  
12 though they might be more informative, they tend to involve more conjuncts than rules with exceptions, which in 12  
13 extreme cases could impact understandability. On the other hand, rules with exceptions can be more general, with 13  
14 fewer conjuncts, which could positively impact understandability. Although, utilizing these rules should involve 14  
15 examining the actual exceptions, which could be a lot of work for an end-user, we believe that these exceptions 15  
16 themselves can be very useful in order to detect biases and investigate the behaviour of the model on outliers. In 16  
17 order to eliminate the complexity involved in the investigation of these exceptions, in the future we plan to further 17  
18 study the area and propose a systematic way of exploiting them. These conclusions were apparent in the explanations 18  
19 generated for the class 3 of CLEVR-Hans3 (Table 6), where the best correct rule was very specific, involved five 19  
20 objects and had a relatively low recall (0.42), while the best rule with exceptions was exactly the ground truth class 20  
21 description and had very high precision (0.99). So in this case a user would probably gain more information about 21  
22 the classifier if they examined the rule with exceptions along with the few false positives, instead of examining the 22  
23 best correct rule, or a set of correct rules. 23

24 Another observation we made, is the fact that some conjuncts were more understandable than others when they 24  
25 were part of explanation rules. For instance in MNIST, knowing a line's location and orientation was imperative 25  
26 for understanding the rule via visualization, while conjuncts involving line intersections and sizes seemed not that 26  
27 important, regardless of metrics. This is something which could be leveraged either in explanation dataset construc- 27  
28 tion (for example domain experts weigh concepts and roles depending on their importance for understandability), 28  
29 or in algorithm design (for example a user could provide as input concepts and roles which they want to appear in 29  
30 explanation rules). We are considering these ideas as a main direction for future work which involves developing 30  
31 strategies for choosing which rules are best to show to a user. 31

32 Finally, in the first experiment (5.1), it is shown that KGrules-H can be used to generate explanations in terms 32  
33 of feature data similarly to other rule-based methods, even if it is not the intended use-case. An interesting com- 33  
34 parison for a user study would be between different vocabularies (for example using the features vs using external 34  
35 knowledge). We note here that the proposed approach can always be applied on categorical feature data, since their 35  
36 transformation to an explanation dataset is straight-forward. This would not be the case if we also had numerical 36  
37 continuous features, in which case we would either require more expressive knowledge to represent these features, 37  
38 or that the continuous features be discretized. Another result which motivates us to explore different knowledge 38  
39 expressivities in the future, was the failure of the algorithms to produce a good (w.r.t. the metrics) explanation for 39  
40 the digit 1 in the MNIST experiment (5.4). Specifically, it was difficult to find a query which only returns images 40  
41 of this digit, since a typical description of a "1" is general and tends to always partially describe other digits. This 41  
42 is something which could be mitigated if we allowed for negation in the generated rules, and this is the second 42  
43 direction which we plan to explore in the future. 43  
44 44  
45 45

## 46 6. Conclusions and Future Work 46

47 47  
48 In this work we have further developed a framework for rule-based *post hoc* explanations of black-box classifiers. 48  
49 The explanations are expressed as Horn rules which use terms from a given vocabulary, instead of the classifier's 49  
50 features. This allows for intuitive explanations even when the feature space of the classifier is too complex to be used 50  
51 within understandable rules (pixels, audio signals etc). The rules are also accompanied by theoretical guarantees 51

regarding their *correctness* for explaining the classifier, given what we call an explanation dataset. The idea of the explanation dataset is at the core of our framework, as it is the probe we use to observe the black-box, by feeding it exemplar data and finding rules which explain its output. The explanation dataset also contains the knowledge from which the semantics of the rules are derived. The problem of finding such rules given an explanation dataset was approached as a search problem in the space of semantic queries, by starting with the most specific queries describing positively classified exemplars, and then progressively merging them using heuristic criteria. The queries are then approximately condensed, converted to rules and are shown to the end-user.

There are multiple directions towards which we plan to extend the framework in the future. First of all, we are currently investigating different strategies for choosing which explanation rules are best to show to a user such that they are both informative and understandable. To do this, we also plan to extend our evaluation framework for real world applications to include user studies. Specifically, we are focusing on decision critical domains in which explainability is crucial, such as the medical domain, and in collaboration with domain experts, we are developing explanation datasets, in addition to a crowd-sourced explanation evaluation platform. There are many interesting research questions which we are exploring in this context, such as what constitutes a good explanation dataset, what is a good explanation, and how can we build the trust required for opaque AI to be utilized in such applications.

Another direction which we are currently exploring involves further investigation of our framework in order to fully exploit its capabilities. In particular, the existence of exceptions might have a large impact on both understandability of the produced explanations, and their ability to approximate the behavior of the black-box model. We believe that the systematic study of exceptions could provide useful information to the end-user regarding biases and outliers, and could drive the way towards incorporating local explanations to our framework. Stemming from this, we want to extend the framework, both in theory and in practice, to incorporate different types of explanations. This includes local explanations which explain individual predictions of the black-box, and counterfactual or contrastive explanations which highlight how a specific input should be modified in order for the prediction of the classifier to change. This extension is being researched with the end-user in mind, and we are exploring the merits of providing a blend of explanations (global, local, counterfactual) to an end-user.

A fourth and final direction to be explored involves extending the expressivity of explanation rules, in addition to that of the underlying knowledge. Specifically, the algorithms developed in this work require that if the knowledge has a non-empty TBox, it has to be eliminated via materialization before running the algorithms. Thus, we are exploring ideas for algorithms which generate explanation rules in the case where the underlying knowledge is represented with DL dialects in which the TBox cannot be eliminated, such as DL\_Lite. Finally, regarding the expressivity of explanation rules, we plan to extend the framework to allow for disjunction, which is a straightforward extension, and for negation, which is much harder to incorporate in the framework while maintaining the theoretical guarantees, which we believe are crucial for building trust with end-users.

## References

- [1] B. Goodman and S.R. Flaxman, European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation", *AI Magazine* **38**(3) (2017), 50–57. doi:<https://doi.org/10.1609/aimag.v38i3.2741>.
- [2] M. Turek, Explainable artificial intelligence (XAI), *Defense Advanced Research Projects Agency* (2018). <https://www.darpa.mil/program/explainable-artificial-intelligence>.
- [3] W.J. Murdoch, C. Singh, K. Kumbier, R. Abbasi-Asl and B. Yu, Definitions, methods, and applications in interpretable machine learning, *Proceedings of the National Academy of Sciences* **116**(44) (2019), 22071–22080. doi:<https://doi.org/10.1073/pnas.1900654116>.
- [4] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti and D. Pedreschi, A Survey of Methods for Explaining Black Box Models, *ACM Computing Surveys* **51**(5) (2019), 93:1–93:42. doi:<https://doi.org/10.1145/3236009>.
- [5] F. Lecue, On the role of knowledge graphs in explainable AI, *Semantic Web* **11** (2020), 41–51. doi:<https://doi.org/10.3233/SW-190374>.
- [6] F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi and P.F. Patel-Schneider (eds), *The Description Logic Handbook: Theory, Implementation, and Applications*, Cambridge University Press, Cambridge, United Kingdom, 2003. doi:<https://doi.org/10.1017/CBO9780511711787>.
- [7] J. van der Waa, E. Nieuwburg, A. Cremers and M. Neerinx, Evaluating XAI: A comparison of rule-based and example-based explanations, *Artificial Intelligence* **291** (2021), 103404. doi:<https://doi.org/10.1016/j.artint.2020.103404>.
- [8] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D.A. Shamma et al., Visual genome: Connecting language and vision using crowdsourced dense image annotations, *International journal of computer vision* **123**(1) (2017), 32–73. doi:<https://doi.org/10.1007/s11263-016-0981-7>.

- [9] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár and C.L. Zitnick, Microsoft coco: Common objects in context, in *Lecture Notes in Computer Science*, Vol. 8693, Springer, 2014, pp. 740–755. doi:[https://doi.org/10.1007/978-3-319-10602-1\\_48](https://doi.org/10.1007/978-3-319-10602-1_48).
- [10] G.A. Miller, WordNet: a lexical database for English, *Communications of the ACM* **38**(11) (1995), 39–41. doi:<https://doi.org/10.1145/219717.219748>.
- [11] R. Speer, J. Chin and C. Havasi, ConceptNet 5.5: An Open Multilingual Graph of General Knowledge, Association for the Advancement of Artificial Intelligence, Washington, DC, USA, 2017, pp. 4444–4451. <https://dl.acm.org/doi/10.5555/3298023.3298212>.
- [12] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. Van Kleef, S. Auer et al., Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia, *Semantic web* **6**(2) (2015), 167–195. doi:<https://doi.org/10.3233/SW-140134>.
- [13] M.Q. Stearns, C. Price, K.A. Spackman and A.Y. Wang, SNOMED clinical terms: overview of the development process and project status, S. Bakken, ed., Hanley & Belfus Inc., Northfield, IL, USA, 2001, pp. 662–666–. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2243297/>.
- [14] G. Futia and A. Vetrò, On the integration of knowledge graphs into deep learning models for a more comprehensible AI—Three Challenges for future research, *Information* **11**(2) (2020), 122. doi:<https://doi.org/10.3390/info11020122>.
- [15] M.K. Sarker, N. Xie, D. Doran, M. Raymer and P. Hitzler, Explaining Trained Neural Networks with Semantic Web Technologies: First Steps, in: *Neural-Symbolic Learning and Reasoning 2017*, T.R. Besold, A.d. Garcez and I. Noble, eds, CEUR Workshop Proceedings, Vol. 2003, CEUR-WS.org, Aachen, Germany, 2017. [https://ceur-ws.org/Vol-2003/NeSy17\\_paper4.pdf](https://ceur-ws.org/Vol-2003/NeSy17_paper4.pdf).
- [16] C. Panigutti, A. Perotti and D. Pedreschi, Doctor XAI: an ontology-based approach to black-box sequential data classification explanations, in: *Proceedings of the 2020 conference on fairness, accountability, and transparency*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 629–639. doi:<https://doi.org/10.1145/3351095.3372855>.
- [17] Z.A. Daniels, L.D. Frank, C.J. Menart, M. Raymer and P. Hitzler, A framework for explainable deep neural models using external knowledge graphs, in: *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II*, Vol. 11413, SPIE, 2020, p. 114131C, International Society for Optics and Photonics. doi:<https://doi.org/10.1117/12.2558083>.
- [18] I. Tiddi and S. Schlobach, Knowledge Graphs as tools for Explainable Machine Learning: a survey, *Artificial Intelligence* **302** (2022), 103627. doi:<https://doi.org/10.1016/j.artint.2021.103627>.
- [19] R. Confalonieri, T. Weyde, T.R. Besold and F. Moscoso del Prado Martín, Using ontologies to enhance human understandability of global post-hoc explanations of black-box models, *Artificial Intelligence* **296** (2021), 103471. doi:<https://doi.org/10.1016/j.artint.2021.103471>. <https://www.sciencedirect.com/science/article/pii/S0004370221000229>.
- [20] F. Croce, G. Cima, M. Lenzerini and T. Catarci, Ontology-based explanation of classifiers, in *CEUR Workshop Proceedings*, Vol. 2578, CEUR-WS.org, 2020. <https://ceur-ws.org/Vol-2578/PIE3.pdf>.
- [21] A. Hogan, E. Blomqvist, M. Cochez, C. d’Amato, G. de Melo, C. Gutiérrez, J.E.L. Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A.N. Ngomo, S.M. Rashid, A. Rula, L. Schmelzeisen, J.F. Sequeda, S. Staab and A. Zimmermann, Knowledge Graphs, *ACM Computing Surveys* **54** (2020). doi:<https://doi.org/10.1145/3447772>.
- [22] B. Mittelstadt, C. Russell and S. Wachter, Explaining explanations in AI, in: *Proceedings of the conference on fairness, accountability, and transparency*, Association for Computing Machinery, New York, NY, USA, 2019, pp. 279–288. doi:<https://doi.org/10.1145/3287560.3287574>.
- [23] V. Praher, K. Prinz, A. Flexer and G. Widmer, On the Veracity of Local, Model-agnostic Explanations in Audio Classification: Targeted Investigations with Adversarial Examples, *arXiv preprint arXiv:2107.09045* (2021). doi:<https://doi.org/10.48550/arXiv.2107.09045>.
- [24] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nature Machine Intelligence* **1**(5) (2019), 206–215. doi:<https://doi.org/10.1038/s42256-019-0048-x>.
- [25] H. Yang, C. Rudin and M. Seltzer, Scalable Bayesian Rule Lists, in: *Proceedings of the 34th International Conference on Machine Learning*, D. Precup and Y.W. Teh, eds, Proceedings of Machine Learning Research, Vol. 70, PMLR, 2017, pp. 3921–3930. <https://proceedings.mlr.press/v70/yang17h.html>.
- [26] Y. Ming, H. Qu and E. Bertini, RuleMatrix: Visualizing and Understanding Classifiers with Rules, *IEEE Transactions on Visualization and Computer Graphics* **25**(1) (2019), 342–352. doi:<https://doi.org/10.1109/TVCG.2018.2864812>.
- [27] J. Lehmann, S. Bader and P. Hitzler, Extracting reduced logic programs from artificial neural networks, *Applied Intelligence* **32**(3) (2010), 249–266. doi:<https://doi.org/10.1007/s10489-008-0142-y>.
- [28] R. Guidotti, A. Monreale, S. Ruggieri, D. Pedreschi, F. Turini and F. Giannotti, Local rule-based explanations of black box decision systems, *arXiv preprint arXiv:1805.10820* (2018). doi:<https://doi.org/10.48550/arXiv.1805.10820>.
- [29] M.T. Ribeiro, S. Singh and C. Guestrin, Anchors: High-precision model-agnostic explanations, in: *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32, Public Knowledge Project, 2018. doi:<https://doi.org/10.1609/aaai.v32i1.11491>.
- [30] M.W. Craven and J.W. Shavlik, Extracting Tree-Structured Representations of Trained Networks, in: *Advances in Neural Information Processing Systems*, Vol. 8, D. Touretzky, M.C. Mozer and M. Hasselmo, eds, MIT Press, 1995, pp. 24–30. <http://papers.nips.cc/paper/1152-extracting-tree-structured-representations-of-trained-networks>.
- [31] Y. Zhou and G. Hooker, Interpreting models via single tree approximation, *arXiv preprint arXiv:1610.09036* (2016). doi:<https://doi.org/10.48550/arXiv.1610.09036>.
- [32] G. Ciravegna, F. Giannini, M. Gori, M. Maggini and S. Melacci, Human-Driven FOL Explanations of Deep Learning, in: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, C. Bessiere, ed., International Joint Conferences on Artificial Intelligence Organization, 2020, pp. 2234–2240, Main track. doi:<https://doi.org/10.24963/ijcai.2020/309>.
- [33] D. Pedreschi, F. Giannotti, R. Guidotti, A. Monreale, S. Ruggieri and F. Turini, Meaningful Explanations of Black Box AI Decision Systems, AAAI Press, Palo Alto, California USA, 2019, pp. 9780–9784. doi:<https://doi.org/10.1609/aaai.v33i01.33019780>.



- [34] A.B. Arrieta, N.D. Rodríguez, J.D. Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila and F. Herrera, Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, *Information Fusion* **58** (2020), 82–115. doi:<https://doi.org/10.1016/j.inffus.2019.12.012>. <https://www.sciencedirect.com/science/article/pii/S1566253519308103>.
- [35] Q.T. Tran, C.Y. Chan and S. Parthasarathy, Query reverse engineering, *The VLDB Journal* **23**(5) (2014), 721–746. doi:<https://doi.org/10.1007/s00778-013-0349-3>.
- [36] D.M.L. Martins, Reverse engineering database queries from examples: State-of-the-art, challenges, and research opportunities, *Information Systems* **83** (2019), 89–100. doi:<https://doi.org/10.1016/j.is.2019.03.002>. <https://www.sciencedirect.com/science/article/pii/S0306437918300978>.
- [37] H. Li, C.-Y. Chan and D. Maier, Query from Examples: An Iterative, Data-Driven Approach to Query Construction, *Proceedings of the VLDB Endowment* **8**(13) (2015), 2158–2169. doi:<https://doi.org/10.14778/2831360.2831369>.
- [38] A. Bonifati, R. Ciucanu and S. Staworko, Learning Join Queries from User Examples, *ACM Transactions on Database Systems* **40**(4) (2016). doi:<https://doi.org/10.1145/2818637>.
- [39] C. Wang, A. Cheung and R. Bodik, Synthesizing Highly Expressive SQL Queries from Input-Output Examples, *SIGPLAN Not.* **52**(6) (2017), 452–466. doi:<https://doi.org/10.1145/3140587.3062365>.
- [40] M. Zhang, H. Elmeleegy, C.M. Procopiuc and D. Srivastava, Reverse Engineering Complex Join Queries, in: *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, SIGMOD '13, Association for Computing Machinery, New York, NY, USA, 2013, pp. 809–820. ISBN 9781450320375. doi:<https://doi.org/10.1145/2463676.2465320>.
- [41] D. Calvanese, G.D. Giacomo, D. Lembo, M. Lenzerini and R. Rosati, Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family, *Journal of Automated Reasoning* **39**(3) (2007), 385–429. doi:<https://doi.org/10.1007/s10817-007-9078-x>.
- [42] B.C. Grau, B. Motik, G. Stoilos and I. Horrocks, Computing Datalog Rewritings Beyond Horn Ontologies, in: *IJCAI*, F. Rossi, ed., AAAI Press / International Joint Conferences on Artificial Intelligence, Menlo Park, California, USA, 2013, pp. 832–838. <https://www.ijcai.org/Proceedings/13/Papers/129.pdf>.
- [43] A. Chortaras, M. Giazizoglou and G. Stamou, Inside the Query Space of DL Knowledge Bases, in: *Proceedings of the 32nd International Workshop on Description Logics*, M. Šimkus and G. Weddell, eds, International Workshop on Description Logics, Vol. 2373, CEUR-WS.org, Aachen, Germany, 2019. <https://ceur-ws.org/Vol-2373/paper-11.pdf>.
- [44] D. Trivela, G. Stoilos, A. Chortaras and G. Stamou, Resolution-based rewriting for Horn-SHIQ ontologies, *Knowledge and Information Systems* **62**(1) (2020), 107–143. doi:<https://doi.org/10.1007/s10115-019-01345-2>.
- [45] A. Petrova, E.V. Kostylev, B.C. Grau and I. Horrocks, Query-Based Entity Comparison in Knowledge Graphs Revisited, in: *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Auckland, New Zealand, October 26-30, 2019, Proceedings, Part I*, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I.F. Cruz, A. Hogan, J. Song, M. Lefrançois and F. Gandon, eds, Lecture Notes in Computer Science, Vol. 11778, Springer, 2019, pp. 558–575. doi:[https://doi.org/10.1007/978-3-030-30793-6\\_32](https://doi.org/10.1007/978-3-030-30793-6_32).
- [46] E. Dervakos, O. Menis-Mastromichalakis, A. Chortaras and G. Stamou, Computing Rule-Based Explanations of Machine Learning Classifiers using Knowledge Graphs, 2022. doi:<https://doi.org/10.48550/arXiv.2202.03971>.
- [47] V. Gutiérrez-Basulto, J.C. Jung and L. Sabellek, Reverse Engineering Queries in Ontology-Enriched Systems: The Case of Expressive Horn Description Logic Ontologies, in: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, International Joint Conferences on Artificial Intelligence Organization, 2018, pp. 1847–1853. doi:<https://doi.org/10.24963/ijcai.2018/255>.
- [48] M. Ortiz, Ontology-Mediated Queries from Examples: a Glimpse at the DL-Lite Case, in: *GCAI 2019. Proceedings of the 5th Global Conference on Artificial Intelligence*, D. Calvanese and L. Iocchi, eds, EPiC Series in Computing, Vol. 65, EasyChair, 2019, pp. 1–14. ISSN 2398-7340. doi:<https://doi.org/10.29007/jhtz>. <https://easychair.org/publications/paper/c3CT>.
- [49] P. Barceló and M. Romero, The Complexity of Reverse Engineering Problems for Conjunctive Queries, *20th International Conference on Database Theory (ICDT 2017)* **68** (2017), 7:1–7:17. ISBN 978-3-95977-024-8. doi:<https://doi.org/10.4230/LIPIcs.ICDT.2017.7>. <http://drops.dagstuhl.de/opus/volltexte/2017/7052>.
- [50] J. Jung, C. Lutz, H. Pulcini and F. Wolter, Logical Separability of Incomplete Data under Ontologies, in: *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, D. Calvanese, E. Erdem and M. Thielscher, eds, IJCAI Organization, 2020, pp. 517–528. doi:<https://doi.org/10.24963/kr.2020/52>.
- [51] M. Arenas, G.I. Diaz and E.V. Kostylev, Reverse Engineering SPARQL Queries, in: *Proceedings of the 25th International Conference on World Wide Web*, J. Bourdeau, J. Hendler, R. Nkambou, I. Horrocks and B.Y. Zhao, eds, WWW '16, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2016, pp. 239–249. doi:<https://doi.org/10.1145/2872427.2882989>.
- [52] G. Cima, F. Croce and M. Lenzerini, Query Definability and Its Approximations in Ontology-Based Data Management, in *CIKM '21*, Association for Computing Machinery, New York, NY, USA, 2021, pp. 271–280. ISBN 9781450384469. doi:<https://doi.org/10.1145/3459637.3482466>.
- [53] G. Gottlob and C.G. Fermüller, Removing Redundancy from a Clause, *Artificial Intelligence* **61**(2) (1993), 263–289. doi:[https://doi.org/10.1016/0004-3702\(93\)90069-N](https://doi.org/10.1016/0004-3702(93)90069-N).
- [54] W.W. Cohen, A. Borgida and H. Hirsh, Computing Least Common Subsumers in Description Logics, in: *Proceedings of the 10th National Conference on Artificial Intelligence*, AAAI Press, Palo Alto, California, USA, 1992, pp. 754–760. <https://aaai.org/papers/00754-computing-least-common-subsumers-in-description-logics/>.
- [55] R. Küsters and R. Molitor, Structural Subsumption and Least Common Subsumers in a Description Logic with Existential and Number Restrictions, *Studia Logica* **81**(2) (2005), 227–259. doi:<https://doi.org/10.1007/s11225-005-3705-5>.
- [56] F. Baader, B. Sertkaya and A. Turhan, Computing the least common subsumer w.r.t. a background terminology, *Journal of Applied Logic* **5**(3) (2007), 392–420. doi:<https://doi.org/10.1016/j.jal.2006.03.002>.

- [57] F.M. Donini, S. Colucci, T.D. Noia and E.D. Sciascio, A Tableaux-based Method for Computing Least Common Subsumers for Expressive Description Logics, in: *Proceedings of the 22nd International Workshop on Description Logics (DL 2009)*, B.C. Grau, I. Horrocks, B. Motik and U. Sattler, eds, CEUR-WS.org, 2009. [http://ceur-ws.org/Vol-477/paper\\_22.pdf](http://ceur-ws.org/Vol-477/paper_22.pdf).
- [58] J. Laguarda, F. Hueto and B. Subirana, COVID-19 artificial intelligence diagnosis using only cough recordings, *IEEE Open Journal of Engineering in Medicine and Biology* **1** (2020), 275–281. doi:<https://doi.org/10.1109/OJEMB.2020.3026928>.
- [59] R. Speer, J. Chin and C. Havasi, ConceptNet 5.5: An Open Multilingual Graph of General Knowledge (2016). doi:<https://doi.org/10.48550/arXiv.1612.03975>.
- [60] G. Diaz, M. Arenas and M. Benedikt, SPARQLByE: Querying RDF Data by Example **9**(13) (2016), 1533–1536–. doi:<https://doi.org/10.14778/3007263.3007302>.
- [61] J. Liartis, E. Dervakos, O. Menis-Mastromichalakis, A. Chortaras and G. Stamou, Semantic Queries Explaining Opaque Machine Learning Classifiers, in: *Proceedings of the Workshop on Data meets Applied Ontologies in Explainable AI (DAO-XAI 2021)*, R. Confalonieri, O. Kutz and D. Calvanese, eds, CEUR Workshop Proceedings, Vol. 2998, CEUR-WS.org, 2021. <http://ceur-ws.org/Vol-2998/paper2.pdf>.
- [62] M.M. Zloof, Query-by-example: the invocation and definition of tables and forms, in *VLDB '75*, Association for Computing Machinery, New York, NY, USA, 1975. doi:<https://doi.org/10.1145/1282480.1282482>.
- [63] R. Kontchakov and M. Zakharyashev, *An Introduction to Description Logics and Query Rewriting*, in: *Reasoning Web. Reasoning on the Web in the Big Data Era: 10th International Summer School 2014, Proceedings*, M. Koubarakis, G. Stamou, G. Stoilos, I. Horrocks, P. Kolaitis, G. Lausen and G. Weikum, eds, Lecture Notes in Computer Science, Vol. 8714, Springer International Publishing, Cham, 2014, pp. 195–244. ISBN 978-3-319-10587-1. doi:[https://doi.org/10.1007/978-3-319-10587-1\\_5](https://doi.org/10.1007/978-3-319-10587-1_5).
- [64] R. Kontchakov, C. Lutz, D. Toman, F. Wolter and M. Zakharyashev, The Combined Approach to Ontology-Based Data Access, T. Walsh, ed., AAAI Press/International Joint Conferences on Artificial Intelligence, Menlo Park, California, USA, 2011, pp. 2656–2661. doi:<http://doi.org/10.5591/978-1-57735-516-8/IJCAI11-442>.
- [65] B. Glimm, Y. Kazakov and T. Tran, Ontology Materialization by Abstraction Refinement in Horn SHOIF, AAAI Press, Palo Alto, California, USA, 2017, pp. 1114–1120. doi:<https://doi.org/10.1609/aaai.v31i1.10691>.
- [66] B. ten Cate and V. Dalmau, The Product Homomorphism Problem and Applications, in: *18th International Conference on Database Theory (ICDT 2015)*, M. Arenas and M. Ugarte, eds, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 31, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2015, pp. 161–176. ISSN 1868-8969. ISBN 978-3-939897-79-8. doi:<https://doi.org/10.4230/LIPIcs.ICDT.2015.161>. <http://drops.dagstuhl.de/opus/volltexte/2015/4983>.
- [67] S.H. Bokhari, On the Mapping Problem, *IEEE Transactions on Computers* **C-30**(3) (1981), 207–214. doi:<https://doi.org/10.1109/TC.1981.1675756>.
- [68] H.M. Grindley, P.J. Artymiuk, D.W. Rice and P. Willett, Identification of Tertiary Structure Resemblance in Proteins Using a Maximal Common Subgraph Isomorphism Algorithm, *Journal of Molecular Biology* **229**(3) (1993), 707–721. doi:<https://doi.org/10.1006/jmbi.1993.1074>. <https://www.sciencedirect.com/science/article/pii/S0022283683710740>.
- [69] J. Xu, GMA: A Generic Match Algorithm for Structural Homomorphism, Isomorphism, and Maximal Common Substructure Match and Its Applications, *Journal of Chemical Information and Computer Sciences* **36**(1) (1996), 25–34. doi:<https://doi.org/10.1021/ci950061u>.
- [70] A. Egozi, Y. Keller and H. Guterman, A Probabilistic Approach to Spectral Graph Matching, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**(01) (2013), 18–27. doi:<https://doi.org/10.1109/TPAMI.2012.51>.
- [71] M. Leordeanu and M. Hebert, A spectral technique for correspondence problems using pairwise constraints, in: *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, Vol. 2, IEEE, 2005, pp. 1482–1489. doi:<https://doi.org/10.1109/ICCV.2005.20>.
- [72] W. Stammer, P. Schramowski and K. Kersting, Right for the Right Concept: Revising Neuro-Symbolic Concepts by Interacting with their Explanations, *arXiv preprint arXiv:2011.12854* (2020). doi:<https://doi.org/10.48550/arXiv.2011.12854>.
- [73] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva and A. Torralba, Places: A 10 million image database for scene recognition, *IEEE transactions on pattern analysis and machine intelligence* **40**(6) (2017), 1452–1464. doi:<https://doi.org/10.1109/TPAMI.2017.2723009>.
- [74] Y. LeCun, L. Bottou, Y. Bengio and P. Haffner, Gradient-Based Learning Applied to Document Recognition, in: *Proceedings of the IEEE*, **86**(11):2278-2324, Vol. 86, G. Setti, ed., IEEE, 1998, pp. 2278–2324. doi:<https://doi.org/10.1109/5.726791>.
- [75] K. He, X. Zhang, S. Ren and J. Sun, Deep Residual Learning for Image Recognition (2015). doi:<https://doi.org/10.48550/arXiv.1512.03385>.
- [76] H. Wang, Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel and X. Hu, Score-CAM: Score-Weighted Visual Explanations for Convolutional Neural Networks, arXiv, 2019. doi:<https://doi.org/10.48550/ARXIV.1910.01279>. <https://arxiv.org/abs/1910.01279>.
- [77] R. Poyiadzi, K. Sokol, R. Santos-Rodríguez, T.D. Bie and P.A. Flach, FACE: Feasible and Actionable Counterfactual Explanations, in *AIES '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 344–350. doi:<https://doi.org/10.1145/3375627.3375850>.
- [78] T. Lindeberg, *Scale-Space*, in: *Wiley Encyclopedia of Computer Science and Engineering*, B.W. Wah, ed., John Wiley & Sons, Ltd, 2008, pp. 2495–2504. ISBN 9780470050118. doi:<https://doi.org/10.1002/9780470050118.ecse609>.
- [79] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L. Li, D.A. Shamma, M.S. Bernstein and L. Fei-Fei, Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations, *International Journal of Computer Vision* **123**(1) (2017), 32–73. doi:<https://doi.org/10.1007/s11263-016-0981-7>.
- [80] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C.L. Zitnick and R.B. Girshick, CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning (2016). doi:<https://doi.org/10.48550/arXiv.1612.06890>.
- [81] F. Baader, I. Horrocks, C. Lutz and U. Sattler, *An Introduction to Description Logic*, Cambridge University Press, 2017. doi:<https://doi.org/10.1017/9781139025355>.
- [82] J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, 2019. doi:<https://doi.org/10.48550/arXiv.1810.04805>.

- 1 [83] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer and V. Stoyanov, RoBERTa: A Robustly Optimized  
2 BERT Pretraining Approach, 2019. doi:<https://doi.org/10.48550/arXiv.1907.11692>.
- 3 [84] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C.H. So and J. Kang, BioBERT: a pre-trained biomedical language representation model for  
4 biomedical text mining, *Bioinformatics* **36** (2020), 1234–1240. doi:<https://doi.org/10.1093/bioinformatics/btz682>.
- 5 [85] G.A. Miller, WordNet: A Lexical Database for English, *Communications of the ACM* **38**(11) (1995), 39–41–.  
6 doi:<https://doi.org/10.1145/219717.219748>.
- 7 [86] M. Alirezaie, M. Långkvist, M. Sioutis and A. Loutfi, A symbolic approach for explaining errors in image classification tasks, 2018.  
8 <https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1233674&dswid=-9184>.
- 9 [87] H. Wang, F. Zhang, X. Xie and M. Guo, DKN: Deep Knowledge-Aware Network for News Recommendation, in: *Proceedings of the 2018*  
10 *World Wide Web Conference on World Wide Web*, International World Wide Web Conferences Steering Committee, Republic and Canton  
11 of Geneva, CHE, 2018, pp. 1835–1844. doi:<https://doi.org/10.1145/3178876.3186175>.
- 12 [88] Q. Ai, V. Azizi, X. Chen and Y. Zhang, Learning heterogeneous knowledge base embeddings for explainable recommendation, *Algorithms*  
13 **11**(9) (2018), 137. doi:<https://doi.org/10.3390/a11090137>.
- 14 [89] V.S. Silva, A. Freitas and S. Handschuh, Exploring knowledge graphs in an interpretable composite approach for text entailment, AAAI  
15 Press, Palo Alto, California, USA, 2019, pp. 7023–7030. doi:<https://doi.org/10.1609/aaai.v33i01.33017023>.
- 16 [90] Y. Zhang, P. Tiffo, A. Leonardis and K. Tang, A Survey on Neural Network Interpretability, *IEEE Transactions on Emerging Topics in*  
17 *Computational Intelligence* **5**(5) (2021), 726–742. doi:<https://doi.org/10.1109/TETCI.2021.3100641>.
- 18 [91] G. Filandrianos, K. Thomas, E. Dervakos and G. Stamou, Conceptual Edits as Counterfactual Explanations, in: *Proceedings of the AAAI*  
19 *2022 Spring Symposium on Machine Learning and Knowledge Engineering for Hybrid Intelligence (AAAI-MAKE 2022)*, A. Martin,  
20 K. Hinkelmann, H.-G. Fill, A. Gerber, D. Lenat, R. Stolle and F. van Harmelen, eds, CEUR Workshop Proceedings, CEUR-WS.org, Palo  
21 Alto, California, USA, 2022. <https://ceur-ws.org/Vol-3121/paper6.pdf>.
- 22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51