

Reasoning over Knowledge Graphs in an Intuitionistic Description Logic^{*}

Bernardo Alkmim, Edward Haeusler, and Daniel Schwabe

Departamento de Informática, Pontifícia Universidade Católica do Rio de Janeiro, Brazil

bpalkmim@gmail.com

hermann@inf.puc-rio.br

dschwabe@inf.puc-rio.br

Abstract. In this work we present a way to model and reason over Knowledge Graphs via an Intuitionistic Description Logic called iALC. We also introduce a Natural Deduction System for iALC to reason over our modelling of the information of the Knowledge Graphs. Furthermore, we apply this modelling to some examples in a context that aims to support the concepts of Trust, Privacy, and Transparency.

Keywords: Description Logic · Intuitionistic Logic · iALC · Knowledge Graphs · Natural Deduction.

1 Introduction

Throughout the years there have been several different ways to represent knowledge, either by trying to mimic human cognitive behaviour or by other approaches. Knowledge Base Systems (KBSs) [13] are one of such methods, derived from Expert Systems, explained in [5] and [10].

One way to represent KBSs is via Knowledge Graphs (KGs), which can be seen in [15] and [18]. With their graph-like structure, KGs offer a versatile way to represent data that postpones the necessity for a schema, which adapts well to different domains. The fact that they are graphs is, as well, another positive point, since this mathematical structure has extensive literature and theory [4] behind it. This is especially useful when one utilises Machine Learning, since KGs can connect it to Knowledge Representation thus imbuing it with tools to solve certain questions in the data, such as Data Insufficiency or Zero-shot Learning.

One other way to represent knowledge is via the usage of Description Logics (DL) [2]. DLs consist in the representation of entities called *concepts*, *roles*, and *individuals*, as well as their relationships, based around *axioms*. This type of Logic comes in different *flavours*, depending on which properties one wishes to utilise. One of the attractive points of DLs is the possibility to make syntactic reasoning (due to being a Logic) over the associated KB, via whichever deduction system is more adequate.

There is, as well, especially for the domain of Law, the option of utilising Deontic Logic, but it is criticised (as in [19], for instance) for not making explicit the models and

^{*} The authors gratefully acknowledge financial support from CNPq.

theories for the making of legal KBSs, among other reasons. By choosing Description Logic, one may utilise a calculus for deductions, similarly to Deontic Logic, but without facing the same modelling issues.

We then end up with two means to represent data: KGs, which are versatile and can represent data more accurately, and DLs, over which one can reason while extracting more direct *meaning* from the reasoning itself at the same time. Due to their equivalence in representation, we are able to reap the good parts of both these methods over the same KBS.

In this article, we present our efforts in translating a representation of a KBS that focuses on the concepts of Trust, Privacy, and Transparency of information from a KG to a DL, namely, iALC, in order to reason over it.

iALC is a Logic that has been used to represent and reason over the domain of Law, and its structure is based on Kelsenian Jurisprudence. Thus, legal individuals can be seen as Valid Legal Statements (VLSs) in iALC, and the notion of precedence between Laws (as per Theory of Legal Order in [3]) can be directly represented in it as a relation between worlds in a Kripke Model. The choice to utilise a Description Logic instead of a type of the more usual Deontic Logic comes from the existence of certain conceptual problems of Deontic Logic, especially when one considers the notion of validity of a normative sentence.

In section 2, we introduce the concept of KGs and the framework which allows them to better deal with concepts of Trust, Privacy, and Transparency of information.

In section 3, we explain the main concepts of the DL we choose to reason with, iALC, and explain why it was chosen to do so.

In section 4, we contextualise the KBS and give some examples of formalizations in iALC and the reasonings made, as well as discuss our decision making throughout this process.

In section 5, we conclude the article and point to future directions of this work.

In the appendix A, we present the Natural Deduction system rules created for iALC to make the deduction itself.

2 Knowledge Graphs

There are multiple different definitions for KGs in the literature and varied implementations of them (all discussed thoroughly in [12]), but they all share a central idea: they utilise graphs to represent data and knowledge. The concept comes from the 70s [16], but the term itself has been popularised by Google in 2012,¹ when they implemented the concept to create a system to support moving from search to question answering.

In general, the nodes in a KG represent entities of the real world and the edges, relations between them. This is, however, a rather simplistic approach that is not much different from a labelled graph. In order to talk about knowledge, it is necessary to have a supplementary *ontology*, or a set of *rules* of the domain to be represented. One may use quantification of variables, as well as separate one big KG into smaller ones, in order to isolate semantically certain parts of the original graph.

¹ <https://www.blog.google/products/search/introducing-knowledge-graph-things-not/>

In [17], the authors present a framework called KG Usage, which allows KGs to support Trust, Transparency, and Privacy concerns in its modelling. This allows for usage in applications for whom decisions based on security or disclosure of information of knowledge is important.

3 iALC

Description Logics (DL) were already proposed for semantic analysis of natural language utterances in [6]. In [7] the DL iALC [9] was used as the basis of a solver/reasoner for the questions in the exam. iALC is based on the canonical ALC, and was by definition made to deal with law systems. In it, valid legal statements (VLS) are not propositions, but individuals in a legal ontology, i. e. one law cannot be true or false, it just either exists or not in a concept. The propositions we consider are some of the *concepts* of the ontology (i. e. concepts of the form $a : C$ or $C \sqsubseteq D$, for a nominal and C, D concepts), which represent the legal systems that can hold different kinds of laws. They also have a precedence relationship, derived by the hierarchy of individual laws of Kelsenian Jurisprudence [14]. As an example, one can imagine Brazilian Law or other similar system. It has the Constitution, comprised of some general principles that must precede every other (usually more specific) law that is based on or created after it. With this notion, one can create different tiers of laws in order to not generate legal antinomies with the laws of higher tiers, and those *higher* laws must precede the *lower* ones.

In [1] iALC is utilised to model and solve questions from the OAB Exam.² In this article, the characteristics of this logic are more thoroughly explained, especially in how they make it well suited to deal with normative sentences, especially laws.

In general, in iALC, legal individuals are specific combinations of laws that represent an artificial legal being, which can validate or not certain concepts. They do not, however, represent fully individuals outside of the legal expressiveness of the logic; they are mere legal projections of real-life concepts. For example, $john : Attorney$ states that the legal individual *john*, which represents a legal document that proves that John is an attorney, is part of the concept *Attorney*, which represents the set of laws relating to attorneys. In general, to simplify, individuals like *john* are assumed to always be, from the start, in a nominal.³ Notice that everything in this example is in the ontology on laws.

iALC is also an intuitionistic logic. Intuitionistic logic is constructive by nature [11]. Intuitionistically, $\neg A$ being true can be understood as there being no way to construct a proof of A (or that, by assuming A to be true, one is led to a contradiction). With this, we have, for instance, that $\neg\neg A$ being true does not imply in A being true, which is a valid implication in classical logic. In fact, intuitionistic logic is weaker than classical logic (see [11]), but it allows for non-conflicting existence of otherwise thought-to-be contradictory logical formulas in Description Logic, particularly. In the domain of

² The Brazilian Bar Exam, which candidates have to pass in order to be able to practice law.

³ A nominal is a concept defined extensionally, consisting of only one individual, by which it is referenced. A nominal for *john* is $\{john\}$. Throughout this article we will omit this notation, referring to the nominals without the brackets.

law, this allows, for instance, the existence of a model representing two distinct legal systems, one in which the death penalty is not allowed, and the other in which it is not. In a regular (i.e. Classical) DL, this cannot be represented directly when modelling, since a Kripke model for a Classical Logic collapses all worlds (the legal individuals) to the same one, causing a contradiction by allowing and prohibiting the death penalty at the same time. This would force us to find a more conflated way to insert Legal Individuals into the logic, at the risk of losing legibility, and, even worse, soundness.

This non-conflicting existence of different sets of rules relates directly to the existence of different KGs in section 4, each with their own set of rules, which may contradict those of other KGs. iALC not only allows for those to coexist, but it also has means to solve those apparent contradictions.

4 Some Examples

The examples utilised here are from [17], further explained in the third section of that article. To summarise, in a fictional context a disaster occurs and a person, Ed, wishes to donate to a foundation in order to give aid to the people affected by it. He, however, wants to be sure that his money will be well spent, so he needs to *trust* the organisation to which he intends to give his money. There is another person, George, who is an officer in a non-governmental organisation (NGO) called ReliefOrg, whom Ed does not trust. Assuming that there is a law that requires NGOs to disclose publicly (*transparency*) their officers, Ed must have access to this information from ReliefOrg. Even so, George may not want to disclose his personal information to others (right to *privacy*), creating a conflict that must be resolved. In this case, since it is in the law that this kind of information must be disclosed, the transparency law has priority.

The approach there postulates the existence of a *trusted graph* containing all the statements an agent will rely upon to decide to take some action. Rules define when a statement should be included in this graph (*trust*), as well as when some action should be allowed (*transparency*) or denied (*privacy*). In the example, there are also other scenarios in this same context that relate the concepts of Trust, Privacy, and Transparency in similar ways.

It is worth noting that this hierarchy of normative sentences, i.e. rules, laws, was the starting point to formalise this information in iALC, since the logic was created to deal with laws from the cradle.

We will now show a modelling of basic information and some examples of rules of this scenario, based on the same rules from the original article, that modelled them in a KG divided in several Knowledge Items, each one being a nanopublication [8].

4.1 Base knowledge

Part of the information is shared by all agents, and is the base knowledge for the scenario, namely: ReliefOrg is a non-profit organisation, AuditInc is a certification agency, George is an officer of ReliefOrg and of AuditInc.

For the base graph knowledge, we have the following formulas:

relief : *NonProfit*^L

$$\begin{aligned}
& \text{audit} : \text{CertAgency}^L \\
& \text{NonProfit} \sqsubseteq \text{Organisation}^L \\
& \text{CertAgency} \sqsubseteq \text{Organisation}^L \\
& \text{george} \text{ officerInOrg } \text{relief} \\
& \text{george} \text{ officerInOrg } \text{audit}
\end{aligned}$$

In the formulas above, we can see that the individuals are represented by nominals in iALC, as would be expected. We utilise a *concept* to represent a category or set. In the first formula, we have that the individual *relief*, an individual representing ReliefOrg, is part of the concept *NonProfit*, a concept for non-profit organisations. One concept can be a subset of another, as in the third formula, which states that non-profit organisations are a subset of organisations in general. When a relation between two individuals is established, it is translated into a *role*. In the fifth formula, for example, the role *officerInOrg* relates the individuals *george* and *relief*.

To model the rules in iALC, we needed to make them somewhat general to be able to chain formulas with one another. So, we used generic labels when modelling e.g. $\text{relief} : \text{NonProfit}^L$, where L is any list of labels for the concept *NonProfit*.

To make the actual reasoning, we developed a Natural Deduction (ND)⁴ system for iALC (whose rules are all in appendix A), which has labels to facilitate the chaining of existential restrictions. Our ND system needs these chainings in order to connect in the same reasoning individuals that are related indirectly (via transitivity) in a fully *TBOX-esque* fashion. For instance, when an *agent* intends to do an *action* to an *organisation* that has a certain *officer*, we need to encapsulate all these different relations between distinct agents in one single formula in order to make more direct and structurally simple reasonings without losing meaning. Luckily, it is possible in DLs via applying restrictions sequentially on concepts. The labels on concepts also aid this chaining by giving the context at all times, thus maintaining the soundness of the system. This chaining of formulas is explained more concretely in section 4.2.

The rules utilised in this document are, mostly, \exists – *intro*, \exists – *elim*, \sqsubseteq – *elim*, and \sqcap – *intro*. If applicable, the versions of the rules are the ones that have nominals.

The rules used have to do with our modelling choices: most of the modelling involved existential restrictions (\exists), subjuncts of concepts (\sqsubseteq), and conjunctions of concepts (\sqcap). Also, since the cases are always instantiated, all the rules involved nominals, to reflect the information as faithfully as possible.

4.2 Rule Ed1 (TRUST)

We begin with the first rule of knowledge graph for the agent Ed: Ed wants to donate to a certified not-for-profit organisation. More precisely, if an organisation *org* has (some) officer *of*, financial record *fin*, *fin* is audited by organisation *aud*, *aud* is certified by organisation *cert*, *cert* is a Certification Agency and the provenance graph for *cert* includes the fact that it is an accredited agency, then Ed’s trusted graph includes the fact

⁴ ND is a deduction system that has many similarities to the way a person makes an actual deduction, which relates closely to a previous usage of iALC: answering questions to a Bar Exam. It is, in a way, a more readable Sequent Calculus. Another reason in its favour is the lack of support in general of inferential tools to Intuitionistic Logics.

that *org* is a non-profit organisation. Some of those sentences can be represented with role assertions, creating the related roles:

fin *auditedBy* *aud*
aud *certifiedBy* *cert*
org *hasFinRecord* *fin*
of *officerInOrg* *org*

These formulas, however, do not fit well into our Natural Deduction system, since they involve information pertinent to the ABOX (assertion) part of iALC, and the ND system reasons over the TBOX (terminological) part of it. Luckily, they are equivalent to a more TBOX-friendly format, in the form of Restrictions of Concepts (in our case, existential restrictions).⁵ We have not yet seen the case in which a universal restriction was necessary, but their modelling would be similar and would chain just as naturally as the existential ones (and iALC supports it if needed). This results in:

$fin : \exists auditedBy.Organisation^L$
 $aud : \exists certifiedBy.CertAgency^L$
 $org : \exists hasFinancialRecord.FinRecord^L$
 $of : \exists officerInOrg.Organisation^L$

The formulas are given generic labels (namely L , representing a list of labels) in order to give them flexibility when used in proofs, because, in those, they need to be able to represent any context in which they appear. Once context is given, L is substituted by an adequate list of labels. This generic list of labels can be interpreted as the formula being valid regardless of the context in which it is inserted. The supposition that financial record *fin* is audited by an organisation, i.e. $\exists auditedBy.Organisation$, can be used in the situation presented, or in any other. If we use fixed labels for this context, the isolated sentence may lose meaning in a different context (maybe one without a certification agency involved, for instance). After the proof, these concrete contexts will be explained more clearly. By chaining some of those in order to relate these different individuals we have:

$fin : \exists auditedBy.\exists certifiedBy.CertAgency^L$

The formula above represents a refinement in $fin : \exists auditedBy.Organisation^L$ via $aud : \exists certifiedBy.CertAgency^L$, which is an existential restriction on *CertAgency*, a subconcept of *Organisation*.

$org : \exists hasFinRecord.\exists auditedBy.\exists certifiedBy.Organisation^L$

This last formula can be read as *org* has a (certain) financial record which is audited by an organisation certified by a certification agency, which is itself an organisation. Representing the entire rule, we have:

$\exists hasFinRecord.\exists auditedBy.\exists certifiedBy.Organisation \sqsubseteq NonProfit^L$

Showing that, for Ed's graph, a company that fits into these criteria must be a non-profit organisation. So, here is the proof for the first rule in our ND system (names were shortened to fit the page):

⁵ The choice here is between universal and existential restrictions, but universal restrictions would imply that, for instance, the organisation *org* would have every financial record (in the world) related to it, which is very rarely the case.

$$\begin{array}{c}
 \frac{cert : Acc \quad Acc \sqsubseteq CA}{cert : CA} \sqsubseteq -e \quad CA \sqsubseteq O^{\exists h, \exists a, \exists c} \sqsubseteq -e \\
 \frac{\frac{\frac{cert : O^{\exists h, \exists a, \exists c}}{aud : \exists c. O^{\exists h, \exists a}} \exists -i}{fin : \exists a. \exists c. O^{\exists h}} \exists -i}{org : \exists h. \exists a. \exists c. O} \exists -i \quad \frac{\exists h. \exists a. \exists c. O \sqsubseteq NP}{org : NP} \sqsubseteq -e
 \end{array}$$

In this proof, the generic label L is substituted by the concrete labels, giving the context needed to the formulas. For instance, $fin : \exists a. \exists c. O^{\exists h}$ has the label $\exists h$, indicating that the restricted concept $\exists a. \exists c. O$ can be even further restricted via the role h (*hasFinancialRecord*), which connects fin to another individual, namely, org . Whereas $org : \exists h. \exists a. \exists c. O$ has no labels, because they are not necessary. In the case of this formula, everything is explicit in the restrictions themselves.

It is worth noting that, in our situation, of can be replaced by *george*, org by *relief*, and $cert$, by *audit*, showing that ed is aware that *relief* is a non-profit organisation.

4.3 Rule Ed2 (Ed does not trust George)

To capture the fact that Ed does not trust George (in the sense of donating to an organisation where George is an officer), we include the following rule. If Ed's trusted graph includes the fact that org is a non-profit organisation, George is officer of org , and Ed intends to take action a of type *Donate* to org , then a is *Denied*.

We can set the formulas to:

$george : \exists officerInOrg. Org^L$

$org : NonProfit^L$

$act : Donate^L$

$act : \exists directedTo. Org^L$

$ed : \exists intends. Donate^L$

(alternatively) $ed : \exists intends. \exists directedTo. Org^L$

We want to get to $act : Denied^L$.

However, before going on to the proof, we need a formalisation of the fact that Ed does not trust George as an officer in an organisation, and that it nullifies his action:

$ed \text{ notTrust } george$, which becomes $ed : \exists notTrust. Person^L$, and, after another transformation:

$ed : \exists notTrust. \exists officerInOrg. Org^L$

We have just inserted a new concept, *Person*, not necessarily present in the context, which was necessary in order to categorise of what concept $george$ is a part. An analogous concept would be *Officer*, but the name would be solely to give semantic information to the reader. To the situation itself the semantics of the concepts are given extensionally.

We also need to account for the fact that the action is directed towards a non-profit organisation:

$(\exists intends. \exists directedTo. NonProfit \sqcap \exists notTrust. \exists officerInOrg)$

$\sqsubseteq \exists intendsTo. Denied^L$

Then, we have the following proof:

$$\frac{\frac{\frac{org : NP^{\exists i, \exists d}}{act : \exists d.NP^{\exists i}} \exists - i}{ed : \exists i.\exists d.NP} \exists - i \quad \frac{ed : \exists nT.\exists o.O}{ed : (\exists i.\exists d.NP \cap \exists nT.\exists o.O)} \cap - i}{\frac{ed : \exists i.D}{act : D^{\exists i}} \exists - e} (\exists i.\exists d.NP \cap \exists nT.\exists o.O) \sqsubseteq \exists i.D^L \sqsubseteq -e$$

4.4 Rule George1 (PRIVACY)

Now, let us focus on the graph of George, the officer, who has a rule that does not allow divulging (reading) any information about him, including his association with an organisation. This rule, which represents his right to *privacy*, and rule 4.5, which represents the *transparency* of a non-profit organisation are apparently in conflict. In the next sections we explain how this conflict is solved. The proofs for the rules of his graph will be similar to those in the graph of Ed.

If *george* is an officer of organisation *org*, action *a* is of type *read*, and Agent *ag* intends *a* Then *a* is *Denied*.

$$\begin{aligned} org &: Org^L \\ george & \text{ officerIn } org, \text{ aka } george : \exists ofIn.Org^L \\ a &: Read^L \\ ag & \text{ intends } a, \text{ aka } ag : \exists intends.Read^L \\ a & \text{ directedTo } george, \text{ aka } a : \exists directedTo.\exists ofIn.Org^L \end{aligned}$$

We also have, by the privacy rule, that $(\exists intends.\exists directedTo.\exists ofIn.Org \cap \exists intends.Read) \sqsubseteq \exists intends.Denied$.

We want to arrive at $a : Denied^L$.

The proof is as follows:

$$\frac{\frac{\frac{org : O^{\exists in, \exists diTo, \exists ofIn}}{g : \exists ofIn.O^{\exists in, \exists diTo}} \exists - i}{a : \exists diTo.\exists ofIn.O^{\exists in}} \exists - i \quad \frac{a : R^{\exists in}}{ag : \exists in.R} \exists - i}{\frac{ag : (\exists in.\exists diTo.\exists ofIn.O \cap \exists in.R)}{ag : \exists in.D} \exists - e} (\exists in.\exists diTo.\exists ofIn.O \cap \exists in.R) \sqsubseteq \exists in.D \sqsubseteq -e$$

To contextualise, *org* is the organisation *relief*, and the agent who intends to read information on ReliefOrg, *ag*, is *ed*.

4.5 Rule Transp (TRANSPARENCY)

The rule for the Transparency ruling over our agent George's privacy rule is: if an organisation *org* is of type *NonProfit* and has officer *of*, action *a* is of type *read* and an agent *ag* intends *a*, then *A* is *Allowed*.

There are also two additional pieces of information to the graph: the primary source to the transparency rule is Non-profit Act, and Non-profit Act is a Law.

$$org : NonProfit^L$$

off $officerIn$ org , which becomes $off : \exists officerIn.NonProfit^L$
 $act : Read^L$
 act $directedTo$ org , which becomes $act : \exists directedTo.NonProfit^L$
 ag $intends$ act , which becomes $ag : \exists intends.Read^L$, and later
 $ag : \exists intends.\exists directedTo.NonProfit^L$

By transparency, $(\exists intends.\exists directedTo.\exists officerIn.NonP \sqcap \exists intends.Read) \sqsubseteq \exists intends.Allowed$. We also may represent $l : NonProfitAct$ as the legal individual for the Non-profit Act. With that, we should have $act : Allowed^L$.

This will not be used in the proof, but it is interesting to state the disjunction between the concepts *Denied* and *Allowed*: $Denied \sqsubseteq \neg Allowed$.

$$\frac{\frac{\frac{org : NP^{\exists in, \exists diTo, \exists ofIn}}{ofr : \exists ofIn.NP^{\exists in, \exists diTo}} \exists -i}{a : \exists diTo.\exists ofIn.NP^{\exists in}} \exists -i}{ag : \exists in.\exists diTo.\exists ofIn.NP} \exists -i \quad \frac{a : R^{\exists in}}{ag : \exists in.R} \exists -i}{ag : (\exists in.\exists diTo.\exists ofIn.NP \sqcap \exists in.R)} \sqcap -i \quad \frac{(\exists in.\exists diTo.\exists ofIn.NP \sqcap \exists in.R) \sqsubseteq \exists in.A}{ag : \exists in.A} \sqsubseteq -e}{a : A^{\exists in}} \exists -e$$

Also, we have a secondary proof to show that the transparency rule (represented by the individual rtr) is preceded by the law l .

$$\frac{rtr : (\exists in.\exists diTo.\exists ofIn.NP \sqcap \exists in.R) \sqsubseteq \exists in.A \quad l \preceq rtr}{l : (\exists in.\exists diTo.\exists ofIn.NP \sqcap \exists in.R) \sqsubseteq \exists in.A} \preceq -elim$$

This Natural Deduction rule, $\preceq -elim$, is a way to introduce the mechanism of precedence between individual laws directly into the Natural Deduction system. The smaller premise, $l \preceq rtr$, indicates the precedence among the individuals concerned. The bigger premise taking to the conclusion tells us that, since a concept is valid for an individual who is preceded by another, the one who precedes (in our case, l) must also validate such concept, due to the intuitionistic nature of the logic, that requires the construction of the validity of a concept to be passed through (in this case, to have been passed through) to following individuals hierarchically.

At this point, we still need to relate rules 4.4 and 4.5, which could pose a problem since they apparently contradict each other. The following rule, then, comes to the rescue.

4.6 MetaRule1 (Precedence)

This rule simply states that Laws precede personal privacy rules from the graphs of agents, such as rule 4.4.

This is simply solved by iALC with $LegalRule \sqsubseteq PersonalPrivRule^L$.

We can, then, conclude that rule 4.5 has precedence over rule 4.4, since the law ruling over the transparency of non-profit organisations has precedence over George's right to privacy (at least when regarding his work at such an organisation).

5 Final Remarks

In this article, we presented the concepts of KGs and showed how they can be used in contexts that involve the concepts of Trust, Privacy, and Transparency, as well as an alternate formalisation for them in iALC, an Intuitionistic Description Logic, to show one way one can reason over them. This modelling and reasoning was proven successful on the context shown and conveyed how we can translate the complexities of KGs via the intuitionistic notion of truth that iALC has.

For future work, we intend to extend this representation in iALC in this context of Trust, Privacy, and Transparency, specially since a lot of the decision making in terms of reasoning comes from the relations between laws that seem conflicting at first glance, but that are solved by mechanisms iALC is equipped to model.

In regards to the more theoretical side of the work, we aim to give a more formal introduction to the Natural Deduction system for iALC, via showing its properties of soundness and completeness, for instance, which are out of the scope of this article, and possibly implement an automated reasoner to it for it to be used in a larger scale.

References

1. Bernardo Alkmim, Edward Hermann Haeusler, and Alexandre Rademaker. Utilizing ialc to formalize the brazilian oab exam. In *Proceedings from the EXplainable AI in Law Workshop (XAILA)*. CEUR-WS, 2018. <http://ceur-ws.org/Vol-2381/>.
2. Franz Baader, Diego Calvanese, Deborah Mcguinness, Daniele Nardi, and Peter Patel-Schneider. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press; 2nd Edition, 01 2007.
3. N. Bobbio and M.C.C.L. dos Santos. *Theory of Legal Order*. Ed. Universidade de Brasília, 1995.
4. J.A. Bondy and U.S.R Murty. *Graph Theory*. Springer Publishing Company, Incorporated, 1st edition, 2008.
5. Bruce G. Buchanan and Edward A. Feigenbaum. Dendral and meta-dendral: Their applications dimension. *Artificial Intelligence*, 11(1):5 – 24, 1978. Applications to the Sciences and Medicine.
6. Cleo Condoravdi, Dick Crouch, Valeria de Paiva, Reinhard Stolle, and Daniel G. Bobrow. Entailment, intensionality and text understanding. In *Proceedings of the HLT-NAACL 2003 Workshop on Text Meaning - Volume 9, HLT-NAACL-TEXTMEANING '03*, pages 38–45, Stroudsburg, PA, USA, 2003. Association for Computational Linguistics.
7. Pedro Delfino, Bruno Cuconato, Edward Hermann Haeusler, and Alexandre Rademaker. Passing the brazilian oab exam: Data. *Legal Knowledge and Information Systems*, page 89, 2017.
8. Paul Groth, Andrew Gibson, and Johannes Velterop. The anatomy of a nano-publication. *Information Services and Use*, 30, 01 2010.
9. Edward Hermann Haeusler and Alexandre Rademaker. On how kelsenian jurisprudence and intuitionistic logic help to avoid contrary-to-duty paradoxes in legal ontologies. *Proc. Journal of Applied Non-Classical Logics*, 2016.
10. Frederick Hayes-Roth, Donald A. Waterman, and Douglas B. Lenat. *Building Expert Systems*. Addison-Wesley Longman Publishing Co., Inc., USA, 1983.
11. A. Heyting. *Intuitionism: An Introduction*. North-Holland Publishing, Amsterdam, Netherlands, 3rd edition, 1956.

12. Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d'Amato, Gerard de Melo, Claudio Gutierrez, José Emilio Labra Gayo, Sabrina Kirrane, Sebastian Neumaier, Axel Polleres, Roberto Navigli, Axel-Cyrille Ngonga Ngomo, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge graphs, 2020.
13. Matthias Jarke, Bernd Neumann, Yannis Vassiliou, and Wolfgang Wahlster. *KBMS Requirements of Knowledge-Based Systems*, page 381–394. Springer-Verlag, Berlin, Heidelberg, 1989.
14. H. Kelsen. *Pure Theory of Law*. Lawbook Exchange, 2005.
15. Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104, 12 2015.
16. E. W. Schneider and VA. Human Resources Research Organization, Alexandria. *Course Modularization Applied [microform] : The Interface System and Its Implications For Sequence Control and Data Analysis / E. W. Schneider*. Distributed by ERIC Clearinghouse [Washington, D.C.], 1973.
17. Daniel Schwabe and Carlos Laufer. Trust and privacy in knowledge graphs. *CoRR*, abs/1903.07673, 2019.
18. Stephan Seufert, Patrick Ernst, Srikanta J. Bedathur, Sarath Kumar Kondreddi, Klaus Berberich, and Gerhard Weikum. Instant espresso: Interactive analysis of relationships in knowledge graphs. In Jacqueline Bourdeau, Jim Hendler, Roger Nkambou, Ian Horrocks, and Ben Y. Zhao, editors, *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11-15, 2016, Companion Volume*, pages 251–254. ACM, 2016.
19. Andre Valente and Joost Breuker. A model-based approach to legal knowledge engineering. *JURIX*, 1992, 01 1992.

A Natural Deduction System for iALC

Let α and β concepts, x and y nominals, δ formula, and R a role. L represents a list of labels (possibly empty). Labels represent either negations of concepts, or universal or existential restrictions on concepts, made *implicit*. They indicate a sort of *context* to the concept they are attached to. L^\forall is a list that restricts all labels in itself to $\forall R$ of some kind, and L^\exists restricts all to $\exists R$. Rules that utilise nominals and involve x and y assume that there is a role connecting them, i.e. xRy . Please note that only the rules related to iALC operators utilised in this article are shown. The n in the names of some rules indicate that those are the versions which utilise nominals.

$$\frac{y : \alpha^{\exists R, L}}{x : (\exists R. \alpha)^L} \exists - \text{intro} \qquad \frac{x : (\exists R. \alpha)^L}{y : \alpha^{\exists R, L}} \exists - \text{elim}$$

$$\frac{\begin{array}{c} [x : \alpha^{L1}] \\ \vdots \\ x : \beta^{L2} \end{array}}{x : (\alpha^{L1} \sqsubseteq \beta^{L2})} \sqsubseteq - \text{intro} - n \qquad \frac{x : \alpha^{L1} \quad x : (\alpha^{L1} \sqsubseteq \beta^{L2})}{x : \beta^{L2}} \sqsubseteq - \text{elim} - n$$

$$\frac{\begin{array}{c} [x : \alpha^L] \\ \vdots \\ x : \perp \end{array}}{x : (\neg\alpha)^{-L}} \neg\text{-intro} - n \qquad \frac{x : \alpha^L \quad x : (\neg\alpha)^{-L}}{x : \perp} \perp\text{-intro} - n$$

$$\frac{x : \alpha^{L^\vee} \quad x : \beta^{L^\vee}}{x : (\alpha \sqcap \beta)^{L^\vee}} \sqcap\text{-intro} - n \qquad \frac{x : (\alpha \sqcap \beta)^{L^\vee}}{x : \alpha^{L^\vee}} \sqcap\text{-elim}_1 - n$$

$$\frac{x : \alpha^L \quad y \preceq x}{y : \alpha^L} \preceq\text{-elim} \qquad \frac{x : (\alpha \sqcap \beta)^{L^\vee}}{x : \beta^{L^\vee}} \sqcap\text{-elim}_2 - n$$