

Jeferson de Oliveira Batista

**Ontologically Correct Taxonomies by
Construction: A Graph Grammar-Based
Approach**

Vitória, ES

2022

Jeferson de Oliveira Batista

Ontologically Correct Taxonomies by Construction: A Graph Grammar-Based Approach

Master's dissertation presented to the Computer Science Graduate Program of the Federal University of Espírito Santo, as partial requirement to obtain the Computer Science Master's Degree.

Federal University of Espírito Santo – UFES

Technology Center

Computer Science Graduate Program

Supervisor: Prof. Dr. João Paulo A. Almeida

Co-supervisor: Prof. Dr. Eduardo Zambon

Vitória, ES

2022

Ficha catalográfica disponibilizada pelo Sistema Integrado de Bibliotecas - SIBI/UFES e elaborada pelo autor

B333o Batista, Jeferson de Oliveira, 1995-
Ontologically Correct Taxonomies by Construction : A
Graph Grammar-Based Approach / Jeferson de Oliveira Batista.
2022.
81 f. : il.

Orientador: João Paulo Andrade Almeida.

Coorientador: Eduardo Zambon.

Dissertação (Mestrado em Informática) - Universidade Federal do Espírito Santo, Centro Tecnológico.

1. Classificação. 2. Ontologia. 3. Gramática de grafos. I. Almeida, João Paulo Andrade. II. Zambon, Eduardo. III. Universidade Federal do Espírito Santo. Centro Tecnológico. IV. Título.

CDU: 004

Jeferson de Oliveira Batista

Ontologically Correct Taxonomies by Construction: A Graph Grammar-Based Approach

Master's dissertation presented to the Computer Science Graduate Program of the Federal University of Espírito Santo, as partial requirement to obtain the Computer Science Master's Degree.

Work approved. Vitória, ES, March 4, 2022:

Prof. Dr. João Paulo A. Almeida
Orientador

Prof. Dr. Eduardo Zambon
Co-orientador

Prof. Dr. Vítor Estêvão Silva Souza
Membro Interno

Prof. Dr. Tiago Prince Sales
Membro Externo

Vitória, ES
2022

Acknowledgements

First of all, I am grateful to God, from whom all good derives.

I am also grateful to my family, which supported me for all these years. To my father, Ailton, for his incentive, our conversations and to be an example for me. To my brother, Jackson, for his advices and affection. And, *in memoriam*, to my mother, Eliane, for all her love, dedication and presence in my hard moments. I am also grateful to my extended family for their love and affection for all these years.

I am thankful, *in memoriam*, to Prof. Ricardo Falbo for he was my first supervisor, provide the first ideas for this work and guided me in the beginning of this journey. I am also thankful to Prof. João Paulo A. Almeida who continued my supervision, for his insights and patience along the path. To my co-supervisor Prof. Eduardo Zambon, I am thankful for his contributions and support along the way. I am also grateful to Prof. Giancarlo Guizzardi, whose work is so important for this dissertation.

I am thankful to my colleagues of NEMO for their companionship and the ideas we shared. I am thankful to my professors for their contribution to the person I am now, and for all the work of the staff of the Federal University of Espírito Santo. To my classmates and other colleagues that I had the pleasure to know and having in my companionship.

Finally, I am grateful to my siblings in Christ that congregate with me, for all their example, words and encouragement along the last years.

Resumo

Taxonomias exercem um papel central em modelagem conceitual de domínio, tendo um impacto direto tanto na Ciência da Computação (em áreas como Representação do Conhecimento, Engenharia de Ontologias e Engenharia de Software) quanto na Ciência da Informação. Apesar disso, há pouca orientação sobre como criar taxonomias de alta qualidade, sendo exceções notáveis a metodologia OntoClean e a linguagem de modelagem conceitual orientada a ontologia OntoUML. Essas técnicas levam em consideração as meta-propriedades ontológicas de rigidez e sortalidade de tipos para estabelecer regras bem fundamentadas sobre a formação de estruturas taxonômicas. A meta-propriedade de rigidez define se um tipo se aplica essencialmente ou contingentemente a suas instâncias, enquanto a sortalidade define se um tipo provê um princípio de identidade uniforme para suas instâncias. Nesta dissertação, mostramos como utilizar as regras formais subjacentes a essas técnicas para construir taxonomias que são corretas por construção. Nós definimos um conjunto de operações que preservam a corretude para sistematicamente introduzir tipos e relações de especialização em estruturas taxonômicas. Além de considerar a micro-teoria ontológica dos tipos subjacente a OntoClean e OntoUML, nós também aplicamos a micro-teoria MLT (Teoria de Multi-Níveis) de tipos de alta ordem que nos permite lidar com taxonomias multi-nível baseadas no *powertype pattern*, nas quais uma entidade pode ser um tipo e uma instância ao mesmo tempo. Para validar nossa proposta, nós formalizamos as operações de construção de modelos como uma gramática de grafos que incorpora ambas as micro-teorias. Uma gramática de grafos é um modo formal de se especificar um grafo inicial e um conjunto de regras de transformação de grafos. Cada grafo representa um modelo, ou seja, uma taxonomia. Uma regra de transformação consiste em precondições que devem ser verdadeiras para um modelo para que a regra seja aplicável, e um conjunto de operações de criação e deleção para vértices e arestas. O conjunto de modelos alcançáveis pela aplicação das regras da gramática é chamado de linguagem da gramática. Aplicamos técnicas de verificação automática sobre a linguagem da gramática para mostrar que a gramática de grafos é correta, ou seja, que todas as taxonomias produzidas pelas regras da gramática são corretas, pelo menos até certo tamanho. Também mostramos que as regras podem gerar todas as taxonomias corretas até certo tamanho (um resultado de completude).

Palavras-chave: taxonomias · modelagem conceitual · ontologias · gramáticas de grafo · corretude por construção

Abstract

Taxonomies play a central role in conceptual domain modeling, having a direct impact in areas such as knowledge representation, ontology engineering, and software engineering, as well as knowledge organization in information sciences. Despite this, there is little guidance on how to build high-quality taxonomies, with notable exceptions being the OntoClean methodology, and the ontology-driven conceptual modeling language OntoUML. These techniques take into account the ontological meta-properties of rigidity and sortality of types to establish well-founded rules on the formation of taxonomic structures. The rigidity meta-property defines whether a type applies essentially or contingently to its instances, while the sortality defines whether a type provides a uniform principle of identity for its instances. In this dissertation, we show how to leverage the formal rules underlying these techniques in order to build taxonomies which are *correct by construction*. We define a set of correctness-preserving operations to systematically introduce types and subtyping relations into taxonomic structures. In addition to considering the ontological micro-theory of enduring types underlying OntoClean and OntoUML, we also employ the MLT (Multi-Level Theory) micro-theory of high-order types, which allows us to address multi-level taxonomies based on the *powertype* pattern, in which an entity can be both a type and an instance at the same time. To validate our proposal, we formalize the model building operations as a graph grammar that incorporates both micro-theories. A graph grammar is a formal way to specify an initial graph and a set of graph transformation rules. Each graph represents a model, in our case, a taxonomy. A transformation rule consists of preconditions that must be true for a model in order to the rule be applicable, and a set of creation and deletion operations for vertices and edges. The set of models reachable applying the grammar rules is called the grammar language. We apply automatic verification techniques over the grammar language to show that the graph grammar is *sound*, i.e., that all taxonomies produced by the grammar rules are correct, at least up to a certain size. We also show that the rules can generate all correct taxonomies up to a certain size (a *completeness* result).

Keywords: taxonomies · conceptual modeling · ontologies · graph grammars · correctness by construction

List of Figures

Figure 1	– Example of GROOVE transformation rule.	18
Figure 2	– An application of the example rule depicted in Figure 1.	19
Figure 3	– Example of a transformation rule with a forbidden absence.	19
Figure 4	– Graphs to illustrate the forbidden absence applicability.	20
Figure 5	– Example of GROOVE type graph.	20
Figure 6	– An example of sortal hierarchy for the kind ‘Person’.	22
Figure 7	– Non-Sortal examples.	23
Figure 8	– A taxonomy for Endurant Types.	24
Figure 9	– Incorrect representation of type ‘Customer’.	25
Figure 10	– An example of multi-level taxonomy in biological domain.	26
Figure 11	– A multi-level model for the biological domain including variants of the powertype pattern.	28
Figure 12	– An incorrect multi-level model present in Wikidata and uncovered in (DADALTO et al., 2021).	29
Figure 13	– Type graph established to reflect the taxonomy for Endurant Types.	32
Figure 14	– Transformation rules to introduce an independent type.	33
Figure 15	– Transformation rule to introduce a SUBKIND type.	33
Figure 16	– Transformation rule to introduce an ANTI-RIGID SORTAL type.	33
Figure 17	– Transformation rules to specialize a CATEGORY or a MIXIN.	34
Figure 18	– Transformation rule to specialize an ANTI-RIGID NON-SORTAL type.	34
Figure 19	– Transformation rules to generalize a SUBKIND or an ANTI-RIGID SORTAL.	35
Figure 20	– Restrictive condition of a KIND specializing another SORTAL.	36
Figure 21	– Restrictive condition of a SORTAL with more than one KIND.	37
Figure 22	– Restrictive condition of a rigid or semi-rigid type specializing an anti-rigid one.	37
Figure 23	– Restrictive condition of a NON-SORTAL type specializing a SORTAL one.	37
Figure 24	– Restrictive condition of a SORTAL without a KIND.	37
Figure 25	– Initial host graph only with one basic type representing the INDIVIDUAL class.	42
Figure 26	– Type graph established to reflect the MLT relations.	42
Figure 27	– Transformation rule to introduce a FIRST-ORDER TYPE.	43
Figure 28	– Transformation rule to introduce a new order.	43
Figure 29	– Transformation rule to introduce a new high-order type.	43
Figure 30	– Transformation rule to introduce a new specialization relation.	44
Figure 31	– Transformation rule to introduce a new instantiation relation.	44
Figure 32	– Restrictive condition of a basic type with more than one powertype.	45
Figure 33	– Restrictive condition of different basic types with the same powertype.	45

Figure 34 – Restrictive condition of unrelated basic types.	46
Figure 35 – Restrictive condition of mutual powertype relation between basic types. . .	46
Figure 36 – Restrictive condition of a type with no basic type as supertype.	46
Figure 37 – Restrictive condition of a type specializing more than one basic type. . . .	47
Figure 38 – Restrictive conditions of categorization between nonadjacent levels.	47
Figure 40 – Restrictive condition of specialization between different levels.	48
Figure 41 – Restrictive condition of a categorizer instance not proper specializing the categorized type.	48
Figure 39 – Restrictive conditions of instantiation between nonadjacent levels.	48
Figure 42 – Restrictive condition of an instance violating the definition of specialization.	49
Figure 43 – Type graph established to reflect the taxonomy for Endurant Types.	54
Figure 44 – Type graph established to reflect the MLT relations.	54
Figure 45 – Transformation rules to introduce a new first-order independent type.	54
Figure 46 – Transformation rules to introduce a new first-order dependent type.	55
Figure 47 – Transformation rule to introduce a new order.	55
Figure 48 – Transformation rules to introduce a new high-order independent type. . . .	56
Figure 49 – Transformation rules to introduce a new high-order dependent type.	56
Figure 50 – Transformation rules to specialize a CATEGORY or a MIXIN.	57
Figure 51 – Transformation rule to specialize an ANTI-RIGID NON-SORTAL.	57
Figure 52 – Transformation rules to generalize a SORTAL.	57
Figure 53 – Transformation rule to classify an ENDURANT TYPE.	58
Figure 54 – Restrictive condition of a type instantiating multiple KINDS.	59

List of Tables

Table 1	– Results of soundness analysis for the endurant types graph grammar.	38
Table 2	– Results of completeness analysis for the endurant types graph grammar. . .	39
Table 3	– Results of completeness analysis for the endurant types graph grammar that can only violate rigidity.	40
Table 4	– Results of completeness analysis for the endurant types graph grammar that can only violate sortality.	40
Table 5	– Results of completeness analysis for the endurant types graph grammar that can only violate KIND constraints.	40
Table 6	– Results of soundness analysis for the MLT-based graph grammar.	49
Table 7	– Results of completeness analysis for the MLT-based graph grammar.	51
Table 8	– Results of completeness analysis for the MLT-based graph grammar that can only violate the stratification by categorization relations.	51
Table 9	– Results of completeness analysis for the MLT-based graph grammar that violates the stratification by specialization relations.	52
Table 10	– Results of completeness analysis for the MLT-based graph grammar with in- stantiation relations violating the categorization and specialization definitions.	52
Table 11	– Results of soundness analysis for the combined grammar.	59

List of abbreviations and acronyms

UML	Unified Modeling Language
UFO	Unified Foundational Ontology
gUFO	gentle Unified Foundational Ontology
MLT	Multi-Level Theory

Contents

1	INTRODUCTION	13
1.1	Context and Motivation	13
1.2	Contribution	14
1.3	Approach	15
1.4	Structure	16
2	BACKGROUND	17
2.1	Graph Grammars	17
2.1.1	Graphs	17
2.1.2	The GROOVE Graph Transformation Tool	18
2.2	Ontological Foundations	20
2.3	The Multi-Level Theory	25
2.4	Related Works	29
3	TAXONOMIES OF ENDURANT TYPES	32
3.1	Graph Transformation Rules for Ontologically Correct Taxonomies	32
3.1.1	Introducing New Types	32
3.1.2	Introducing Dependent Types	33
3.1.3	Introducing Specializations for Existing Non-Sortal Types	34
3.1.4	Introducing Generalizations for Existing Sortal Types	34
3.1.5	Summary of the Grammar Rules	35
3.2	Formal Verification	36
3.2.1	Verifying Soundness	37
3.2.2	Verifying Completeness	38
3.2.3	Verification Scope Matters	41
4	MULTI-LEVEL TAXONOMIES	42
4.1	Graph Transformation Rules for Correct Multi-Level Taxonomies	42
4.1.1	Introducing a New First-Order Type	42
4.1.2	Introducing a New Order	43
4.1.3	Introducing a New High-Order Type	43
4.1.4	Introducing Specialization Relations	44
4.1.5	Introducing Instantiation Relations	44
4.1.6	Summary of the Grammar Rules	44
4.2	Formal Verification of MLT Constraints	45
4.2.1	Constraints Related to Basic Types	45

4.2.2	Constraints Related to Level Stratification	46
4.2.3	Constraints Related to Categorization and Specialization	48
4.2.4	Verifying Soundness	49
4.2.5	Verifying Completeness	50
5	JOINING THE FOUNDATIONAL THEORIES	53
5.1	Graph Transformation Rules to Build Ontologically Correct Multi-Level Taxonomies	53
5.1.1	Introducing First-Order New Types	54
5.1.2	Introducing First-Order Dependent Types	54
5.1.3	Introducing a New Order	55
5.1.4	Introducing High-Order New Types	55
5.1.5	Introducing High-Order Dependent Types	56
5.1.6	Introducing Specializations for Existing Non-Sortal Types	56
5.1.7	Introducing Generalizations for Existing Sortal Types	57
5.1.8	Classifying Types	57
5.1.9	Summary of the Grammar Rules	58
5.2	Formal Verification	59
6	FINAL CONSIDERATIONS	60
6.1	Limitations	61
6.2	Future Work	62
	BIBLIOGRAPHY	63
	APPENDICES	67

1 Introduction

In this chapter, we present an overview of this work, briefly presenting taxonomies, their importance, and the gap in representation systems to guide the construction of high-quality taxonomies. Then, we present some representation systems that address the issues of taxonomy quality partially, motivating this work by highlighting their current limitations. Finally, we describe our contribution, showing how taxonomies can be thought as the application of certain *modeling patterns* captured in graph grammars, and present the structure of this work.

1.1 Context and Motivation

Taxonomies are structures connecting types via *subtyping*, i.e., type specialization relations. These structures are fundamental for conceptual domain modeling, and have a central organizing role in areas such as knowledge representation, ontology engineering, object-oriented modeling, as well as in knowledge organization in information sciences (e.g., in the construction of vocabularies and other lexical resources). Despite taxonomies' key role in all these areas, there is little guidance in the literature on how to build high-quality taxonomies, i.e., taxonomies that take into account ontological distinctions.

A notable exception is OntoClean (GUARINO; WELTY, 2004), a pioneering methodology providing a number of guidelines for diagnosing and repairing taxonomic relations that are inconsistent from an ontological point of view. These guidelines are grounded on a number of *formal meta-properties*, i.e., properties characterizing types. Derived from these meta-properties, the methodology offers a number of formal rules governing how types characterized by different meta-properties can be associated to each other in well-formed taxonomies.

OntoClean has been successfully employed to evaluate and suggest repairs to several important resources, e.g., WordNet (OLTRAMARI et al., 2002). However, being a methodology, OntoClean does not offer a representation mechanism for building taxonomies according to its prescribed rules. To address this problem, a UML profile (GUIZZARDI et al., 2004) was proposed containing modeling distinctions extending the meta-properties and rules of OntoClean. This UML profile would later become the basis of the OntoUML modeling language (GUIZZARDI, 2005), incorporating syntactic rules to prevent the construction of incorrect taxonomies in conceptual models. In (GUIZZARDI et al., 2018), the language has its full formal semantics defined in terms of a (proved-consistent) ontological theory, and its abstract syntax is defined in terms of an extension of the UML 2.0 metamodel, redesigned to reflect the ontological distinctions and axiomatizations put forth by that ontological theory.

As argued in (RUY et al., 2017), instead of leveraging on this axiomatization by proposing

methodological rules (OntoClean) or semantically motivated syntactical constraints (OntoUML), a representation system based on this ontological theory could employ a more productive strategy, leveraging on the fact that formal constraints from the theory impose a correspondence between each particular type of type (characterized by ontological meta-properties) and certain modeling *structures* (or modeling *patterns*). In other words, a representation system grounded on this ontological theory is a *pattern language*, i.e., a system whose basic building blocks are not low-granularity primitives such as types and relations, but higher-granularity patterns *formed* by types and relations.

In this work, we also employ an additional foundational theory called MLT (Multi-Level Theory) (ALMEIDA et al., 2018; CARVALHO; ALMEIDA, 2018), which allows us to address taxonomies that also include high-order types, i.e., *multi-level* taxonomies. As shown in (FONSECA et al., 2021; BRASILEIRO et al., 2016a; DADALTO et al., 2021), multi-level taxonomies are observed at scale in practice and suffer from a large number of modeling errors that can be attributed to the inadequate use of subclassing and instantiation in tandem. By incorporating the rules of MLT in a graph grammar, we can prevent such errors from occurring.

1.2 Contribution

This work contributes to the foundations of rigorous conceptual modeling by identifying the set of rules that should be considered as primitives in the design of correct taxonomies (including multi-level ones) and formalizing it as a graph grammar that, starting from an initial taxonomy, allows us to build only correct taxonomies. To put this more precisely: in the area of formal verification, statements about a system are usually split between *soundness* and *completeness* properties. The soundness of a modeled system ensures that only desirable models are possible. In our setting, this means that only correct taxonomies can be part of the grammar language. On the other hand, *completeness* ensures that if a desirable system configuration can exist “in the real world”, then a corresponding model is reachable in the formalization. In our setting, this means that any correct taxonomy can be created using the proposed graph grammar.

The identification of rules in the design of correct taxonomies is done in a metamodel-independent way, so the results presented here can be incorporated into different modeling languages (e.g., ORM (HALPIN; MORGAN, 2010), whose development was already influenced by the ontological distinctions underlying OntoUML) as well as different tools used by different communities (e.g., as a modeling plugin to Semantic Web tools such as Protégé¹). The quality of the set of taxonomy building rules is ensured with some activities of formal verification regarding their *soundness* and *completeness*.

¹ <<https://protege.stanford.edu/>>

1.3 Approach

The set of model building operations for the construction of correct taxonomies is formalized in this work as a graph grammar. A graph grammar is a formal way to specify an initial graph and a set of graph transformation rules. Each graph represents a model, in our case, a taxonomy. A transformation rule consists of preconditions that must be true for a model in order to the rule be applicable, and a set of creation and deletion operations for vertices and edges. The set of models reachable applying the grammar rules is called the grammar language.

In (ZAMBON; GUIZZARDI, 2017), the authors used the GROOVE graph transformation tool (GHAMARIAN et al., 2012) to make a first attempt to formalize a representation system grounded on the underlying theory of enduring types from OntoUML as a true *pattern grammar*, i.e., as a graph grammar that reflects this theory, with the grammar capturing the language patterns and their possible relations as graph transformation rules.

Later, in (BATISTA et al., 2021), we used GROOVE to define a graph grammar over a smaller scope of the theory of enduring types from OntoUML. By employing the state space exploration mechanism supported by GROOVE, we managed to detect important omissions in the rule set proposed in (ZAMBON; GUIZZARDI, 2017) and, inspired by that work, to gradually define our own graph grammar.

Being able to employ automatic verification tools and techniques over the grammar language is one major advantage of capturing ontological patterns as a graph grammar (BATISTA et al., 2021). We were able to show that the graph grammar presented in (BATISTA et al., 2021) is *sound* and *complete* up to a certain size, i.e., that all models produced by the grammar rules are correct (*soundness*); and that the rules can generate all correct models (*completeness*). Such formal graph transformation rules thus comprise a framework that allows the user to build models which are *correct by construction*. This framework could be incorporated in a graphical tool that, by selecting a transformation rule, the entities in which the selected rule could be applied would be highlighted; and that, by selecting an entity, the applicable rules would be highlighted. Such a graphical tool could be used to prototype taxonomies or participate in the construction of conceptual models involving concepts beyond the scope of the present work.

Here, we expand the work started in (BATISTA et al., 2021) with considerations about our verification scope, and also employing a foundational theory called MLT (Multi-Level Theory) (ALMEIDA et al., 2018; CARVALHO; ALMEIDA, 2018), which allows us to address taxonomies that also include high-order types, i.e., *multi-level* taxonomies, incorporating the rules of MLT in a graph grammar. This new MLT-based grammar is then combined with the original OntoUML graph grammar from (BATISTA et al., 2021), to provide a comprehensive pattern grammar for multi-level ontology-based conceptual modeling.

1.4 Structure

The remainder of this dissertation is structured as follows. In Chapter 2, we present the required theoretical foundations and discuss related work. In Chapter 3, we present and assess the first graph grammar, which corresponds to the theory of endurant types underlying OntoClean and OntoUML. In Chapter 4, we address the graph grammar corresponding to the MLT Multi-Level Theory, which is again formalized and assessed. In Chapter 5, both graph grammars are combined, leading us to correct ontology-based multi-level taxonomies. Finally, Chapter 6 presents our concluding remarks.

2 Background

This chapter covers the background elements employed in this work, including graph grammars (Section 2.1), the ontological distinctions underlying the Unified Foundational Ontology (UFO) that form the first micro-theory employed (Section 2.2), the rules in the MLT Multi-Level Theory that form the second micro-theory (Section 2.3) and related work (Section 2.4).

2.1 Graph Grammars

Graph transformation (or *graph rewriting*) (HECKEL, 2006) has been advocated as a flexible formalism, suitable for modeling systems with dynamic configurations or states. This flexibility is achieved by the fact that the underlying data structure, that of graphs, is capable of capturing a broad variation of systems. Some areas where graph transformation is being applied include visual modeling of systems, the formal specification of model transformations, and the definition of graph languages, to name a few (ZAMBON, 2013; GHAMARIAN et al., 2012).

The core concept of graph transformation is the rule-based modification of graphs, where each application of a rule leads to a graph transformation step. A transformation rule specifies both the necessary preconditions for its application and the rule effect (modifications) on a *host graph*. The modified graph produced by a rule application is the result of the transformation.

In this work, we use graph transformations to formally model the operations for the construction of a taxonomy. A set of graph transformation rules can be seen as a declarative specification of how the construction can evolve from an initial state, represented by an initial host graph. This combination of a rule set plus an initial graph is called a *graph grammar*, and the (possibly infinite) set of graphs reachable from the initial graph constitute the *grammar language*.

2.1.1 Graphs

Graphs and diagrams provide a simple and powerful approach to a variety of problems that are typical to computer science in general, and software engineering in particular. In fact, for most activities in the software process, a variety of visual notations have been proposed, including state diagrams, Structured Analysis, control flow graphs, architectural description languages, function block diagrams, and the UML family of languages. These notations produce models that can be easily seen as graphs and thus graph transformations are involved,

either explicitly or behind the scenes, when specifying how these models should be built and interpreted, and how they evolve over time and are mapped to implementations. At the same time, graphs provide a universally adopted data structure, as well as a model for the typology of object-oriented, component-based and distributed systems. Computations in such systems are therefore naturally modelled as graph transformations, too (HECKEL, 2006).

Graphs provide the most basic mathematical model for entities and relations. A graph consists of a set of vertices V and a set of edges E such that each edge e in E has a source and a target vertex $s(e)$ and $t(e)$ in V , respectively (HECKEL, 2006). The vertices (or the edges) can be labeled, having a label that suggests an interpretation for the graph with respect to corresponding real-world entities.

2.1.2 The GROOVE Graph Transformation Tool

Graphs in GROOVE consist of nodes (vertices), which can be typed, and edges. An edge is a directed labeled arrow between two nodes.

Graphs are transformed by applying rules. A rule consists of the following:

- A pattern that must be present in the host graph in order for the rule to be applicable;
- Subpatterns that must be absent in the host graph in order for the rule to be applicable;
- Elements (nodes and edges) to be deleted from the graph;
- Elements (nodes and edges) to be added to the graph.

All these elements are combined into a single graph; colors and shapes are used to distinguish them. Alternatively, one may think in terms of *application conditions* and *modifications*: of the former, we distinguish positive (which must be present in order to apply a rule) and negative (which must be absent in order to apply a rule) ones, whereas of the latter, we distinguish deletion and creation of elements (GHAMARIAN et al., 2012).




Figure 1 – Example of GROOVE transformation rule.

- The black (continuous thin) “reader” elements, in Figure 1 the nodes **A** and **B** and the edge labeled ‘parent’ between them, must be present and are preserved – in fact, they form a positive application condition;

- The red (dashed thick) “embargo” elements must be absent in the graph, in Figure 1 the nodes **A** and **C** and the edges labeled ‘child’ from the node **B** to them – in fact, each connected subgraph of embargo elements forms a negative application condition;
- The blue (dashed thin) “eraser” elements, in Figure 1 the node **D** and the edge labeled ‘child’ from the node **B** to it, must be present and are deleted;
- The green (continuous thick) “creator” elements, in Figure 1 the node **C** and the edge labeled ‘child’ from the node **B** to it, are created.

For the rule depicted in Figure 1 to be applicable, we need to have a model with a node **B** with an edge labeled ‘parent’ pointing to a node **A** and an edge labeled ‘child’ pointing to a node **D**. Besides, an edge labeled ‘child’ from this node **B** to either a node **A** or a node **C** must be absent. When the rule is applied, both the edge labeled ‘child’ from node **B** to node **D** and the node **D** are deleted. And a node **C** and an edge labeled ‘child’ from the node **B** to this new node **C** is created. An application of the rule depicted in Figure 1 is illustrated in Figure 2.




Figure 2 – An application of the example rule depicted in Figure 1.

The label of a given edge in an application condition may end with one of the symbols ‘*’ and ‘+’. The symbol ‘*’ is used to represent a condition with zero or more edges with the same label, while the symbol ‘+’ is used to represent a condition with one or more edges with the same label. Finally, the symbol ‘!’ at the beginning of an edge label is used in GROOVE as a negation. If used in a positive application condition, it represents the absence of an edge with that label and its use is equivalent to the use of the same label in a negative application condition. If the symbol ‘!’ is used in a negative application condition, it represents the *forbidden absence* (not the presence) of an edge with that label. To illustrate the concept of the forbidden absence of an edge, consider the transformation rule shown in Figure 3, that is applicable in the graphs of Figures 4(a, b), but not of Figure 4(c). Note that an edge ‘children’ from the node **A** to a node **D** must be present to the absence of an edge ‘children’ from the node **B** to a node **D** be forbidden.




Figure 3 – Example of a transformation rule with a forbidden absence.




Figure 4 – Graphs to illustrate the forbidden absence applicability.

The core functionality of GROOVE is to recursively apply one of all rules from a predefined set (the graph transformation system) to a given start graph, and to all graphs generated by such applications. This results in a *state space* consisting of the generated graphs (GHAMARIAN et al., 2012). The combination of a start graph with a graph transformation system is called a *graph grammar*.

A graph grammar may include *type graphs* that determine the allowed combinations of node types and edges (GHAMARIAN et al., 2012). In these type graphs, as usual, UML arrows with a closed head connect subtypes to their supertypes (the arrowhead pointing to the supertype). The Figure 5 shows the type graph used in the grammar with the rules presented in Figures 1 and 3. The types **A**, **B**, **C** and **D** are defined as specializations of the type **Node**, and only the ‘child’ and ‘parent’ relations can be present between **Nodes**.




Figure 5 – Example of GROOVE type graph.

There are graphs in GROOVE that present the same elements of graphs representing transformation rules, but that do not have *modification* elements, only *application conditions*. It is said that these graphs represent *graph conditions*. These graph conditions, as the name points out, are used to verify whether the graphs of a given state space satisfy certain conditions or not.

2.2 Ontological Foundations

In this section, we present some ontological distinctions that are, alongside with multi-level modeling concepts, the basis for this dissertation. These notions, and the constraints governing their definitions and relations, correspond to a fragment of the foundational ontology underlying OntoUML, which incorporates and extends the theory of types underlying OntoClean (GUIZZARDI, 2005; GUIZZARDI et al., 2015). For an in-depth discussion, philo-

sophical justification, empirical support, and full formal characterization of these notions, see (GUIZZARDI, 2005; GUIZZARDI et al., 2018).

Types represent properties that are shared by a set of possible instances. The set of properties shared by those instances is termed the *intension* of a type; the set of instances that share those properties (i.e., the instances of that type) is termed the *extension* of that type. Types can change their extension across different circumstances, either because things come in and out of existence, or because things can acquire and lose some of those properties captured in the intension of that type.

Taxonomic structures capture *subtyping* relations among types, both from *intensional* and *extensional* points of view. In other words, subtyping is thus a relation between types that govern the relation between the possible instances of those types. So, if type *B* is a subtype of *A*, then we have that: (i) it is necessarily the case that all instances of *B* are instances of *A*, i.e., in all possible circumstances, the extension of *B* subsets the extension of *A*; and (ii) all properties captured by the *intension* of *A* are included in the *intension* of type *B*, i.e., *B*'s are *A*'s and, therefore, *B*'s have all properties that are properties defined for type *A*.

Suppose all instances that exist in a domain of interest are *endurants* (GUIZZARDI et al., 2018). Endurants roughly correspond to what we call *objects* in ordinary language, i.e., things that (in contrast to occurrences, events) endure in time changing their properties while maintaining their identity. Examples include you, each author of this dissertation, Mick Jagger, the Moon, the Federal University of Espírito Santo.

Every endurant in our domain belongs to one **KIND**. In other words, central to any domain of interest we will have a number of object kinds, i.e., the genuine fundamental types of objects that exist in that domain. The term “kind” is meant here in a strong technical sense, i.e., by a kind, we mean a type capturing *essential* properties of the things it classifies. In other words, the objects classified by that kind could not possibly exist without being of that specific kind (GUIZZARDI et al., 2018).

Kinds tessellate the possible space of objects in that domain, i.e., all objects belong to exactly one kind and do so necessarily. Typical examples of kinds include ‘Person’, ‘Organization’, and ‘Car’. We can, however, have other static subdivisions (or subtypes) of a kind. These are naturally termed **SUBKINDS**. As an example, the kind ‘Person’ can be specialized in the (biological) subkinds ‘Man’ and ‘Woman’.

Endurant kinds and subkinds represent essential properties of objects. They are examples of **RIGID TYPES** (GUIZZARDI et al., 2018). Rigid types are those types that classify their instances necessarily, i.e., their instances must instantiate them in every possible circumstance in which they exist. We have, however, types that represent *contingent* or *accidental* properties of endurants termed **ANTI-RIGID TYPES** (GUIZZARDI et al., 2018). For example, in the way that ‘being a living person’ captures a cluster of contingent properties of a person, that ‘being a

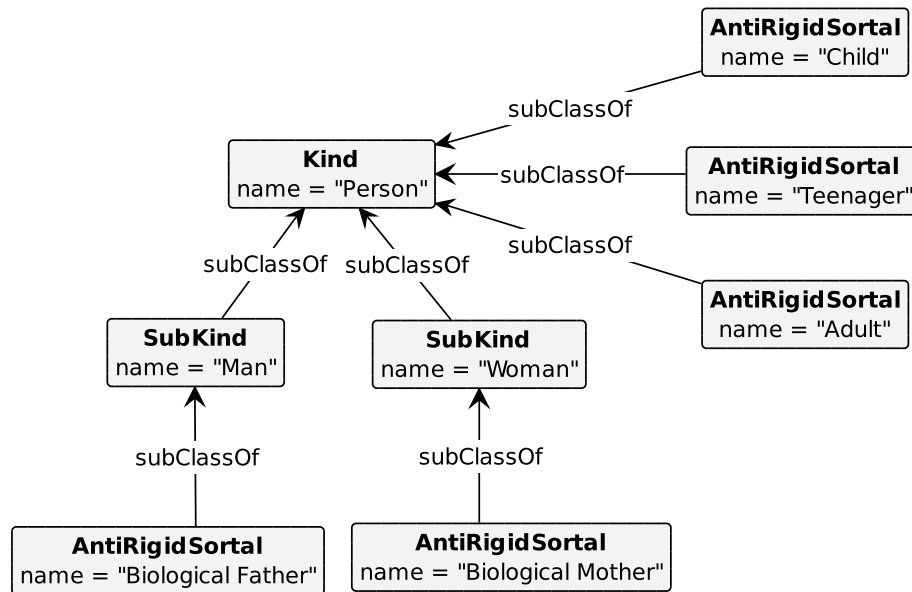


Figure 6 – An example of sortal hierarchy for the kind ‘Person’.

puppy’ captures a cluster of contingent properties of a dog, or that ‘being a husband’ captures a cluster of contingent properties of a person participating in a marriage.

Kinds, subkinds, and the anti-rigid types specializing them are categories of enduring SORTALS. In the philosophical literature, a sortal is a type that provides a uniform principle of identity, persistence, and individuation for its instances (GUZZARDI et al., 2018). To put it simply, a sortal is either a kind (e.g., ‘Person’) or a specialization of a kind (e.g., ‘Student’, ‘Teenager’, ‘Woman’), i.e., it is either a type representing the essence of what things are or a sub-classification applied to the entities that “have that same type of essence”, be it rigid, i.e., a SUBKIND, or anti-rigid, i.e., an ANTI-RIGID SORTAL. Figure 6 presents an example of sortal hierarchy for the kind ‘Person’. The figure depicts this hierarchy as a directed graph, revealing the abstract syntax of taxonomies as adopted in this work. Each node represents a class, whose ontological nature in UFO is represented in bold (here, ‘Kind’, ‘SubKind’ and ‘AntiRigidSortal’). Edges are labeled to identify the relation between the elements; in this figure all edges represent specializations and are labeled ‘subClassOf’.

In general, types that represent properties shared by entities of *multiple kinds* are termed NON-SORTALS, i.e., non-sortals are types whose extension possibly intersect with the extension of more than one kind. Non-sortals too can also be further classified depending on whether the properties captured in their intension are essential (i.e., rigid) properties or not.

Now, before we proceed, we should notice that the logical negation of rigidity is not anti-rigidity but *non-rigidity*. If being rigid for a type *A* means that all instances of *A* are *necessarily* instances of *A*, the negation of that (i.e., non-rigidity) is that there is at least one instance of *A* that can *cease to be* an instance of *A*; anti-rigidity is much stronger than that, it means that *all* instances of *A* can cease to be instances of *A*, i.e., *A*’s intension describes

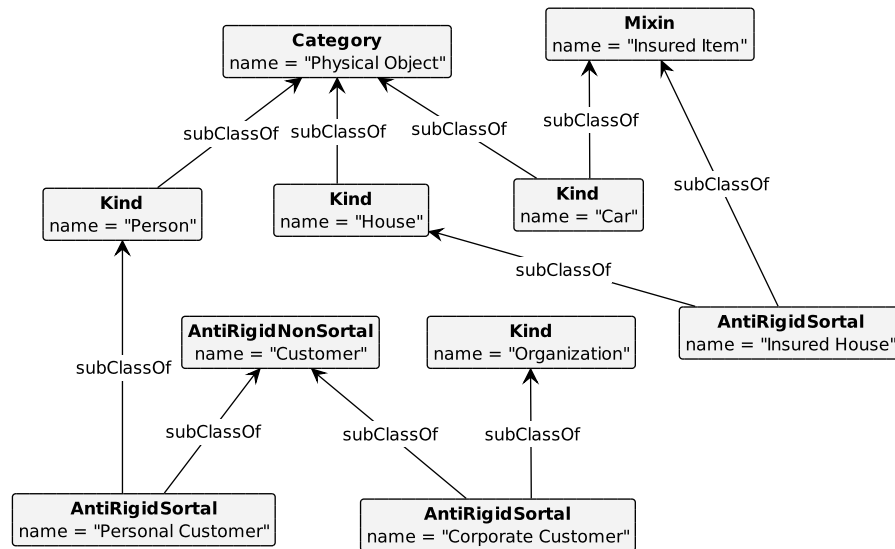


Figure 7 – Non-Sortal examples.

properties that are contingent for all its instances. Finally, we call a type *A* *semi-rigid* iff it is non-rigid but not anti-rigid, i.e., if it describes properties that are essential to some of its instances but contingent to some other instances. Because non-sortal types are dispersive (HIRSCH, 1992), i.e., they represent properties that behave in very different ways with respect to instances of different kinds, among non-sortal types, we have: those that describe properties that apply *necessarily to the instances of all kinds* it classifies (i.e., Rigid Non-Sortals, which are termed CATEGORIES); those that describe properties that apply *contingently to the instances of all kinds* it classifies (ANTI-RIGID NON-SORTALS); those that describe properties that apply *necessarily to the instances of some of the kinds it classifies* but that also apply *contingently to the instances of some other kinds it classifies* (i.e., Semi-Rigid Non-Sortals, termed MIXINS). An example of a category is ‘Physical Object’ representing properties of all kinds of entities that have masses and spatial extensions (e.g., people, cars, watches, buildings); an example of an anti-rigid non-sortal is ‘Customer’ representing contingent properties for all its instances (i.e., no customer is necessarily a customer), which can be of the kinds ‘Person’ and ‘Organization’; an example of a mixin is ‘Insured Item’, which describes properties that are essential to entities of given kinds (e.g., suppose that cars are necessarily insured) but which are contingent to things of other kinds (e.g., houses can be insured but they are not necessarily insured). Figure 7 presents examples of the different types of non-sortals discussed.

Figure 8 represents (with a UML class diagram) the typology of enduring types generated by the possible values of the two properties, sortality and rigidity. As usual, UML arrows with a closed head connect subtypes to their supertypes (the arrowhead pointing to the supertype). Two subtyping relations joined in their arrowheads form a *generalization set*, which here we assume to tessellate the extension of the supertype (pointed to by the joint arrowhead), i.e., these are disjoint and complete generalization sets. There are two superimposed generalization trees, one formed by first considering the *sortality* meta-property (in red) and the other considering

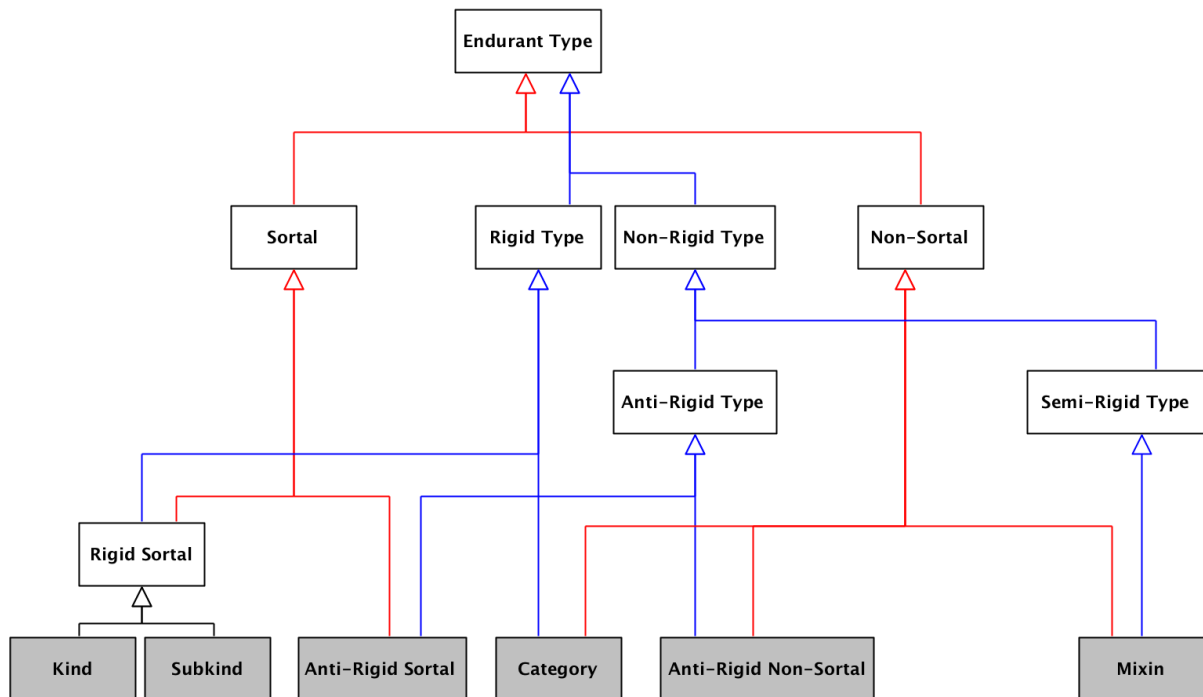


Figure 8 – A taxonomy for Endurant Types.

first the *rigidity* meta-property (in blue). As a result of the combination of these two meta-properties, we have the following six (exhausting and mutually disjoint) types of types (i.e., metatypes): KINDS, SUBKINDS, ANTI-RIGID SORTALS, CATEGORIES, ANTI-RIGID NON-SORTALS, and MIXINS (shaded in Figure 8).

The ontological meta-properties that characterize these different types of types also impose constraints on how they can be combined to form taxonomic structures (GUIZZARDI et al., 2018). As we have already seen, since kinds tessellate our domain and, because all sortals are either kinds or specializations thereof, we have that: (i) no kind can specialize another kind; and (ii) every sortal that is not a kind specializes a unique kind. In other words, every sortal hierarchy has a unique kind at the top. Moreover, from these, we have that any type that is a supertype of a kind must be a non-sortal. But also that, given that every specialization of a kind is a sortal, non-sortals cannot specialize sortals. Finally, given the formal definitions of rigidity (including anti-rigidity), it just follows logically that anti-rigid types (sortals or not) cannot be supertypes of semi-rigid and rigid types (sortals or not – see proof in (GUIZZARDI et al., 2018)). For example, if we determine that ‘Customer’ applies contingently to persons in the scope of business relationships, then a taxonomy in which a rigid type ‘Person’ specializes an anti-rigid type ‘Customer’, as shown in Figure 9, is logically incorrect. Intuitively, a person will be at the same time required through the specialization to be statically classified as a ‘Customer’ while at the same time, being defined dynamically classified as a ‘Customer’, in virtue of the definition of that type. So, either: (i) the definition of ‘Customer’ should be revised to capture only essential properties, becoming a rigid type and thus solving the incorrect specialization problem; or (ii)

a different organization of the taxonomy is required, with ‘Personal Customer’ as an anti-rigid specialization of the rigid type ‘Person’, ‘Corporate Customer’ an anti-rigid specialization of ‘Organization’, and ‘Customer’ as an anti-rigid supertype of ‘Personal Customer’ and ‘Corporate Customer’ as shown in Figure 7 (this solution is identified as the ‘roles with disjoint allowed types’ design pattern in (GUIZZARDI, 2005)).




Figure 9 – Incorrect representation of type ‘Customer’.

2.3 The Multi-Level Theory

So far, we have covered conventional taxonomies built up by establishing specialization relations between types. However, there are several knowledge domains in which types are also considered *instances of* other types. For example, in the biological domain, types of animals such as ‘Dog’ and ‘Cat’ may be considered instances of ‘Species’, and types such as ‘Greyhound’ and ‘Siamese Cat’ may be considered instances of ‘Breed’ specializing ‘Dog’ and ‘Cat’ respectively. In these domains, meta-types or high-order types appear (such as ‘Species’ and ‘Breed’). The taxonomy can thus be considered a *multi-level* one, with a classification level of types whose instances are individuals (in this example, ‘Dog’ and ‘Cat’) and higher classification levels, with types whose instances are types (in this example, ‘Species’ and ‘Breed’). This example in the biological domain is shown in Figure 10. Other examples of multiple classification levels come from domains such as that of organizational roles (or professional positions) (CARVALHO; ALMEIDA, 2015), software engineering (GONZALEZ-PEREZ; HENDERSON-SELLERS, 2006) and product types (NEUMAYR; GRÜN; SCHREFFL, 2009).

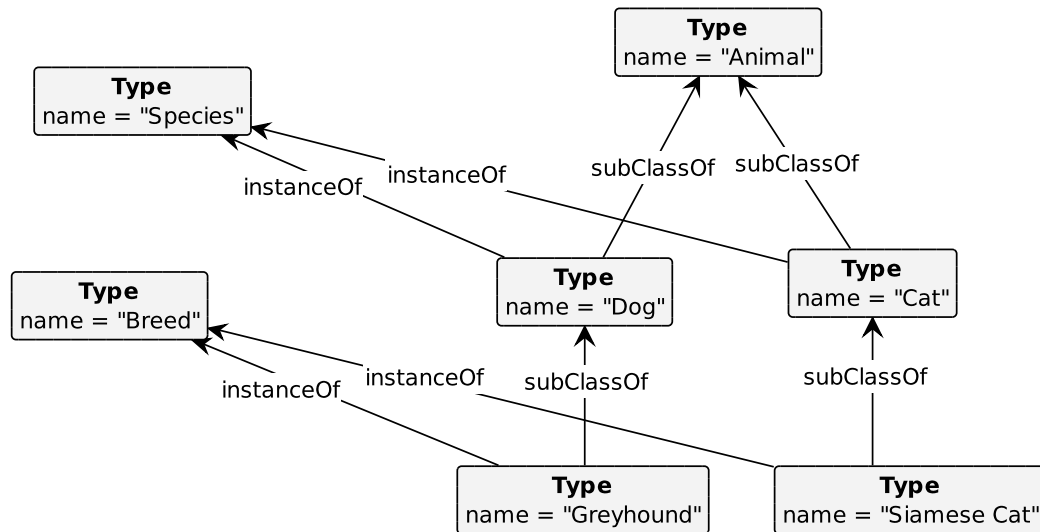


Figure 10 – An example of multi-level taxonomy in biological domain.

Over the last decades, the importance of this modeling phenomenon has justified a number of works under the banner of “multi-level modeling” (ATKINSON; KÜHNE, 2001; GONZALEZ-PEREZ; HENDERSON-SELLERS, 2006; NEUMAYR; GRÜN; SCHREFL, 2009; CARVALHO; ALMEIDA, 2018; ALMEIDA; FRANK; KÜHNE, 2018). Techniques for multi-level modeling must provide modeling concepts to deal with types in various classification levels and the relations that may occur between those types. These approaches embody conceptual notions that are key to the representation of multi-level models, such as the existence of entities that are simultaneously types and instances, the iterated application of instantiation across an arbitrary number of levels, etc (ALMEIDA et al., 2018).

These fundamental notions for multi-level modeling have been captured formally in the MLT Multi-Level Theory, described in various publications over the last years (ALMEIDA; FRANK; KÜHNE, 2018; CARVALHO et al., 2017; CARVALHO; ALMEIDA, 2018). Similarly to the enduring types theory described in Section 2.2, MLT was formalized (CARVALHO; ALMEIDA, 2018) and applied in the design of a well-founded profile for UML (CARVALHO; ALMEIDA; GUIZZARDI, 2016). The theory was also applied fruitfully to uncover problematic taxonomies in Wikidata (BRASILEIRO et al., 2016a; DADALTO et al., 2021), to design the ML2 multi-level modeling language embodying the theory’s rules as syntactic constraints (FONSECA et al., 2021), to support multi-level models using Semantic Web technologies (BRASILEIRO et al., 2016b), etc.

The following key definitions are proposed for MLT (ALMEIDA et al., 2018): *Individuals* are those entities which cannot possibly play the role of type in the instantiation relation. Examples include ALBERT EINSTEIN, LAIKA THE SOVIET SPACE DOG, THE EARTH. *First-order types* are those types whose instances are individuals. Examples include PERSON, DOG, PLANET, CAR. *Second-order types* are those types whose instances are first-order types. Examples include SPECIES, BREED, but also ASTRONOMICAL OBJECT TYPE, CAR MODEL. *Third-order types* are those

types whose instances are second-order types (e.g., TAXONOMIC RANK whose instances may include SPECIES and BREED), and so on. The topmost order can be established as required by applications, and the scheme can thus be extended to cope with an arbitrary number of levels.

Classification levels in MLT are generated by the iterative application of the notion of *powertype* in line with the definition of Cardelli (CARDELLI, 1988). A type pt is *powertype* of a (base) type t iff all instances of pt are (improper) specializations of t and all possible specializations of t are instances of pt . Powertypes in this sense are analogous to powersets. The powerset of a set A is a set that includes as members all subsets of A (including A itself). Using this definition, we can clarify how classification levels are related: if INDIVIDUAL is the type that classifies all possible individuals, then the type that classifies all first-order types—FIRST-ORDER TYPE—can be defined as the *powertype* of INDIVIDUAL. SECOND-ORDER TYPE can be defined as the *powertype* of FIRST-ORDER TYPE, and so on. The types defined in this way, i.e., INDIVIDUAL, FIRST-ORDER TYPE, SECOND-ORDER TYPE, etc., are called *basic types* in MLT. It follows from Cardelli’s definition of powertype that a powertype is unique for a base type, i.e., a base type has one and only one powertype. Further, the base type is unique for a given powertype (for theorems and their proofs see (CARVALHO; ALMEIDA, 2018)).

Using the definition of basic types, MLT partitions a domain taxonomy into strictly stratified levels by establishing that all domain types are specializations of these basic types. This amounts to enforcing the *strict metamodeling principle* (ATKINSON; KÜHNE, 2000). Since basic types form a line connected by powertype relations, a first-order type (such as ANIMAL) is at the same time an instance of the basic type FIRST-ORDER TYPE and a specialization of INDIVIDUAL (since its instances are individuals). A second-order type (such as SPECIES) is at the same time an instance of the basic type SECOND-ORDER TYPE and a specialization of FIRST-ORDER TYPE (since its instances are first-order types), and so on for higher-order types.

MLT also accounts for an important variant of the powertype pattern proposed by (ODELL, 1994) (that inspired the homonymous construct in UML class diagrams). Odell stated simply that a *powertype* is a type whose instances are subtypes of a base type. This means, that, differently from Cardelli, not all subtypes of the base type are required to be instances of the powertype. (In fact, as pointed out by (HENDERSON-SELLERS, 2012), the relation defined by Odell is misnamed *powertype* since, in fact, it is analogous to a *subset* of the *powerset*.) MLT calls the relation between an Odell powertype and its base type *categorization*. For example, we can say that ANIMAL SPECIES *categorizes* ANIMAL; this is because some specializations of ANIMAL are instances of SPECIES (DOG, CAT), but not all of them are (e.g., FEMALE DOG, SIAMESE CAT, while specializations of ANIMAL, are not instances of SPECIES). Figure 11 enriches the model of Figure 10 revealing the ‘First-Order Type’ and ‘Individual’ basic types as well as the categorization relations. (The ‘instanceOf’ edges between the various first-order types specializing ‘Individual’ are omitted for the sake of readability.)

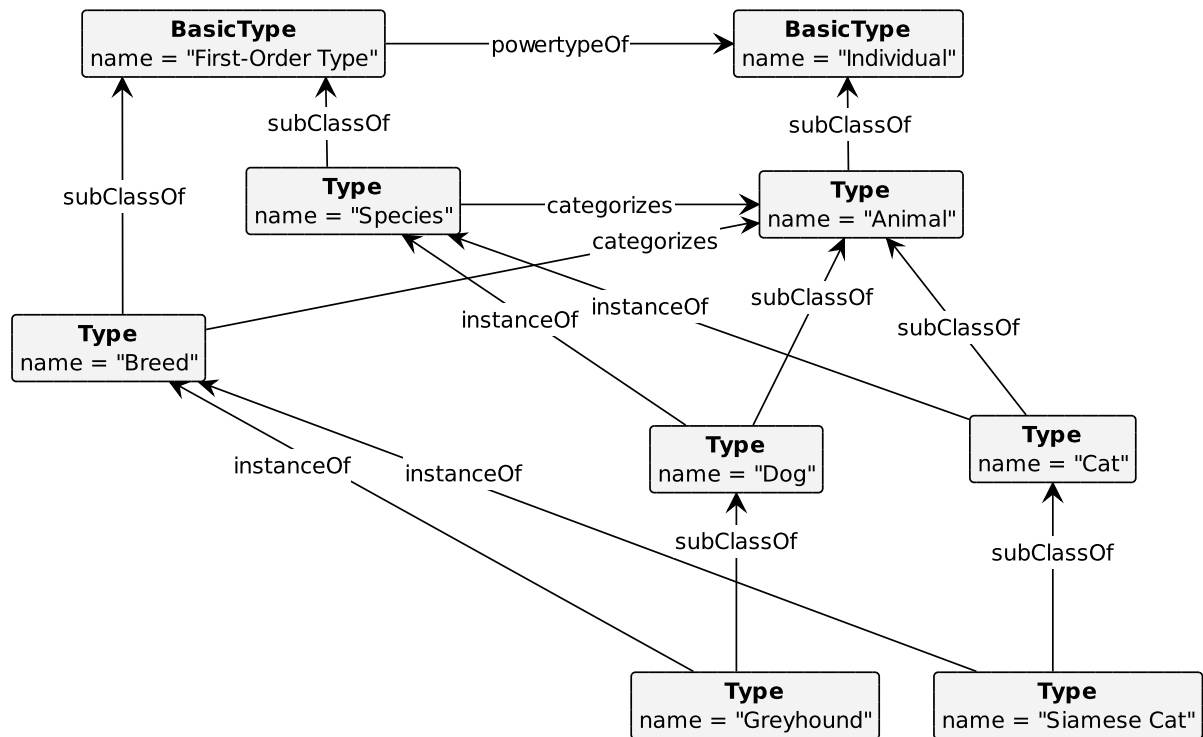


Figure 11 – A multi-level model for the biological domain including variants of the powertype pattern.

The consequences of the theory can be used to identify sound multi-level structures, including the following derived rules (these rules are formally proven theorems that follow from the MLT definitions and axioms (CARVALHO; ALMEIDA, 2018) as summarized in (ALMEIDA et al., 2018)):

- The *instance of* relation in MLT is irreflexive, antisymmetric and anti-transitive. Further, it only relates entities of adjacent levels.
- Every type belongs to a specific order, specializing one and only one basic type.
- *Specialization* (whether proper or not) cannot cross level boundaries (i.e., a first-order type can only be specialized by first-order types, a second-order type can only be specialized by second-order types, and so on).
- Both relations *is powertype of* and *categorizes* can only be applied between adjacent levels, with the base type one order lower than the high-order type.
- The *powertype* of a base type (in Cardelli's sense) is unique for that base type;
- Types that *categorize* a base type always specialize the base type's *powertype*.

Violation of these rules and basic definitions of MLT lead to problematic taxonomies, as illustrated in Figure 12 with a multi-level taxonomy found in Wikidata. As discussed in

(DADALTO et al., 2021), this taxonomy is problematic since ‘Mayor’ is at the same time an instance of ‘Position’ and a specialization of ‘Position’ (through ‘Public Office’). Note the logical contradiction: instantiation places ‘Position’ and ‘Mayor’ in adjacent levels (e.g., ‘Mayor’ as a first-order type and ‘Position’ as a second-order type), while specialization requires them to be at the same level. Further empirical evidence discussed in (DADALTO et al., 2021; BRASILEIRO et al., 2016a; FONSECA et al., 2021) shows that this is a large-scale problem in practice. Our objective here is then to incorporate the basic rules underlying MLT in our graph grammar, and thus rule out problematic multi-level taxonomies by construction.

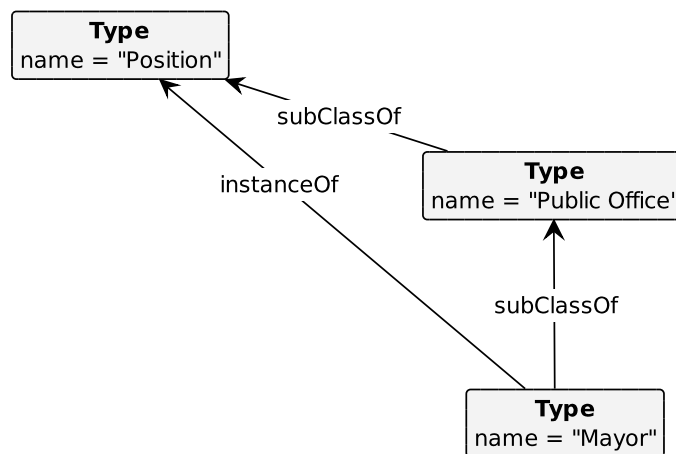


Figure 12 – An incorrect multi-level model present in Wikidata and uncovered in (DADALTO et al., 2021).

2.4 Related Works

Graph grammars and graph transformation (GT) have long been advocated as a suitable formalism for the specification, analysis and verification of models from many distinct domains (HECKEL; TAENTZER, 2020). One of the reasons for this proposition is the smaller syntactic gap between some modeling languages and graph rewriting. Since several languages are graphical, GT rules fit naturally within the model syntax; much more so, for instance, than when compared with text-based formalisms. For instance, Triple Graph Grammars (TGGs), a specialized form of graph grammar, have been applied successfully for several years in a range of application scenarios including: model generation, conformance testing, bidirectional model transformation, and incremental model synchronization (ANJORIN; LEBLEBICI; SCHÜRR, 2015). In fact, model transformation is a very active field of research where the modeling and GT communities intersect. As a particular example, one case study in (GHAMARIAN et al., 2012) describes how GT rules can be used to transform from BPMN to BPEL models. This formalization identified inconsistencies on the original transformation semantics, which was only informally specified.

Several other GT tools have been developed to perform some kind of model transforma-

tion or analysis. For instance, Fujaba (NICKEL; NIERE; ZÜNDORF, 2000) is a tool with support for model-based software engineering and re-engineering, employing UML class diagrams and specialized activity diagrams (called story diagrams in the tool), that are based on graph transformations. Viatra2 (VARRÓ; BALOGH, 2007) and Henshin (ARENDDT et al., 2010) are model transformation tools that are part of the Eclipse Modeling Framework (EMF), and thus can handle Ecore models.

Closer to conventional programming languages and methods, GrGen.NET (JAKUMEIT; BUCHWALD; KROLL, 2010) is a GT-based tool designed for compiler optimization and code refactoring. Similarly, the domain-specific language Chart (MOL; RENSINK; HUNT, 2012) adds GT-based functionality to existing Java programs. The approach relies on a set of annotations to identify the intended graph structure, as well as on user methods to manipulate that structure, within the user's own Java class declarations. The advantage of the approach is that it allows any Java program to be enhanced, non-invasively, with declarative graph rules that can later be used for program analysis. Along the lines of programming language semantics, a type graph (a sort of graph grammar metamodel) for Java was defined in (RENSINK; ZAMBON, 2009) and later refined in (ZAMBON, 2013). Subsequently, the GROOVE tool was used to define a complete formalization for the control flow semantics of Java (ZAMBON; RENSINK, 2018).

As a final note, it is worth mentioning that the GT rules used in GROOVE have the same expressive power than First-Order Logic (FOL) (RENSINK, 2004). Thus, any ontological theory based on FOL can, in principle, be properly supported/formalized by a GROOVE graph grammar.

In the work described in (RUY et al., 2017), the authors argued that instead of proposing methodological rules or semantically motivated syntactical constraints based on the axiomatization of ontological theories, a representation system to build models following ontological restrictions could employ a more productive strategy, leveraging on the fact that formal constraints from the theory impose a correspondence between each particular type of type (characterized by ontological meta-properties) and certain modeling *structures* (or modeling *patterns*). The authors present an approach to derive conceptual ontology patterns from ontologies of different generality levels, ranging from foundational to domain ontologies. Then, they present guidelines that describe how these derived patterns can be applied in combination for building reference domain ontologies in a reuse-oriented process.

The work discussed here is related to other efforts that employ foundational theories in conceptual modeling tasks, detecting important omissions, by employing the state space exploration mechanism supported by GROOVE, in the rule set proposed in (ZAMBON; GUIZZARDI, 2017), where the authors made a first attempt to formalize a representation system to build models following ontological restrictions, capturing modeling patterns derived from the underlying theory of enduring types from OntoUML and their possible relations as graph transformation rules. This work also extends the work started in (BATISTA et al., 2021), adding

to its considerations about the verification scope of the grammar proposed and all the analysis to the construction of multi-level taxonomies. As discussed in (GUIZZARDI, 2005), UFO's typology of enduring types was originally inspired by the typology of types on which the OntoClean methodology is based. A pioneering and among the most important methodologies for the construction of ontologically correct taxonomies, OntoClean is extended in important ways by the graph grammar proposed here: (i) firstly, because OntoClean does not provide any concrete modeling mechanism for building taxonomies, but only a set of methodological guidelines and only for taxonomy *evaluation*. Our proposal instead formalizes a grammar with concrete ontological patterns as modeling constructs; (ii) secondly, OntoClean is focused (even if not explicitly) on object types (also called substantial types) as opposed to more generally addressing taxonomies formed by other enduring types (e.g., types whose instances are dependent entities such as symptoms, marriages, enrollments, etc. (GUIZZARDI et al., 2018)); (iii) finally, OntoClean does not address the construction of higher-order types. In summary, our proposal provides for a concrete modeling mechanism that supports the construction of taxonomic structures that go much beyond what is covered by OntoClean. The first two mentioned points also apply to the graph grammar proposed in (BATISTA et al., 2021).

The aforementioned points (ii-iii) but also (i) (regarding the pattern-based representation support) are also true for theories of type structures that extend OntoClean (e.g. (WELTY; ANDERSEN, 2005; SEYED; SHAPIRO, 2011; SEYED, 2012)). In particular, they also apply to a series of proposals of Order-Sorted Logics extended with ontological predicate meta-hierarchies. These logics play an important role in knowledge representation and, more generally, in symbolic artificial intelligence by supporting the construction and formal verification of ontologically-informed taxonomies, as well as computationally tractable automated reasoning with them. These include: (KANEIWA; MIZOGUCHI, 2005; KANEIWA, 2011), which propose new order-sorted modal logics and tableau calculus to check the (un)satisfiability and validity of sorted modal formulas; (KANEIWA; NGUYEN, 2009; KANEIWA, 2004), which propose order-sorted Horn-calculus combining ontologies and logic programming.

3 Taxonomies of Endurant Types

This chapter presents our first proposed graph grammar that takes into account the ontological constraints imposed by the meta-properties of categories of enduring types (Section 3.1). This graph grammar comprises the identified transformation rules for the construction of ontologically correct taxonomies of enduring types. It is made an assessment of the proposed graph grammar regarding its *soundness* and *completeness* (Section 3.2).

3.1 Graph Transformation Rules for Ontologically Correct Taxonomies

The grammar described in this section was created with GROOVE (GHAMARIAN et al., 2012), a graph transformation tool suitable for the (partial) enumeration of a grammar language, which the tool calls the *state space exploration* of the graph grammar. This grammar has a type graph shown in Figure 13 and established to reflect the taxonomy presented in Figure 8. The grammar starts from an empty host graph. We will use the GROOVE tool later to demonstrate that our proposed taxonomy grammar are *sound* and *complete*, at least up to a certain taxonomy size.

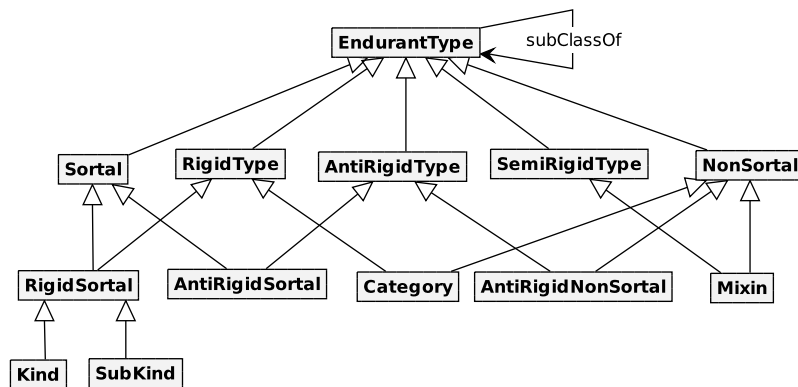


Figure 13 – Type graph established to reflect the taxonomy for Endurant Types.

3.1.1 Introducing New Types

We start by defining transformation rules to introduce a new type in the taxonomy. Types for four of the leaf ontological metatypes given in Figure 8 can be introduced in the taxonomy without being related with a previously introduced type: these include all KINDS and all the non-sortals: CATEGORIES, MIXINS and ANTI-RIGID NON-SORTALS.

Figure 14 shows the four rules that introduce these ‘independent’ types, using the GROOVE visual notation for presenting rules. Each rule is formed only by a green thick lined

box representing the type that will be *created* during rule application. A type has an ontological metatype (the label inside the box). No rule in Figure 14 has preconditions. Therefore, types for these four ontological metatypes can be introduced without requiring the existence of other types or relations in the taxonomy.




Figure 14 – Transformation rules to introduce an independent type.

3.1.2 Introducing Dependent Types

In contrast to non-sortals and kinds, SUBKINDS and ANTI-RIGID SORTALS have preconditions upon their introduction.

In the case of SUBKINDS, their introduction requires the existence of a previous sortal, from which the subkind will inherit a principle of identity. In addition, this sortal must be rigid, to respect the ontological principle that a rigid type cannot specialize an anti-rigid one. These preconditions for the introduction of a new SUBKIND are captured in the rule shown in Figure 15. The *existing* RIGID SORTAL is shown as a gray box in the figure. The green thick “subClassOf” arrow states that a new direct subtyping relation will be introduced in the model.




Figure 15 – Transformation rule to introduce a SUBKIND type.

In the case of an ANTI-RIGID SORTAL type, the only precondition is the existence of a previous sortal, from which the newly introduced ANTI-RIGID SORTAL will inherit a principle of identity. This rule is shown in Figure 16. Differently from a SUBKIND, an ANTI-RIGID SORTAL can specialize any SORTAL (and not only rigid ones).




Figure 16 – Transformation rule to introduce an ANTI-RIGID SORTAL type.

3.1.3 Introducing Specializations for Existing Non-Sortal Types

Having defined rules for the introduction of types, we proceed with rules to insert subtyping relations between two types already present in the model. We start with `CATEGORY` and `MIXIN` specializations, as both of these ontological metatypes have meta-properties that allow their types to be specialized in any `ENDURANT TYPE`, without breaking formal ontology principles.

Figure 17(a) shows a rule that creates a subtyping relation between an existing `CATEGORY` supertype and an existing `ENDURANT` subtype. The red dashed arrow in the figure prevents the introduction of a circularity in the subtyping relations. Note that circularity of specializations may be tolerated in taxonomies structured with *improper specialization relations*, such as the case of `rdfs:subClassOf` used in RDF and OWL. However, for the purposes of this work, we represent only *proper specializations*, i.e., those in which the subtype is different from its supertype (in that some possible instances of the supertype are not instances of the subtype). Red elements in GROOVE rules indicate *forbidden patterns*, i.e., elements that, if present, prevent the rule application. The label “subClassOf*” indicates “subClassOf” paths of any size (including zero, which would amount to the equality between the related classes.) Figure 17(b) shows the analogous rule for the specialization of a `MIXIN`.

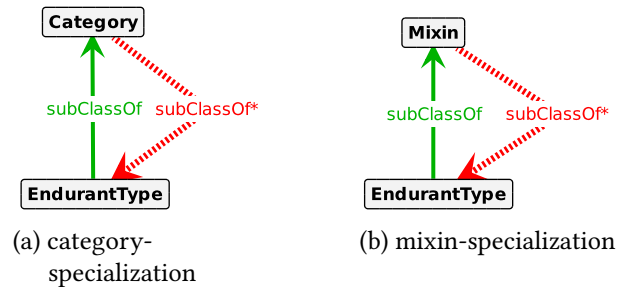


Figure 17 – Transformation rules to specialize a `CATEGORY` or a `MIXIN`.

Finally, the rule depicted in Figure 18 allows the specialization of an `ANTI-RIGID NON-SORTAL` by another `ANTI-RIGID TYPE`.




Figure 18 – Transformation rule to specialize an `ANTI-RIGID NON-SORTAL` type.

3.1.4 Introducing Generalizations for Existing Sortal Types

`KIND` types appear on the top of `SORTAL` hierarchies because kinds provide a principle of identity for all their instances. By definition, kinds cannot specialize other kinds. Therefore,

they can only specialize NON-SORTAL types, more specifically CATEGORIES and MIXINS. These specializations can already be constructed with the rules presented in Section 3.1.3.

SUBKIND types, on the other hand, carry a principle of identity from their supertypes and, ultimately, from exactly one KIND type. The rule shown in Figure 19(a) properly captures this restriction. If there is a RIGID SORTAL distinct from a SUBKIND and not specializing the latter (as defined by the “subClassOf*” red dashed edge) and both carry a principle of identity from the same KIND, then a direct subtyping relation can be created between the two. The black thin lined edges with labels “subClassOf*” and “subClassOf+” indicate that, for the rule to be applied, a specialization relation from the new supertype and from the SUBKIND to the same KIND must already be present, or at least that the new (direct) supertype of the SUBKIND is its own KIND. Subkinds can also specialize any rigid or semi-rigid non-sortal, but these cases are already covered by the rules presented in Section 3.1.3. A similar construction for ANTI-RIGID SORTAL types can be seen in Figure 19(b).

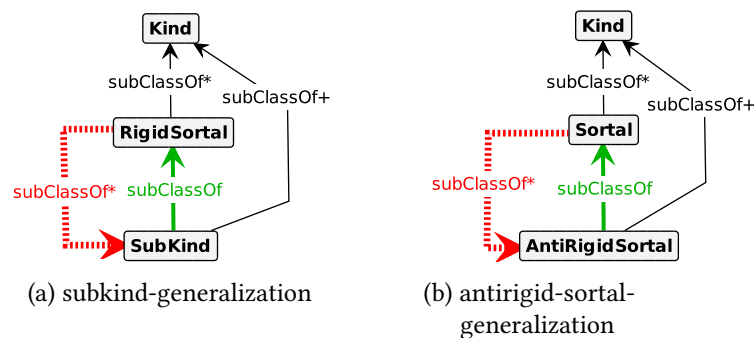


Figure 19 – Transformation rules to generalize a SUBKIND or an ANTI-RIGID SORTAL.

3.1.5 Summary of the Grammar Rules

The grammar to build ontologically correct taxonomies is structured into the following operations:

- The introduction of enduring types that do not require a previously existing type in the taxonomy (kinds and non-sortals);
- The introduction of sortals that require a previously existing sortal in the taxonomy to inherit a principle of identity;
- The specialization of existing non-sortals in other enduring types;
- The generalization of existing sortals in other sortals of the same kind.

3.2 Formal Verification

We use the GROOVE graph transformation tool to carry out a formal verification of the graph grammar presented in Section 3.1. To do so, we employ verification conditions in GROOVE, which formally define the ontological constraints described in Section 2.2, and allow us to perform an analysis over the graph state model (a state in this case corresponds to a *taxonomy shape*, representing an equivalence class of isomorphic taxonomies). We then use the state space exploration functionality of the tool to examine the taxonomy shapes that the grammar generates.

As stated in Section 3.1, our objective with the verification is two-fold: to demonstrate the *soundness* and *completeness* of the proposed graph grammar. Soundness ensures that the grammar rules only produce correct taxonomies, *i.e.*, those that do not invalidate well-formedness constraints. Completeness ensures that any and all correct taxonomies can be produced by a sequence of rule applications. (Certain caveats concerning taxonomy size apply to our formal verification tasks, these will be discussed in Sections 3.2.1 and 3.2.2.)

A *graph condition* in GROOVE is represented diagrammatically in the same way as transformation rules, albeit without creator (green thick lined) elements. A graph condition is satisfied by a taxonomy model if all reader (black thin lined) elements of the condition are present in the model, and all forbidden (red dashed) elements are absent.

Figure 20 shows our first graph condition, capturing the restriction that KINDS must appear at the top of sortal hierarchies, hence not specializing another SORTAL. It is important to note that restrictions are stated *positively* but are checked *negatively*. Thus, the condition in Figure 20 characterizes an undesired model violation (a KIND specializing a SORTAL), and therefore, by verifying that such condition never occurs in any taxonomy model, we can determine the grammar is sound. This same rationale holds for all other conditions shown in this section.




Figure 20 – Restrictive condition of a KIND specializing another SORTAL.

Figure 21 formalizes a second restrictive condition, stating that a SORTAL cannot inherit its principle of identity from more than one KIND. A third condition, shown in Figure 22, captures the situation in which the *rigidity* meta-property is contradicted, that is, when a rigid or semi-rigid type specializes an anti-rigid one. Similarly, the fourth restrictive condition, depicted in Figure 23, represents the situation in which the *sortality* meta-property is contradicted, that is, when a NON-SORTAL type specializes a SORTAL one.




Figure 21 – Restrictive condition of a SORTAL with more than one KIND.

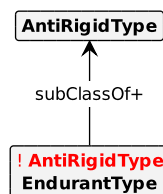


Figure 22 – Restrictive condition of a rigid or semi-rigid type specializing an anti-rigid one.




Figure 23 – Restrictive condition of a NON-SORTAL type specializing a SORTAL one.

To specify a fifth and final restrictive condition, we consider that all SORTALS ultimately should specialize (or be) a KIND, from which they inherit a principle of identity. The violating situation, in which a SORTAL does not specialize a KIND, is shown in Figure 24.



Figure 24 – Restrictive condition of a SORTAL without a KIND.

Note that each restrictive condition emerges as a direct consequence of an ontological restriction, what gives reliance to the assumption that the graph conditions are correct.

3.2.1 Verifying Soundness

The first step in verifying the soundness of the graph grammar proposed is to enumerate its language, i.e., construct all possible taxonomy shapes reachable by any sequence of rule applications. Subsequently, the graph conditions just presented are checked against these constructed shapes. If any model triggers one or more graph conditions, then we know the model violates some ontological restrictions, and therefore it is incorrect. Consequently, the

goal of the soundness analysis is to verify that no taxonomy shape in the language is incorrect. To perform the grammar state space exploration we use the GROOVE Generator, a command-line tool designed for this task. Details of GROOVE usage can be found at the tool manual (available at <https://sourceforge.net/projects/groove/>). Additional case studies that illustrate the tool functionalities are presented in (GHAMARIAN et al., 2012).

A major caveat in the first step above is that the grammar language is *infinite*, thus preventing a complete enumeration in a finite amount of time. To cope with this situation, we need to perform a *bounded* exploration with the GROOVE tool. In this setting, our bound N is the number of types present in a taxonomy shape. When performing the exploration, the tool managed to generate a total of 2,123,196 taxonomy shapes up to a bound $N = 6$, with a breakdown of this total per bound value shown in Table 1. The table also shows that our soundness goal was validated (at least up to $N = 6$), with no taxonomy shapes being flagged as incorrect by the graph conditions.

# types (N)	Produced taxonomy shapes	Incorrect taxonomy shapes
1	4	0
2	21	0
3	160	0
4	2,032	0
5	46,885	0
6	2,074,094	0

Table 1 – Results of soundness analysis for the enduring types graph grammar.

Given the inherently exponential growth of the number of possible taxonomy shapes with respect to bound N , it was not possible to continue the exploration for $N = 7$ and beyond due to memory limitations (the execution was halted after several million models partially produced.) This *state space explosion* is a common problem for all explicit state model checkers, such as GROOVE (GHAMARIAN et al., 2012).

To support that the soundness results in Table 1 are significant, we rely on the *small scope hypothesis*, which basically claims most design errors can be found in small counterexamples (GAMMAITONI; KELSEN; MA, 2018). Experimental results suggest that exhaustive testing within a small finite domain does indeed catch all type system errors in practice (ROBERSON et al., 2008), and many case studies using the tool Alloy have confirmed the hypothesis by performing an analysis in a variety of scopes and showing, retrospectively, that a small scope would have sufficed to find all the bugs discovered (JACKSON, 2019).

3.2.2 Verifying Completeness

The verification described in the previous subsection assures that all taxonomy shapes produced are correct, but does nothing to persuade us that any and all possible correct taxonomy

shapes can be produced. To provide this kind of assurance is the goal of the completeness verification described in this subsection.

To perform the completeness analysis we need to consider not only correct taxonomy shapes but also the incorrect ones. To this end, we developed another, completely permissible, graph grammar that allows the creation of both correct and incorrect models. The grammar is quite simple, with six rules for the unrestricted creation of the leaf types of types in Figure 8, and one rule allowing the introduction of a subtyping relation between any two enduring types. This completely permissive grammar is presented in Appendix A.

The results of the exploration with this new permissible grammar are presented in Table 2. As expected, the rate of growth in this scenario is even steeper, given that more models can be produced. The tool was able to perform a bounded exploration up to $N = 5$, with larger bounds exceeding the available memory. The second column of Table 2 lists all shapes created with the new grammar, both correct and incorrect. We again use the same graph conditions to flag violations of ontology restrictions in the models. If a taxonomy shape triggers *any* of the graph conditions, then it is considered incorrect. Conversely, if *no* graph condition is triggered by a model, then it certainly describes a correct taxonomy shape. The last two columns in the table summarize this classification.

# types (N)	All taxonomy shapes	Incorrect shapes	Correct shapes
1	6	2	4
2	57	36	21
3	956	796	160
4	30,741	28,709	2,032
5	1,958,538	1,911,653	46,885

Table 2 – Results of completeness analysis for the enduring types graph grammar.

The completeness goal can be verified by a comparison between the **Correct shapes** column of Table 2 and the **Produced taxonomy shapes** column of Table 1. It can be seen immediately that all values up to $N = 5$ match. Given that the permissible grammar produces all possible models (correct and incorrect), this allows us to conclude that the taxonomy grammar of Section 3.1 produces *all* correct taxonomy shapes up to that size, and *only* the correct ones.

Partial Completeness Results

To strengthen the evidence of completeness to our graph grammar for enduring types, we developed three additional permissible graph grammars, of which one, that is presented in Appendix B, does not necessarily respect the meta-property of rigidity, whose constraint is presented in Figure 22; another, that is presented in Appendix C, does not necessarily respect the meta-property of sortality, whose constraint is presented in Figure 23; and a last, that is presented in Appendix D, does not necessarily respect the constraints related to KINDS, shown in Figures 20, 21 and 24.

The results of the exploration with these three permissive grammars are presented in Tables 3, 4 and 5. The tool was able to perform a bounded exploration up to $N = 6$, for all these three permissive grammars, with larger bounds exceeding the available memory. We again use the same graph conditions to flag violations of ontology restrictions in the models. Note that the set of incorrect models produced by each of these grammars is disjoint from the others, since the constraint not enforced by one grammar is always respected in all others.

# types (N)	All taxonomy shapes	Incorrect shapes	Correct shapes
1	4	0	4
2	24	3	21
3	231	71	160
4	4,005	1,973	2,032
5	132,672	85,787	46,885
6	8,649,770	6,575,676	2,074,094

Table 3 – Results of completeness analysis for the enduring types graph grammar that can only violate rigidity.

# types (N)	All taxonomy shapes	Incorrect shapes	Correct shapes
1	4	0	4
2	24	3	21
3	235	75	160
4	4,217	2,185	2,032
5	146,398	99,513	46,885
6	10,055,218	7,981,124	2,074,094

Table 4 – Results of completeness analysis for the enduring types graph grammar that can only violate sortality.

# types (N)	All taxonomy shapes	Incorrect shapes	Correct shapes
1	4	0	4
2	27	6	21
3	277	117	160
4	4,685	2,653	2,032
5	139,253	92,368	46,885
6	7,634,331	5,560,237	2,074,094

Table 5 – Results of completeness analysis for the enduring types graph grammar that can only violate KIND constraints.

These tables 3, 4 and 5 show that *none* of the 14,543,863 taxonomy shapes (the same 2,123,196 correct ones presented in Table 1, and 12,420,667 incorrect ones) is an incorrect taxonomy shape not produced by the graph grammar presented in 3.1, providing additional evidence that the taxonomy grammar of Section 3.1 produces *only* correct shapes, and *all* of

them. In other words, these tables provide additional evidence that our enduring types graph grammar is not only *sound*, but also *complete*.

3.2.3 Verification Scope Matters

The small scope hypothesis asserts that a problem in the grammar would be revealed in violating taxonomy shapes of small size. To lend credence to this hypothesis, we need to consider the nature of well-formedness violations (according to the restrictive conditions shown in Figures 20–24). For example, in the case of rigidity and sortality, violations arise from specializations relating classes with particular meta-properties: rigid classes cannot specialize anti-rigid classes (Figure 22) and non-sortals cannot specialize sortals (Figure 23). The violations of rules concerning rigidity and sortality can therefore arise with taxonomies of size 2 (e.g., a single rigid class specializing a single anti-rigid class and a single non-sortal specializing a single sortal). Taxonomies of size 2 also reveal violations of the rules involving kinds (a kind specializing another kind or any other sortal, Figure 20). These sorts of violations are all covered by our analysis (and are included in the 36 incorrect taxonomy shapes of size 2 in the second row of Table 2). So, a problem involving these violations would have been revealed if the proposed grammar could produce them. Other problems involving kinds are manifested with taxonomies of size 3, e.g., when a sortal specializes two kinds (Figure 21), or even size 1, in a taxonomy with a single sortal other than a kind (Figure 24). Given the nature of the restrictive conditions, they can only be violated in taxonomies larger than 3 classes either due to: (i) the occurrence of a violating fragment of size 1–3 (corresponding to one of the Figures 20–24) or (ii) due to the transitivity of the `subClassOf` relation. Problems involving transitivity would already be revealed with taxonomies of size 3 (in most cases) and 5 (in the case of Figure 21). In conclusion, all restrictive conditions can be possibly violated with taxonomy shapes of size 5 or less, which gives us confidence that the exhaustive exploration up to that size (1,958,538 taxonomy shapes as reported in Table 2) is relevant to our verification goals.

With respect to completeness, there is an inherent limitation of the exhaustive enumeration of states as proposed here. In our future work, we intend to address this limitation by exploring other techniques, such as other ways to partially explore the grammars state space or the use of proof assistants.

4 Multi-Level Taxonomies

This chapter presents our second proposed graph grammar that takes into account the constraints imposed by the Multi-Level Theory (Section 4.1) on types. This graph grammar comprises the identified transformation rules for the construction of correct multi-level taxonomies, i.e., taxonomies involving types of different classification levels. We have made an assessment of the proposed graph grammar regarding its *soundness* and *completeness* (Section 4.2).

4.1 Graph Transformation Rules for Correct Multi-Level Taxonomies

Our construction of multi-level taxonomies starts from a host graph that contains only the basic type representing the class of *individuals*, as shown in Figure 25.



Figure 25 – Initial host graph only with one basic type representing the INDIVIDUAL class.

The type graph shown in Figure 26 is used.




Figure 26 – Type graph established to reflect the MLT relations.

Taxonomies are then further built up with operations that introduce first-order types by specializing this first basic type (Section 4.1.1); or that introduce new basic types for additional levels of classification (or orders – Section 4.1.2). Then, high-order types can be introduced as (Odell) powertypes of existing (lower-order) types (Section 4.1.3). Finally, specialization and instantiation between existing elements can be introduced (Sections 4.1.4 and 4.1.5).

4.1.1 Introducing a New First-Order Type

The order of a particular basic type is inferred by the length of the *powertype relation* chain. The last basic type in this chain—the one that is not a powertype—is always the basic type representing the type of all individuals (and is introduced initially in the host graph as shown in Figure 25). By definition, the specializations of this basic type are first-order types,

which are introduced exactly as direct or indirect (proper) specializations of the basic type at the tail of the sequence of basic types. This rule is shown in Figure 27. The forbidden fragment guarantees the selection of the right basic type for specialization (the basic type representing the INDIVIDUAL class).




Figure 27 – Transformation rule to introduce a FIRST-ORDER TYPE.

4.1.2 Introducing a New Order

The model can be extended to accommodate higher levels of types by creating the *powertype* of the basic type of the highest order present in the model thus far, i.e., the powertype with no existing powertype in the model (see Figure 28).



Figure 28 – Transformation rule to introduce a new order.

4.1.3 Introducing a New High-Order Type

New high-order types are then introduced in the model as Odell powertypes categorizing a base type at one order below. As shown in Figure 29, the categorizer specializes the basic type that is the powertype of the basic type at the order below (to which the categorized base type belongs). Note that, as all specializations of a basic type are instances of its powertype (which is the basic type of the order above), we omit the instantiations of that higher-order basic type, as they are redundant.

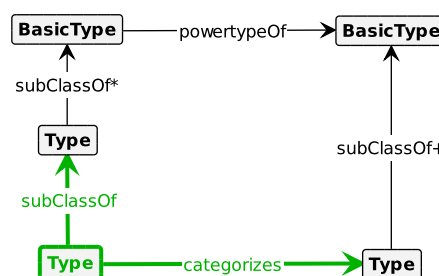


Figure 29 – Transformation rule to introduce a new high-order type.

4.1.4 Introducing Specialization Relations

As stated by MLT, a (proper) specialization relation can only hold between two types of the same order, i.e., types that specialize the same basic type. In addition, to enforce the semantics of specialization, we require the subtype to have no instances when the *subClassOf* relation is first introduced in the model. This prevents the situation in which an entity is an *instance of* a subtype without being an *instance of* the supertype. All these restrictions are formalized in the graph transformation rule given in Figure 30.

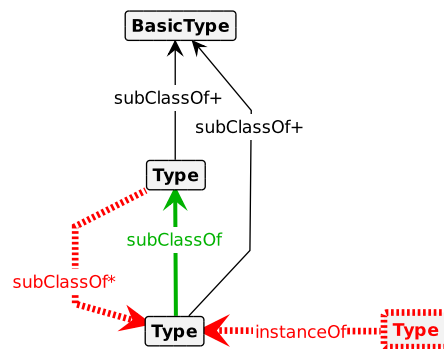


Figure 30 – Transformation rule to introduce a new specialization relation.

4.1.5 Introducing Instantiation Relations

A high-order type can be instantiated by a type that specializes the base type it categorizes, as shown by the rule in Figure 31. The new *instance of* the high-order type must be an *instance of* all its supertypes that are not basic types, to ensure that the semantics of specialization is enforced. The ‘!’ in the “!instanceOf” red dashed arrow in the rule indicates a *forbidden absence* of an *instance of* relation.

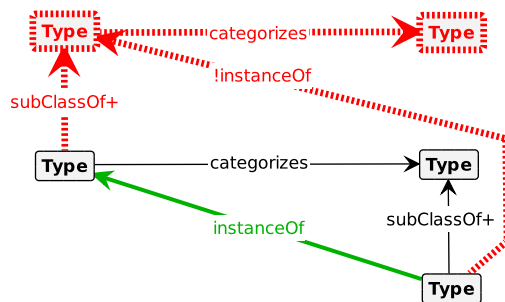


Figure 31 – Transformation rule to introduce a new instantiation relation.

4.1.6 Summary of the Grammar Rules

The grammar to build correct multi-level taxonomies is structured into the following operations:

- The introduction of new first-order types, specializing the basic type representing the

INDIVIDUAL class;

- The introduction of a basic type representing a new order;
- The introduction of new high-order types, specializing some high-order basic type and categorizing some type at one level below;
- The introduction of specialization relations between types of the same order;
- The introduction of instantiation relations between a specialization of a type and its (higher-order) categorizer.

4.2 Formal Verification of MLT Constraints

The formal verification of the graph grammar presented in Section 4.1 follows the same process discussed in Section 3.2 for the enduring types grammar. In order to use GROOVE to analyze the new grammar language, we again need verification conditions, this time to formally define the constraints that follow from the MLT axiomatization, as summarized in Section 2.3. These conditions are discussed in the following subsections, with the verification results presented subsequently.

Once more, as in 3.2, the fact that the restrictive conditions are direct consequences of the MLT theorems gives reliance that they are correct.

4.2.1 Constraints Related to Basic Types

We start by presenting constraints on how basic types can be related. Figure 32 shows our first graph condition regarding basic types, capturing the restriction that a basic type cannot have more than one powertype.




Figure 32 – Restrictive condition of a basic type with more than one powertype.

Conversely, the restriction in Figure 33 shows that different basic types cannot have the same powertype.




Figure 33 – Restrictive condition of different basic types with the same powertype.

The rule presented in Figure 34 captures the restriction that different basic types must be, directly or indirectly, related by the powertype relation, since they have different orders. It is important to note once again that all restrictions in this section are stated *positively* but are checked *negatively*. Thus, the condition shown in Figure 34 describes the undesired case where two basic types in the model are not in a transitive powertype relation. Therefore, by checking in the verification phase that such case never occurs, we ensure that all type orders are related.




Figure 34 – Restrictive condition of unrelated basic types.

Finally, the condition in Figure 35 states that two basic types cannot be mutually connected by the powertype relation. Since a base type is an instance of a Cardelli powertype, basic types thus related would in fact instantiate each other, which would contradict their purpose to establish the basic structure for the stratified level scheme.




Figure 35 – Restrictive condition of mutual powertype relation between basic types.

4.2.2 Constraints Related to Level Stratification

Figure 36 shows a key graph condition related to level stratification, capturing the restriction that every type must be either a basic type or a specialization of a basic type.




Figure 36 – Restrictive condition of a type with no basic type as supertype.

The graph condition depicted in Figure 37 captures the restriction that no type can specialize more than one basic type, otherwise the type would belong to more than one order or level.



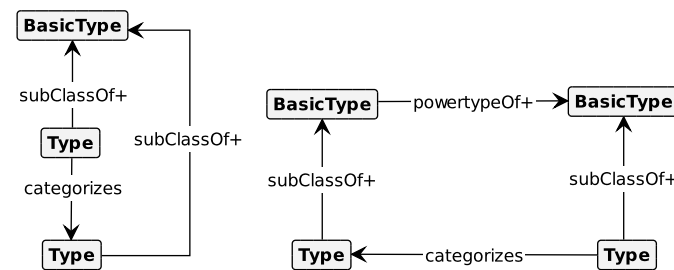


Figure 37 – Restrictive condition of a type specializing more than one basic type.


Three more conditions related to level stratification are presented in Figures 38(a, b, c), indicating that a *categorization* relation cannot exist between types that are not in adjacent levels. In other words, a *categorization* cannot cross multiple levels (Figure 38a) or occur between two types in the same level (Figure 38b), and a type cannot categorize a type at a higher order (Figure 38c). Types can only categorize lower-order types one level below.



(a) categorization-across-multiple-levels



(b) intra-level-categorization



(c) reverse-categorization

Figure 38 – Restrictive conditions of categorization between nonadjacent levels.

Analogously to categorization, instantiation cannot cross multiple levels or occur between two types in the same level, and a type cannot instantiate another at a lower order. These restrictions are captured by the graph conditions shown in Figures 39(a, b, c). As can be seen, the conditions in Figures 38 and 39 differ only in the “categorizes” and “instanceOf” relations and their directions.

The last condition concerning level stratification is shown in Figure 40, capturing the restriction that a type cannot specialize another type in a different level (whether directly or indirectly).




Figure 40 – Restrictive condition of specialization between different levels.

4.2.3 Constraints Related to Categorization and Specialization

The MLT theory summarized in Section 2.3 defines one more condition with respect to categorization, that stems directly from the definition of an Odell powertype. This condition is shown in Figure 41, requiring that all instances of a categorizer must be a proper specialization of the categorized type.



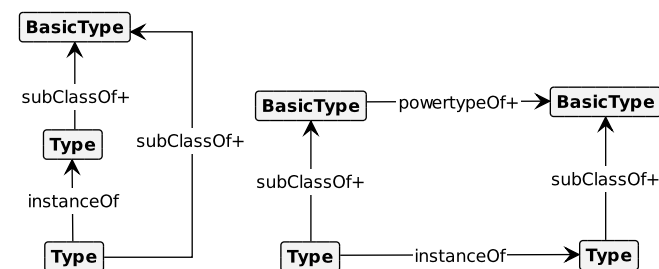


Figure 41 – Restrictive condition of a categorizer instance not proper specializing the categorized type.

A similar condition follows straightforwardly from the semantics of specialization. This condition is depicted in Figure 42, where an instance of a type must instantiate all of its (direct or indirect) supertypes.



(a) instantiation-across-multiple-levels



(b) intra-level-instantiation

(c) reverse-instantiation

Figure 39 – Restrictive conditions of instantiation between nonadjacent levels.




Figure 42 – Restrictive condition of an instance violating the definition of specialization.

4.2.4 Verifying Soundness

Similarly to what was reported in Section 3.2.1 for the enduring types grammar, in order to verify that the graph grammar proposed for multi-level taxonomies is sound, we need to enumerate its language, constructing all possible taxonomy shapes reachable by any sequence of rule applications. Subsequently, the graph conditions presented for multi-level modeling are checked against these constructed shapes. If a model triggers any of the graph conditions, then the taxonomy shape is incorrect, because it violates some well-formedness rules from the theory. Consequently, the goal of the soundness analysis is to verify that no taxonomy shape in the language is incorrect.

As done in Section 3.2.1, given that the grammar language is *infinite*, we perform a *bounded* exploration with the GROOVE tool, with bound N being the number of types present in a taxonomy shape. The tool managed to generate a total of 149,282 taxonomy shapes up to a bound $N = 7$, with a breakdown of this total per bound value shown in Table 6. The table also shows that the soundness goal was validated (at least up to $N = 7$), with no shapes being flagged as incorrect by the graph conditions.

# types (N)	Produced taxonomy shapes	Incorrect taxonomy shapes
2	2	0
3	5	0
4	22	0
5	196	0
6	3,685	0
7	145,372	0

Table 6 – Results of soundness analysis for the MLT-based graph grammar.

Once more, the exponential growth of the number of possible taxonomy shapes w.r.t. bound N prevented an enumeration for larger values, a common limiting factor when performing an exhaustive explicit enumeration of a grammar language (GHAMARIAN et al., 2012). For the MLT-based grammar, an enumeration for $N = 8$ and beyond was not possible due to time and memory limitations. Nevertheless, to support the significance of the soundness results presented in Table 6, we again rely on the *small scope hypothesis*, as discussed in Section 3.2.3, which requires violations to be detectable in small taxonomy shapes. Concerning this matter, the restrictive conditions on basic types (Figures 32–35) can be violated with taxonomy shapes

of size 2 and 3. Conditions on stratification (Figures 36–40) can be violated with taxonomy shapes of size 5 or less. Conditions on categorization and specialization (Figures 41 and 42) can be violated with taxonomy shapes of size 5 (2 of which are basic types).

4.2.5 Verifying Completeness

For the completeness analysis of the MLT-based grammar we repeat the same procedures described in Section 3.2.2, where we consider both correct and incorrect taxonomy shapes. As before, this requires the elaboration of an additional, permissible graph grammar, that allows the creation of correct and incorrect models. The construction of taxonomies with this grammar starts from a host graph that contains only the basic type representing the class of *individuals*, as shown in Figure 25. This grammar is presented in Appendix E and is composed by five rules:

- A rule to create new levels of classification, as in Figure 28;
- A rule to create a new non-basic type as a specialization of any existing type;
- A rule to introduce a categorizer of any non-basic type as a specialization of any existing type;
- A rule to introduce a specialization relation between any two non-basic types, avoiding circularity;
- And a rule to introduce an instantiation relation between any two non-basic types, avoiding circularity.

The exploration results for this permissible grammar are presented in Table 7. The second column lists the number of taxonomy shapes created with the permissible grammar, both correct and incorrect. As we can see, at least up to $N = 5$, the number of correct taxonomy shapes matches those produced by our grammar (second column of Table 6). Even more acutely than in the case of the completeness analysis of the UFO-based grammar presented earlier, there is an exponential growth in the number of shapes, fueled by the additional relations introduced in a multi-level taxonomy (instantiation and categorization). This suggests further work is required to support a completeness claim. Note that given the formation constraints of the grammar, a very large portion of the taxonomy shapes generated by the permissive (control) grammar is incorrect. This shows that much is required from a free-form modeler to produce a correct model on their own by employing what we call – in comparison with a high-level pattern language – ‘low-level’ primitives; the liberty in an unrestricted setting is so vast that some mistakes are bound to occur, especially in large constructs. Therefore, tool automation and assistance is of the utmost importance for a model designer.

# types (N)	All taxonomy shapes	Incorrect shapes	Correct shapes
2	2	0	2
3	23	18	5
4	4,064	4,042	22
5	6,480,294	6,480,098	196

Table 7 – Results of completeness analysis for the MLT-based graph grammar.

Partial Completeness Results

Further evidence of the completeness of our MLT-based grammar can be given following a strategy similar to that given in Section 3.2.2. To this end, we define three additional graph grammars to partially verify the completeness of the grammar presented in Section 4.1. The first one produces taxonomy shapes in that categorization relations do not necessarily respect the stratification of classification levels, but specialization relations do. In other words, in this grammar, that is presented in Appendix F, categorization relations do not necessarily occur between types of adjacent levels, i.e., the constraints presented in Figures 38(a, b, c) are not enforced, but the constraint presented in Figure 40 is. The second grammar, that is presented in Appendix G, produces taxonomy shapes in that specialization relations do not necessarily respect the stratification, i.e., the constraint presented in Figure 40 is not enforced. Finally, a third grammar, that is presented in Appendix H, produces taxonomy shapes in that instantiation relations do not necessarily respect the definitions of categorization and specialization, i.e., the constraints presented in Figures 41 and 42 are not enforced.

Table 8 summarizes the results obtained with the permissible grammar that can violate the stratification by categorization relations, but not by specialization relations. The Table 9 summarizes the results obtained with the permissible grammar that can violate the stratification by specialization relations. And, finally, the Table 10 summarizes the results obtained with a permissible grammar where instantiation relations can violate the definitions of categorization and specialization.

# types (N)	All taxonomy shapes	Incorrect shapes	Correct shapes
2	2	0	2
3	5	0	5
4	22	0	22
5	201	5	196
6	3,909	224	3,685
7	168,078	22,706	145,372

Table 8 – Results of completeness analysis for the MLT-based graph grammar that can only violate the stratification by categorization relations.

# types (<i>N</i>)	All taxonomy shapes	Incorrect shapes	Correct shapes
2	2	0	2
3	5	0	5
4	25	3	22
5	348	152	196
6	22,641	18,956	3,685
7	8,340,140	8,194,768	145,372

Table 9 – Results of completeness analysis for the MLT-based graph grammar that violates the stratification by specialization relations.

# types (<i>N</i>)	All taxonomy shapes	Incorrect shapes	Correct shapes
2	2	0	2
3	5	0	5
4	23	1	22
5	218	22	196
6	4,417	732	3,685
7	190,919	45,547	145,372

Table 10 – Results of completeness analysis for the MLT-based graph grammar with instantiation relations violating the categorization and specialization definitions.

Each one of these three permissive grammars generated all the 149,282 taxonomy shapes produced by our MLT-based grammar, and no more correct models, what can be seen comparing the Column **Produced taxonomy shapes** of the Table 6 with the Column **Correct shapes** of Tables 8, 9 and 10. These three permissive grammars generated a total of 8,432,398 taxonomy shapes (8,283,116 of them are incorrect ones). As in 3.2.2, this provides further evidence in favor of the completeness of our MLT-based grammar.

5 Joining the Foundational Theories

Having designed and assessed two independent graph grammars each corresponding to a foundational theory, in this chapter, we are ready to combine both grammars into a single rule set, which is presented in Section 5.1. Subsequently, in Section 5.2, we reuse the graph conditions for the grammar of enduring types (Section 3.2) and the grammar of multi-level taxonomies (Section 4.2), to show that the results of soundness still carry over to the combined grammar. In other words, the combined grammar consolidates rules that guarantee the production of correct *ontologically well-founded multi-level taxonomies*.

As discussed in (GUIZZARDI et al., 2015), the ontological distinctions among enduring types discussed in Section 2.2 can also apply to higher-order types. For example, *Bird Species* can be conceived as a second-order KIND, whose instances are individual bird species such as *American Eagle* or *Emperor Penguin*, which instantiate that type necessarily. In their turn, *Endangered Bird Species* or *Extinct Bird Species* are anti-rigid second-order types (phases) that can contingently classify instances of *Bird Species*, and *Recognized Bird Species* is an anti-rigid second-order type (a role) played by instances of *Bird Species* when officially recognized. Hence, a combined theory is required to take into account these ontological distinctions for higher-order types in multi-level taxonomies. It is important to note that the rules of our combined grammar cannot be used to model the categories of enduring types, presented in 2.2 in the Figure 8. The rules can only be used to model domain types.

5.1 Graph Transformation Rules to Build Ontologically Correct Multi-Level Taxonomies

Our construction of ontology-based multi-level taxonomies starts from the same host graph of Section 4.1 (Figure 25). This initial host graph contains only a single basic type representing the class of *individuals*.

The type graphs presented in Figures 43 and 44 are used. The dashed lined node type **Type** is abstract, in the sense that it cannot be directly instantiated. In practice, only the six leaf types of types in Figure 43 can be directly instantiated.

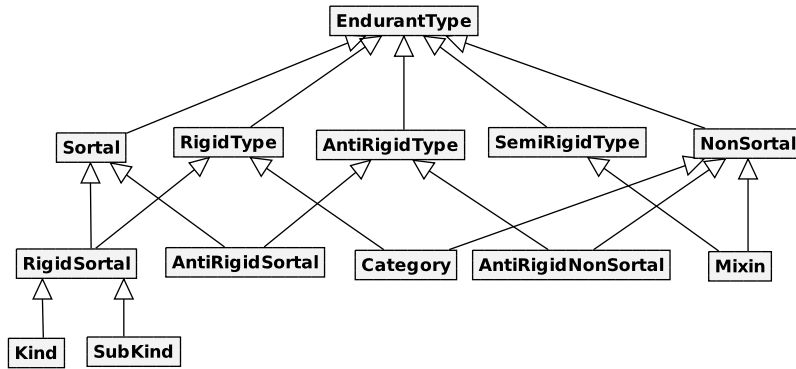


Figure 43 – Type graph established to reflect the taxonomy for Endurant Types.




Figure 44 – Type graph established to reflect the MLT relations.

5.1.1 Introducing First-Order New Types

The enduring independent types originally presented in Section 3.1.1, i.e., KINDS, CATEGORIES, MIXINS and ANTI-RIGID NON-SORTALS, are now introduced in the first level as a direct specialization of the basic type representing the INDIVIDUAL class. The amalgamation of the rules given in Figures 14(a-d) and 27 lead to the new rules shown in Figures 45(a-d).

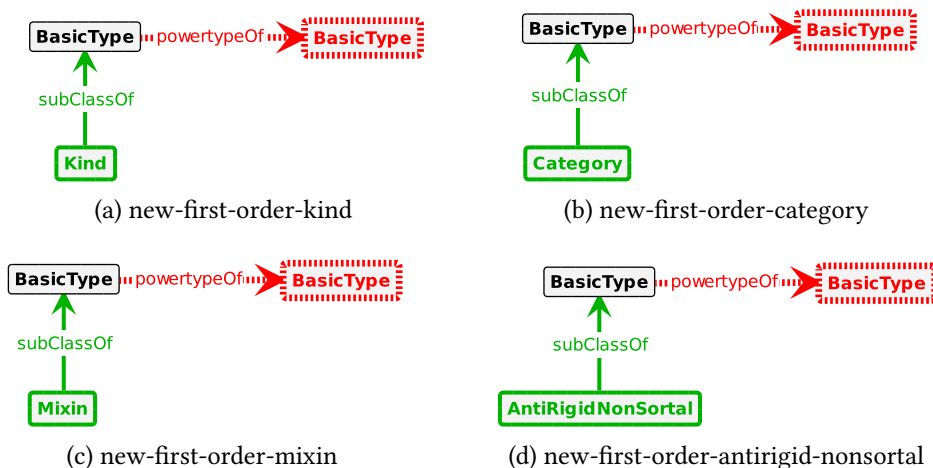


Figure 45 – Transformation rules to introduce a new first-order independent type.

5.1.2 Introducing First-Order Dependent Types

Similarly to the independent ones, the enduring dependent types from Section 3.1.2, i.e., SUBKINDS and ANTI-RIGID SORTALS, are introduced in the first level as a specialization of

some SORTAL type from which they inherit a principle of identity. The new rules depicted in Figures 46(a, b) result from the combination of the rules originally given in Figures 15, 16 and 27. In case of a SUBKIND, its principle of identity is inherited from a RIGID-SORTAL. In case of an ANTI-RIGID SORTAL, it is inherited from any SORTAL.

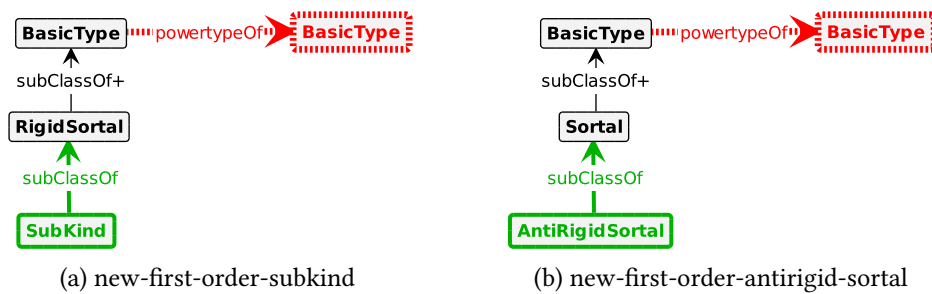


Figure 46 – Transformation rules to introduce a new first-order dependent type.

5.1.3 Introducing a New Order

A new order in the taxonomy is introduced in the same way as discussed in Section 4.1.2, via the creation of a *powertype* of a basic type in the highest order so far. This is accomplished by the rule shown in Figure 47, which is the same as the original rule in Figure 28, being repeated here just for convenience.




Figure 47 – Transformation rule to introduce a new order.

5.1.4 Introducing High-Order New Types

Endurant independent types are introduced in high levels as a direct specialization of a basic type representing some high-order, categorizing any ENDURANT TYPE at one order below, as shown in Figures 48(a-d). These new rules result from the combination of the ones presented in Figures 14(a-d) and 29. Once more the instantiations of basic types remain implicit as they can be inferred from the powertype declaration: all subtypes of the base type are, by definition, instances of the powertype.

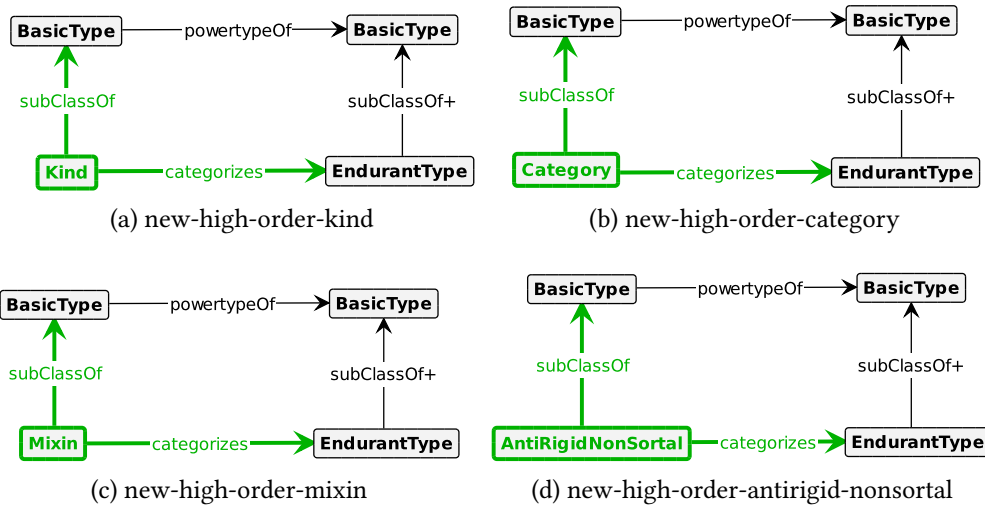


Figure 48 – Transformation rules to introduce a new high-order independent type.

5.1.5 Introducing High-Order Dependent Types

Endurant dependent types are introduced in high levels as a specialization of some SORTAL type from which they inherit a principle of identity, respecting the *rigidity* ontological principle, and categorizing any ENDURANT TYPE at one order below. This is performed by the rules shown in Figures 49(a, b), which are the product of the amalgamation of the original rules presented in Figures 15, 16 and 29.

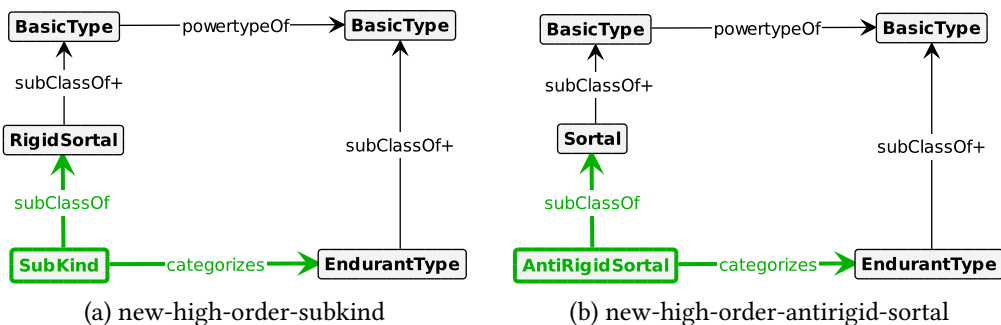


Figure 49 – Transformation rules to introduce a new high-order dependent type.

5.1.6 Introducing Specializations for Existing Non-Sortal Types

Proceeding with the merging of the rules in Sections 3.1.3 and 4.1.4, CATEGORIES and MIXINS can be specialized in any ENDURANT TYPE without instances at the same level, and ANTI-RIGID NON-SORTALS can be specialized in an ANTI-RIGID TYPE without instances at the same level. Two new rules are shown in Figures 50(a, b), resulting from combining the ones from Figures 17(a, b) and 30. Additionally, the new rule depicted in Figure 51 stems from the merging of original rules from Figures 18 and 30.

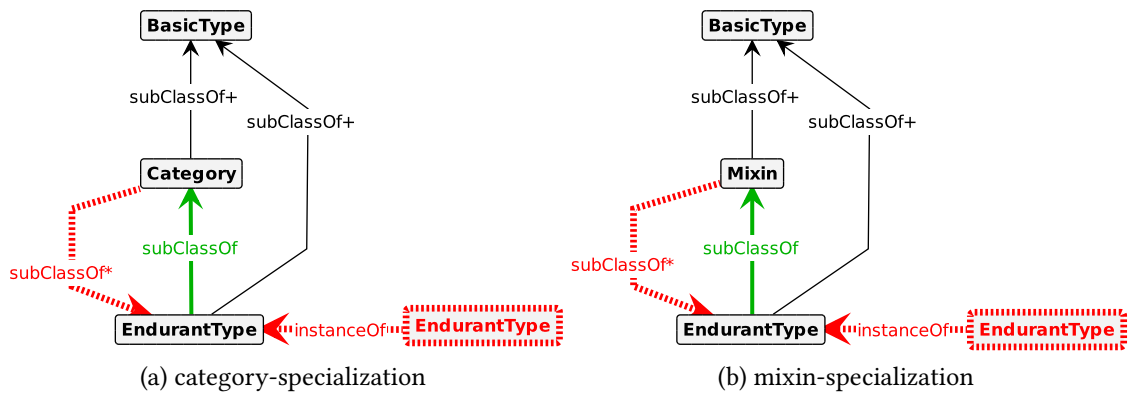


Figure 50 – Transformation rules to specialize a CATEGORY or a MIXIN.

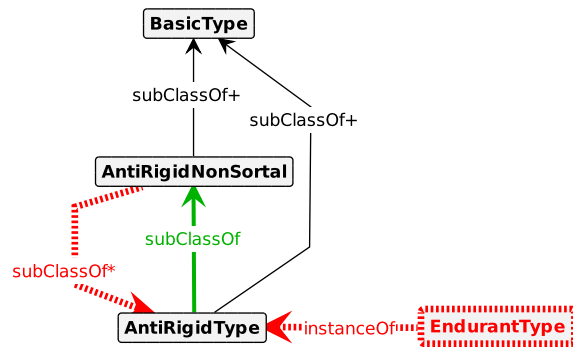


Figure 51 – Transformation rule to specialize an ANTI-RIGID NON-SORTAL.

5.1.7 Introducing Generalizations for Existing Sortal Types

In line with the discussion from Sections 3.1.4 and 4.1.4, SORTAL types without instances can be generalized accordingly to the rules in Figures 52(a, b), which are formed from the original ones in Figures 19(a, b) and 30.

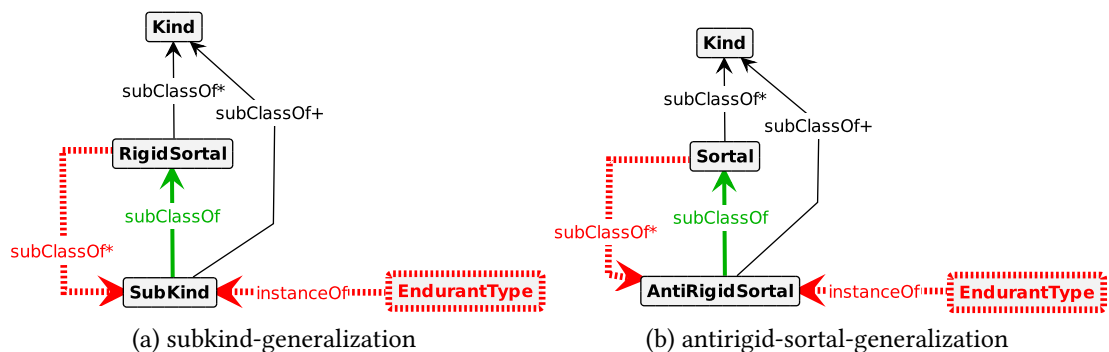


Figure 52 – Transformation rules to generalize a SORTAL.

5.1.8 Classifying Types

Finally, to introduce a new instantiation relation similar to the one described in Section 4.1.5, we adapt the original rule from Figure 31 to consider only the instantiation of

ENDURANT TYPES, preventing the instantiation of two different KINDS by the same type, yielding the new rule shown in Figure 53.




Figure 53 – Transformation rule to classify an ENDURANT TYPE.

5.1.9 Summary of the Grammar Rules

The grammar to build ontologically correct multi-level taxonomies is structured into the following operations:

- The introduction of new first-order types specializing directly the basic type representing the INDIVIDUAL class (introduction of first-order kinds and non-sortals);
- The introduction of new first-order sortals that require an existing first-order sortal to inherit a principle of identity;
- The introduction of a basic type representing a new order;
- The introduction of new high-order types specializing directly some high-order basic type and categorizing some type at one level below (introduction of high-order kinds and non-sortals);
- The introduction of new high-order sortals that require an existing high-order sortal at the same level to inherit a principle of identity, categorizing some type at one level below;
- The specialization of existing non-sortals in other endurant types;
- The generalization of existing sortals in other sortals of the same kind;
- The introduction of instantiation relations between a specialization of a type and its (higher-order) categorizer.

5.2 Formal Verification

We perform the same verification process as previously described for each individual micro-theory grammar, in order to ensure that the soundness results still hold for the combined grammar described in this chapter. Luckily, given that the ontological restrictions from each micro-theory are orthogonal, our new set of well-formedness restrictions can be composed by just taking the graph conditions presented for the grammar of enduring types (Section 3.2) and the grammar of multi-level taxonomies (Section 4.2) in tandem, except for one restriction, which is of a type instantiating simultaneously two different KINDS. This last restriction, that only appears when both micro-theories are considered in tandem, is presented in Figure 54.




Figure 54 – Restrictive condition of a type instantiating multiple KINDS.

We once more run a language enumeration, this time to analyze the combined graph grammar proposed in this chapter. When executing, the GROOVE tool managed to generate a total of 2,389,713 taxonomies up to a bound $N = 7$. The results obtained with this experiment are summarized in Table 11. Once again, we can see that the soundness property holds (at least up to the reachable bound.)

# types (N)	Produced taxonomy shapes	Incorrect taxonomy shapes
2	5	0
3	26	0
4	202	0
5	2,514	0
6	55,592	0
7	2,331,374	0

Table 11 – Results of soundness analysis for the combined grammar.

As discussed in Section 4.2.5 for the MLT-based grammar, a completeness analysis beyond $N = 5$ with the proposed strategy is not feasible, and further work is required to support a completeness claim for the joint grammar, possibly involving partial exploration of the state space or proof assistants.

6 Final Considerations

We developed a formally specified, ontologically well-founded, metamodel-independent and sound graph grammar for the elaboration of multi-level taxonomies. This expands on the current state-of-the-art modeling methods, by proposing a novel technique that leads to the development of (multi-level) taxonomies that are correct by construction. We have accomplished that by leveraging on a *typology of enduring types* and a *theory of high-order types*. The typology of enduring types is part of the Unified Foundational Ontology (UFO) (GUIZZARDI et al., 2015), which underlies the Ontology-Driven Conceptual Modeling language OntoUML (GUIZZARDI et al., 2018). The theory of high-order types was conceived as a foundation for multi-level models (ALMEIDA et al., 2018; CARVALHO; ALMEIDA, 2018) and was also applied successfully in a number of initiatives, including the definition of a well-founded multi-level modeling language (FONSECA et al., 2021).

The original theory of UFO puts forth a number of ontological distinctions based on formal meta-properties. As a result of the logical characterization of these meta-properties, we have that certain structures (patterns) are imposed on the language primitives representing these distinctions (RUY et al., 2017). We have identified a set of primitive operations on taxonomic structures that incorporates the ontological distinctions of UFO and multi-level modeling concepts from MLT, and that guarantees the soundness of the generated taxonomies. This forms the basis for the systematic design of such structures at a higher level of abstraction.

Given the limitations of metamodels as a mechanism for representing the abstract syntax of a language, these structures were not treated as first-class citizens before and have remained hidden in the abstract syntax of the original OntoUML proposal (GUIZZARDI, 2005). This work addresses this exact problem. By leveraging on that theory, we propose a *pattern grammar* (formalized as a graph grammar) that embeds these distinctions and multi-level modeling concepts, ensuring *by design* the construction of taxonomic structures that abide by the formal constraints governing their relations. The work proposed here is inspired by the work in (ZAMBON; GUIZZARDI, 2017) and advances the work initiated in (BATISTA et al., 2021). For example, by employing the state exploration mechanism supported by GROOVE, we managed to detect important omissions in the rule set proposed in that first work.

As we have shown, the ratio between incorrect and correct taxonomies is enormous, and worse still, grows exponentially as the number of elements in the model grows. This provides support for our claim that the graph grammar patterns identified here are indeed ‘higher-level’ constructs when compared to the direct creation of modeling elements and relations in conventional (free-form) modeling environments.

Another important aspect is that our proposal captures the representation consequences

of ontology theories in a way that is metamodel-independent. For this reason, these results can be carried out to other languages and platforms. In particular, we are currently developing a plugin for Protégé that, among other things, implements the primitive operations proposed in this dissertation. This plugin is intended to be used in tandem with the gUFO ontology (a lightweight implementation of UFO) (ALMEIDA et al., 2019). In that implementation, these operations take the form of ontology patterns to be applied, to support its users in modeling consistent Semantic Web ontologies.

All the graph grammars proposed in this dissertation for the construction of correct taxonomies, in addition to our completely permissive grammars for enduring types and the MLT, can be obtained at <<https://github.com/nemo-ufes/ufo-mlt-taxonomy-graph-grammar>>. All of our permissive graph grammars, including those used to partially verify the completeness of our proposed grammars, are presented in the Appendix.

6.1 Limitations

As a first limitation of our work, we can point that the grammar language generated by all of our graph grammars are *infinite*, i.e., these grammars potentially generate an infinite number of graph state models. To cope with this situation, using model checking as in this work, we need to limit the exploration of the state space, what is done here by the number of types present in each taxonomy shape. This way, we cannot conclusively demonstrate any *soundness* or *completeness* claim for the graph grammars proposed here, as maybe it could be done using proof assistants, for example. Given the inherently exponential growth of the number of possible taxonomy shapes with respect to the number of types, a great amount of additional memory and processing time is necessary to a little increase in the number of types used as limit. This *state space explosion* is a common problem for all explicit state model checkers, such as GROOVE (GHAMARIAN et al., 2012).

Nonetheless, to support our *soundness* claims for the graph grammars we propose, we rely on the *small scope hypothesis*, which basically claims most design errors can be found in small counterexamples (GAMMAITONI; KELSEN; MA, 2018), and we present an analysis of the nature of well-formedness violations according to our restrictive conditions, showing that these violations could appear in taxonomy shapes with a number of types less or equal the limits established in this work. With respect to *completeness* claims, we developed additional permissive graph grammars, each of which generates the correct models and a set of incorrect models disjoint from the others, providing additional evidence that our graph grammars are not only *sound*, but also *complete*.

6.2 Future Work

In our future work, we intend to address the aforementioned scalability limitations by exploring other techniques, such as different ways to partially explore the grammars state space or the use of proof assistants to analyse properties of the grammars in order to define and prove theorems about them.

In this work, we did some choices about what to include in the scope from the MLT and the typology of endurant types of UFO. We have focused on the essential aspects of taxonomic structures, leaving out some concepts for the sake of grammar simplicity and formal verification tractability. For instance, we did not include different forms of *powertype* and *categorization*, *subordination*, *generalization set declarations* and its constraints (*disjointness* and *completeness*), *relationships* and *existential dependence*. In future works, these and other concepts from the two micro-theories could be taken into account.

Bibliography

- ALMEIDA, J. P. A. et al. Multi-Level Conceptual Modeling: Theory and Applications. In: *Proceedings of the XI Seminar on Ontology Research in Brazil and II Doctoral and Masters Consortium on Ontologies (Ontobras 2018)*. [S.l.]: CEUR, 2018. p. 26–41. Cited 5 times on pages [14](#), [15](#), [26](#), [28](#), and [60](#).
- ALMEIDA, J. P. A.; FRANK, U.; KÜHNE, T. Multi-Level Modelling (Dagstuhl Seminar 17492). *Dagstuhl Reports*, v. 7, p. 18–49, 2018. ISSN 21925283. Cited in page [26](#).
- ALMEIDA, J. P. A. et al. *gUFO: a lightweight implementation of the Unified Foundational Ontology (UFO)*. 2019. Cited in page [61](#).
- ANJORIN, A.; LEBLEBICI, E.; SCHÜRR, A. 20 years of triple graph grammars: A roadmap for future research. *Electron. Commun. Eur. Assoc. Softw. Sci. Technol.*, v. 73, 2015. Cited in page [29](#).
- ARENDDT, T. et al. Henshin: advanced concepts and tools for in-place EMF model transformations. In: *MODELS'10*. [S.l.: s.n.], 2010. p. 121–135. Cited in page [30](#).
- ATKINSON, C.; KÜHNE, T. Meta-level independent modelling. In: *International Workshop on Model Engineering at 14th European Conference on Object-Oriented Programming*. [S.l.: s.n.], 2000. v. 12, p. 16. Cited in page [27](#).
- ATKINSON, C.; KÜHNE, T. The essence of multilevel metamodeling. In: SPRINGER. *International Conference on the Unified Modeling Language*. [S.l.], 2001. p. 19–33. Cited in page [26](#).
- BATISTA, J. O. et al. Building correct taxonomies with a well-founded graph grammar. In: *15th International Conference on Research Challenges in Information Science (RCIS 2021)*. [S.l.]: Springer, 2021. (Lecture Notes in Business Information Processing, v. 415), p. 506–522. Cited 4 times on pages [15](#), [30](#), [31](#), and [60](#).
- BRASILEIRO, F. et al. Applying a Multi-Level Modeling Theory to Assess Taxonomic Hierarchies in Wikidata. In: *Proceedings of the 25th International Conference Companion on World Wide Web*. [S.l.]: International World Wide Web Conferences Steering Committee, 2016. p. 975–980. Cited 3 times on pages [14](#), [26](#), and [29](#).
- BRASILEIRO, F. et al. Expressive Multi-level Modeling for the Semantic Web. In: *15th International Semantic Web Conference, Kobe, Japan, October 17–21, 2016, Proceedings, Part I*. [S.l.]: Springer, 2016. p. 53–69. Cited in page [26](#).
- CARDELLI, L. Structural subtyping and the notion of power type. In: *Proceedings of the 15th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. [S.l.: s.n.], 1988. p. 70–79. Cited in page [27](#).
- CARVALHO, V. et al. Multi-level ontology-based conceptual modeling. *Data & Knowledge Engineering*, v. 109, p. 3–24, 2017. ISSN 0169023X. Cited in page [26](#).
- CARVALHO, V.; ALMEIDA, J. P. A.; GUIZZARDI, G. Using a Well-Founded Multi-Level Theory to Support the Analysis and Representation of the Powertype Pattern in Conceptual Modeling.

In: *28th International Conference on Advanced Information Systems Engineering (CAiSE 2016)*. [S.l.]: Springer, 2016. p. 309–324. Cited in page 26.

CARVALHO, V. A.; ALMEIDA, J. P. A. A Semantic Foundation for Organizational Structures: A Multi-level Approach. In: *2015 IEEE 19th International Enterprise Distributed Object Computing Conference*. [S.l.]: IEEE, 2015. p. 50–10. Cited in page 25.

CARVALHO, V. A.; ALMEIDA, J. P. A. Toward a well-founded theory for multi-level conceptual modeling. *Software and Systems Modeling*, Springer, v. 17, p. 205–231, 2018. ISSN 16191366. Cited 6 times on pages 14, 15, 26, 27, 28, and 60.

DADALTO, A. A. et al. Type or Individual? Evidence of Large-Scale Conceptual Disarray in Wikidata. In: *40th International Conference on Conceptual Modeling (ER 2021)*. [S.l.]: Springer, 2021. (Lecture Notes in Computer Science, v. 13011), p. 1–10. Cited 4 times on pages 7, 14, 26, and 29.

FONSECA, C. et al. Multi-level conceptual modeling: Theory, language and application. *Data & Knowledge Engineering*, Elsevier BV, v. 134, jul. 2021. ISSN 0169023X. Disponível em: <<https://doi.org/10.1016/j.datak.2021.101894>>. Cited 4 times on pages 14, 26, 29, and 60.

GAMMAITONI, L.; KELSEN, P.; MA, Q. Agile validation of model transformations using compound f-alloy specifications. *Science of Computer Programming*, Elsevier, v. 162, p. 55–75, 2018. Cited 2 times on pages 38 and 61.

GHAMARIAN, A. H. et al. Modelling and analysis using GROOVE. *International journal on software tools for technology transfer*, Springer, v. 14, n. 1, p. 15–40, 2012. Cited 9 times on pages 15, 17, 18, 20, 29, 32, 38, 49, and 61.

GONZALEZ-PEREZ, C.; HENDERSON-SELLERS, B. A powertype-based metamodelling framework. *Software & Systems Modeling*, v. 5, n. 1, p. 72–90, Apr 2006. ISSN 1619-1374. Disponível em: <<https://doi.org/10.1007/s10270-005-0099-9>>. Cited 2 times on pages 25 and 26.

GUARINO, N.; WELTY, C. A. An overview of OntoClean. In: *Handbook on ontologies*. [S.l.]: Springer, 2004. p. 151–171. Cited in page 13.

GUIZZARDI, G. *Ontological foundations for structural conceptual models*. [S.l.]: University of Twente, 2005. (Telematica Institute Fundamental Research Series, 15). Cited 6 times on pages 13, 20, 21, 25, 31, and 60.

GUIZZARDI, G. et al. Towards an Ontological Analysis of Powertypes. In: *Proceedings of the Joint Ontology Workshops 2015 Episode 1: The Argentine Winter of Ontology*. CEUR-WS.org, 2015. (CEUR Workshop Proceedings, v. 1517). Disponível em: <http://ceur-ws.org/Vol-1517/JOWO-15_FOfAI_paper_7.pdf>. Cited in page 53.

GUIZZARDI, G. et al. Endurant types in ontology-driven conceptual modeling: Towards OntoUML 2.0. In: SPRINGER. *International Conference on Conceptual Modeling*. [S.l.], 2018. p. 136–150. Cited 6 times on pages 13, 21, 22, 24, 31, and 60.

GUIZZARDI, G. et al. Towards ontological foundations for conceptual modeling: The unified foundational ontology (UFO) story. *Applied ontology*, IOS Press, v. 10, n. 3-4, p. 259–271, 2015. Cited 2 times on pages 20 and 60.

- GUIZZARDI, G. et al. An ontologically well-founded profile for UML conceptual models. In: SPRINGER. *International Conference on Advanced Information Systems Engineering*. [S.l.], 2004. p. 112–126. Cited in page 13.
- HALPIN, T.; MORGAN, T. *Information modeling and relational databases*. [S.l.]: Morgan Kaufmann, 2010. Cited in page 14.
- HECKEL, R. Graph transformation in a nutshell. *Electronic notes in theoretical computer science*, Elsevier, v. 148, n. 1, p. 187–198, 2006. Cited 2 times on pages 17 and 18.
- HECKEL, R.; TAENTZER, G. *Graph Transformation for Software Engineers: With Applications to Model-Based Development and Domain-Specific Language Engineering*. [S.l.]: Springer, 2020. 309 p. ISBN 978-3-030-43915-6. Cited in page 29.
- HENDERSON-SELLERS, B. *On the Mathematics of Modelling, Metamodelling, Ontologies and Modelling Languages*. [S.l.]: Springer Publishing Company, Incorporated, 2012. Cited in page 27.
- HIRSCH, E. *The concept of identity*. [S.l.]: Oxford University Press, 1992. Cited in page 23.
- JACKSON, D. Alloy: a language and tool for exploring software designs. *Communications of the ACM*, ACM New York, NY, USA, v. 62, n. 9, p. 66–76, 2019. Cited in page 38.
- JAKUMEIT, E.; BUCHWALD, S.; KROLL, M. GrGen.NET. *International Journal on Software Tools for Technology Transfer*, v. 12, p. 263–271, 2010. Cited in page 30.
- KANEIWA, K. Order-sorted logic programming with predicate hierarchy. *Artificial Intelligence*, Elsevier, v. 158, n. 2, p. 155–188, 2004. Cited in page 31.
- KANEIWA, K. Existential rigidity and many modalities in order-sorted logic. *Knowledge-Based Systems*, Elsevier, v. 24, n. 5, p. 629–641, 2011. Cited in page 31.
- KANEIWA, K.; MIZOGUCHI, R. An order-sorted quantified modal logic for meta-ontology. In: SPRINGER. *International Conference on Automated Reasoning with Analytic Tableaux and Related Methods*. [S.l.], 2005. p. 169–184. Cited in page 31.
- KANEIWA, K.; NGUYEN, P. H. Decidable order-sorted logic programming for ontologies and rules with argument restructuring. In: SPRINGER. *International Semantic Web Conference*. [S.l.], 2009. p. 328–343. Cited in page 31.
- MOL, M. de; RENSINK, A.; HUNT, J. J. Graph transforming Java data. In: *FASE*. [S.l.: s.n.], 2012. p. 209–223. Cited in page 30.
- NEUMAYR, B.; GRÜN, K.; SCHREFL, M. Multi-level domain modeling with m-objects and m-relationships. In: AUSTRALIAN COMPUTER SOCIETY, INC. *Proceedings of the Sixth Asia-Pacific Conference on Conceptual Modeling-Volume 96*. [S.l.], 2009. p. 107–116. Cited 2 times on pages 25 and 26.
- NICKEL, U.; NIERE, J.; ZÜNDORF, A. The FUJABA environment. *Proceedings of the 2000 International Conference on Software Engineering. ICSE 2000 the New Millennium*, p. 742–745, 2000. Cited in page 30.
- ODELL, J. Power types. *J. Object Oriented Program.*, v. 7, p. 8–12, 1994. Cited in page 27.

- OLTRAMARI, A. et al. Restructuring WordNet's top-level: The OntoClean approach. *LREC2002, Las Palmas, Spain*, v. 49, 2002. Cited in page 13.
- RENSINK, A. Representing first-order logic using graphs. In: *ICGT*. [S.l.: s.n.], 2004. p. 319–335. Cited in page 30.
- RENSINK, A.; ZAMBON, E. A type graph model for Java programs. In: *FMOODS/FORTE*. [S.l.: s.n.], 2009. p. 237–242. Cited in page 30.
- ROBERSON, M. et al. Efficient software model checking of soundness of type systems. *ACM Sigplan Notices*, ACM New York, NY, USA, v. 43, n. 10, p. 493–504, 2008. Cited in page 38.
- RUY, F. B. et al. From reference ontologies to ontology patterns and back. *Data & Knowledge Engineering*, Elsevier, v. 109, p. 41–69, 2017. Cited 3 times on pages 13, 30, and 60.
- SEYED, A. P. Integrating ontoclean's notion of unity and identity with a theory of classes and types. In: IOS PRESS. *Formal Ontology in Information Systems: Proceedings of the Seventh International Conference (FOIS 2012)*. [S.l.], 2012. v. 239, p. 205. Cited in page 31.
- SEYED, A. P.; SHAPIRO, S. C. A method for evaluating ontologies. Citeseer, 2011. Cited in page 31.
- VARRÓ, D.; BALOGH, A. The model transformation language of the VIATRA2 framework. *Sci. Comput. Program.*, v. 68, p. 214–234, 2007. Cited in page 30.
- WELTY, C.; ANDERSEN, W. Towards ontoclean 2.0: A framework for rigidity. *Applied Ontology*, IOS Press, v. 1, n. 1, p. 107–116, 2005. Cited in page 31.
- ZAMBON, E. *Abstract Graph Transformation – Theory and Practice*. [S.l.]: University of Twente, 2013. (Centre for Telematics and Information Technology). Cited 2 times on pages 17 and 30.
- ZAMBON, E.; GUIZZARDI, G. Formal definition of a general ontology pattern language using a graph grammar. In: IEEE. *2017 Federated Conference on Computer Science and Information Systems (FedCSIS)*. [S.l.], 2017. p. 1–10. Cited 3 times on pages 15, 30, and 60.
- ZAMBON, E.; RENSINK, A. Recipes for coffee: compositional construction of Java control flow graphs in GROOVE. In: *Principled Software Development*. [S.l.: s.n.], 2018. p. 305–323. Cited in page 30.

Appendices

APPENDIX A – Completely Permissive Endurant Types Graph Grammar




Figure 55 – Transformation rules to introduce an enduring type.




Figure 56 – Transformation rule to introduce a specialization relation.

APPENDIX B – Rigidity Permissive Endurant Types Graph Grammar




Figure 57 – Transformation rules to introduce an independent type.



Figure 58 – Transformation rule to introduce a SUBKIND type.




Figure 59 – Transformation rule to introduce an ANTI-RIGID SORTAL type.

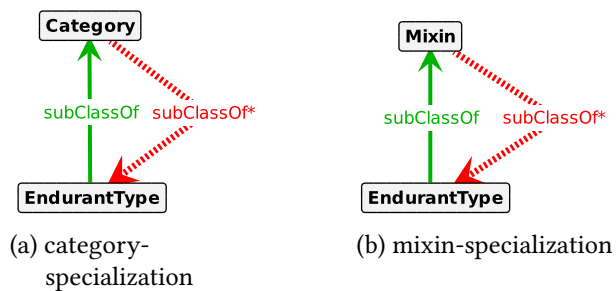


Figure 60 – Transformation rules to specialize a CATEGORY or a MIXIN.




Figure 61 – Transformation rule to specialize an ANTI-RIGID NON-SORTAL type.

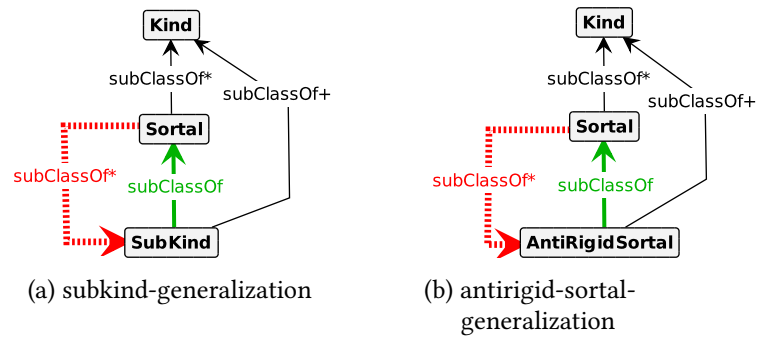


Figure 62 – Transformation rules to generalize a SUBKIND or an ANTI-RIGID SORTAL.

APPENDIX C – Sortality Permissive Endurant Types Graph Grammar




Figure 63 – Transformation rules to introduce an independent type.




Figure 64 – Transformation rule to introduce a SUBKIND type.




Figure 65 – Transformation rule to introduce an ANTI-RIGID SORTAL type.

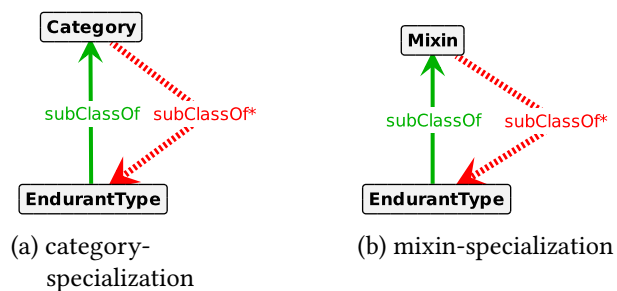


Figure 66 – Transformation rules to specialize a CATEGORY or a MIXIN.




Figure 67 – Transformation rule to specialize an ANTI-RIGID NON-SORTAL type.

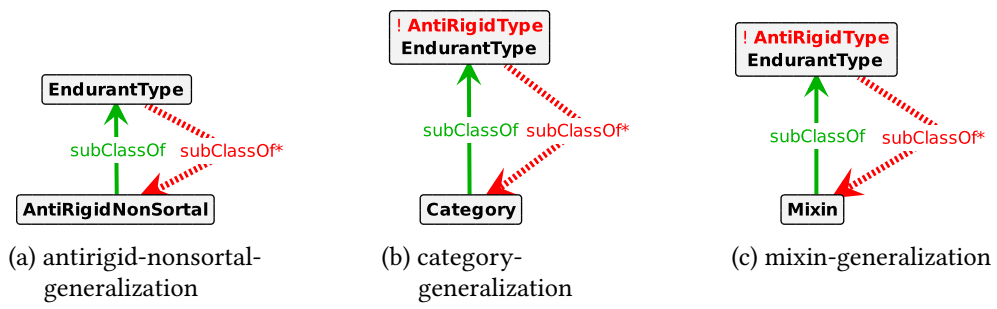


Figure 68 – Transformation rules to generalize NON-SORTALS.

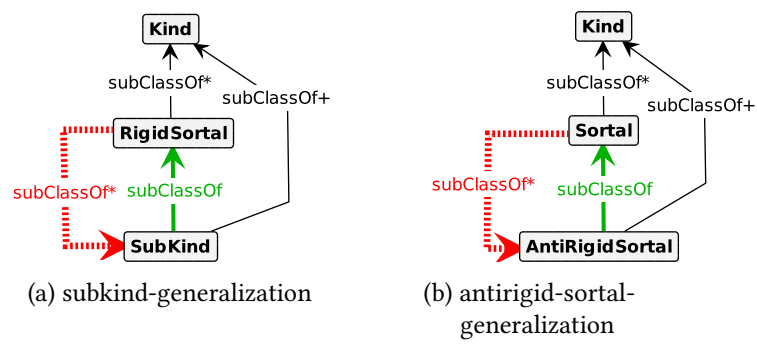


Figure 69 – Transformation rules to generalize a SUBKIND or an ANTI-RIGID SORTAL.

APPENDIX D – Permissive Endurant Types Graph Grammar Regarding Kind Constraints




Figure 70 – Transformation rules to introduce an independent type.

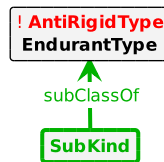


Figure 71 – Transformation rule to introduce a SUBKIND type.




Figure 72 – Transformation rule to introduce an ANTI-RIGID SORTAL type.

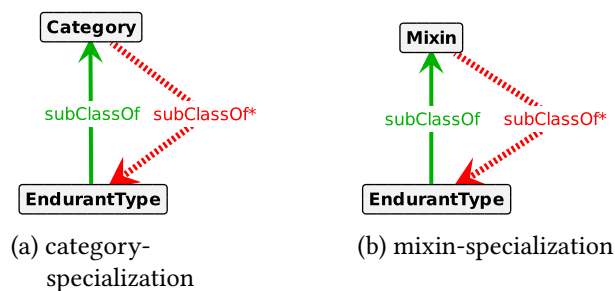


Figure 73 – Transformation rules to specialize a CATEGORY or a MIXIN.




Figure 74 – Transformation rule to specialize an ANTI-RIGID NON-SORTAL type.




Figure 75 – Transformation rules to generalize a RIGID or an ANTI-RIGID SORTAL.

APPENDIX E – Completely Permissive Multi-Level Taxonomies Graph Grammar




Figure 76 – Transformation rule to introduce a non-basic type.




Figure 77 – Transformation rule to introduce a new order.



Figure 78 – Transformation rule to introduce a new categorizer.




Figure 79 – Transformation rule to introduce a new specialization relation.




Figure 80 – Transformation rule to introduce a new instantiation relation.

APPENDIX F – Multi-Level Taxonomies

Graph Grammar with Categorization

Breaking Stratification



Figure 81 – Transformation rule to introduce a first-order type.




Figure 82 – Transformation rule to introduce a new order.

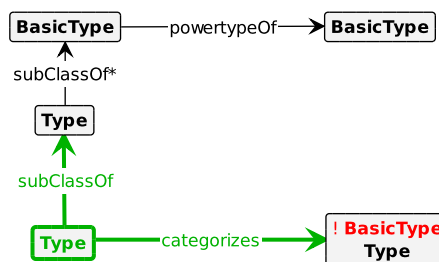


Figure 83 – Transformation rule to introduce a new categorizer.

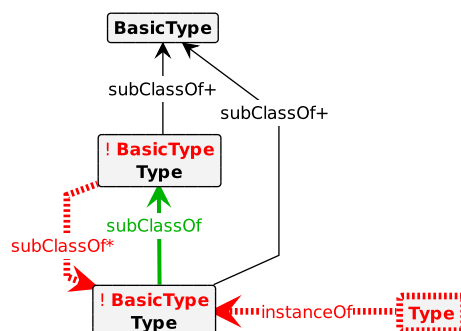


Figure 84 – Transformation rule to introduce a new specialization relation.

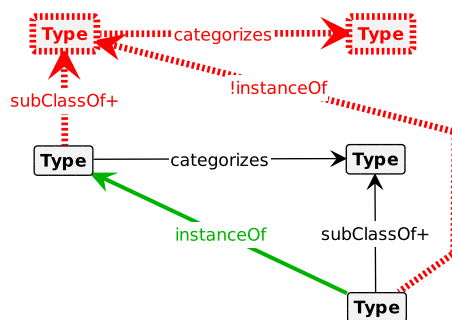


Figure 85 – Transformation rule to introduce a new instantiation relation.

APPENDIX G – Multi-Level Taxonomies

Graph Grammar with Specialization

Breaking Stratification



Figure 86 – Transformation rule to introduce a first-order type.




Figure 87 – Transformation rule to introduce a new order.

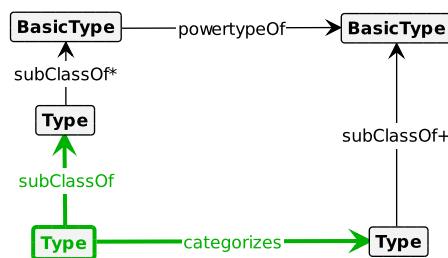


Figure 88 – Transformation rule to introduce a new high-order type.




Figure 89 – Transformation rule to introduce a new specialization relation.

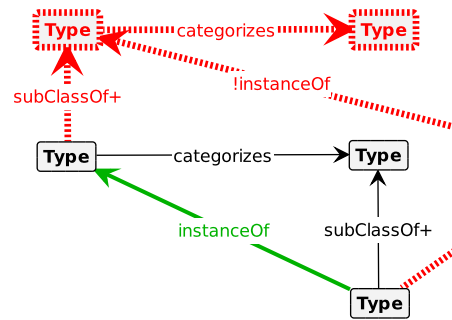


Figure 90 – Transformation rule to introduce a new instantiation relation.

APPENDIX H – Permissive Multi-Level Taxonomies Graph Grammar Regarding Instantiation Constraints



Figure 91 – Transformation rule to introduce a first-order type.




Figure 92 – Transformation rule to introduce a new order.




Figure 93 – Transformation rule to introduce a new high-order type.




Figure 94 – Transformation rule to introduce a new specialization relation.

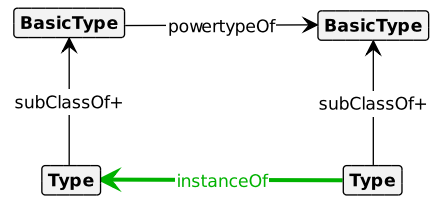


Figure 95 – Transformation rule to introduce a new instantiation relation.